Retrospective Search: Exploration and Ambition on Uncharted Terrain*

Can Urgun[†] Leeat Yariv[‡]

November 25, 2022

Abstract. We provide a framework for analyzing search across correlated observations. The agent—an online customer, a drug company, a politician—tracks innovations over a Brownian path. The agent chooses the speed and length of search and retrospectively picks the best innovation when search is completed. We show that the optimal search speed is U-shaped: it is most intense when approaching a breakthrough or when nearing search termination. Unlike search across independent samples, search optimally stops when observations are sufficiently *discouraging*, following a drawdown stopping boundary. We also show the tractability and features of optimal contracts with retrospective searchers.

Keywords: Retrospective Search, Correlated Samples, Optimal Stopping, Search Speed

JEL codes: C73, D81, D83, O35

*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, Penn State University, 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.

†Princeton University; curgun@princeton.edu

‡Princeton University, CEPR, and NBER; lyariv@princeton.edu

1 Introduction

Many economic decisions are preceded by a discovery process yielding the alternative deemed best. When excavation teams search for minerals, they seek the most promising tract on their path; When pharmaceutical companies introduce a new drug, they seek the most potent dosage; When consumers shop online, they seek the best product for their needs, or the lowest price for their desired item. Since search often entails a cost, such discovery processes entail two fundamental questions. First, when should search stop and what outcomes would it deliver? Second, how rapid should search be depending on where discoveries stand? The canonical models used to answer these key questions—starting from the landmark studies of Stigler (1961), McCall (1970), and Mortensen (1970)—assume independent search outcomes. Independence simplifies dramatically the analysis of search: in its basic form, the distribution of possible discoveries is assumed fixed. Even when search is carried out sequentially, 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. Geological survey teams face correlation across observations driven by the similarity of adjacent plots. The adjustment of drug dosages in phase I/II clinical trials commonly follows a pre-determined sequence, where proximal dosages are understood to produce similar results. A governmental policymaker may revise a program incrementally until the right feature combination meets the challenge. A consumer looking for the best product—an item on Amazon.com, a Neflix movie, or an Airbnb rental unit—often begins search with a simple query, then considers related suggestions that pop up, until making a choice. Without the technical convenience afforded by independence, models allowing for correlation, as in these examples, assume short-lived or myopic agents and a fixed search speed.[1] We provide a framework for studying search by long-lived agents when observations are correlated over time, and when search speed can be adjusted dynamically. We fully characterize the optimal search policy and analyze how it affects agency relationships.

We consider an agent searching over a pre-prescribed path in continuous time. Indeed, the search for minerals most commonly occurs over a one-dimensional path, often termed a *vein* (U.S. Department of Agriculture, 1995); the adjustment of drug dosages in phase I/II clinical trials commonly follows a pre-determined sequence of proximal dosages (Thall and Cook, 2004); and an online shopper may follow a suggested path of related items, or an ordered sequence of retailers (De los Santos, Hortaçsu, and Wildenbeest, 2012). In our model, Brownian motion governs the value of each ordered discovery, as in Callan-

[1] For instance, see Callander (2011), Garfagnini and Strulovici (2016), and our literature review below.

der (2011). That is, the realized Brownian path maps each ordered discovery to its value. This structure naturally introduces intertemporal correlation between discoveries: one moment's observation forms the expectation for any future moment. Recent papers provide empirical evidence for Gaussian correlation in the case of oil exploration (Hodgson, 2021) and of consumer product search (Hodgson and Lewis, 2021). Jovanovic and Rob (1990) develops an axiomatic foundations for the Brownian search evolution in the innovation context.[2]

While searching, the agent can continuously adjust her search speed, the distance traversed on the realized path per unit of time. Search is costly and depends on the speed at which the agent traverses the path, with higher speeds associated with greater costs. For example, a mineral excavation team may need to hire personnel in order to cover ground faster without sacrificing samples' precision and drug companies can adjust costly sample sizes over time to learn the efficacy of different dosages faster or slower. In our model, whenever the agent stops searching, she can recall all her observations retrospectively and pick the highest, which constitutes her search outcome. The agent's payoff—her objective—is the discounted maximal observation from her search net of the cumulative discounted search costs spent.[3] We call such a search process *retrospective search*. The special case of retrospective search in which search speed is fixed exogenously is a direct analog of the classical Stigler (1961) setting, but with correlated samples.

We start by characterizing the agent's optimal policy. We show that the agent optimally chooses a search speed that depends on the distance between the current discovery and her best discovery yet. The agent's optimal search speed is U-shaped. The optimal search speed is highest when approaching a transition point: either accomplishing a new peak discovery—a *breakthrough*—or when nearing search termination—a *breakdown*. The optimal search speed depends on search costs, with more log-convex costs associated with slower search speeds. More patient agents are less responsive to the relative quality of the current discovery when choosing speed. In fact, absent discounting, a constant search speed is optimal.

A large literature in operation management focuses on product development speed and its link to outcomes (see the meta-studies of Chen, Damanpour, and Reilly, 2010 and Cankurtaran, Langerak, and Griffin, 2013). Our analysis suggests that whenever discounting is non-trivial, optimal search speeds are adjusted as new discoveries are made, increasing when search nears a breakthrough or a breakdown. Consequently, naively considering

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

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

correlations between outcomes and a snapshot of development speeds may yield misleading results.

The observation that, with discounting, the optimal speed increases as the agent approaches a turning point—a breakthrough or a breakdown—is reminiscent of the vast literature on the so-called *goal-gradient hypothesis*. First introduced by Hull (1932), and more recently revived by Kivetz, Urminsky, and Zheng (2006), the goal-gradient hypothesis suggests a tendency for individuals to increase effort as a goal approaches. 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.

Certainly, many search units consider long activity horizons and arguably exhibit less pronounced discounting. In such settings, our results show that the optimal search speed is constant and need not respond to the ebb and flow of instantaneous discoveries. Evidence from a variety of environments in which discounting may not be pronounced is in line with this insight. For example, while mineral exploration teams can choose the speed of mine digging on their path, standard practices since the 1950s dictate constant and pre-prescribed speed and depth of sampled wells, independent of prior observations, see Zhilkin (1961). Similarly, in clinical trials for new drugs, research suggests the use of a constant sample size for each consecutive dosage, see Brock, Billingham, Copland, Siddique, Sirovica, and Yap (2017).

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 cutoff for stopping evolves over time: as the agent accumulates greater discoveries, she becomes more demanding and stops more willingly.

The optimal stopping policy of a retrospective searcher stands in stark contrast to optimal search policies when observations are independent. With independence, a searching agent optimally satisfices, stopping when an observation's value is *sufficiently high*. Intuitively, an observation today foretells nothing about the future and continuation values are therefore constant over time. Since search is costly, when observing a sufficiently high value, the agent should stop. With correlation, the agent optimally stops when an observation's value is *sufficiently low*: such an observation foretells poor observations in the immediate future and indicates a long duration till additional breakthroughs are made.

Our characterization of the optimal stopping boundary is consistent with many search strategies already utilized in the field. For instance, it suggests that when prospecting for, say, oil, hitting wells on the path that are worse than the most promising wells already identified should limit additional excavation. Similarly, when looking for the optimal drug dosage in clinical trials, evidence of severe side effects should limit exploration of larger dosages—see Thall and Cook (2004) for a detailed algorithm specifying a stopping policy in line with the one we suggest. Our characterization also offers caution for empirical studies estimating search behaviors *assuming* independence of observations, see the influential study of De los Santos et al. (2012). In particular, agents halting their search with poor observations, or utilizing recall, need not indicate naïve search heuristics, but can instead be explained through an underlying correlation in agents' search path.

In the last part of the paper, we embed our retrospective-search model within a principal-agent interaction. Many search processes, particularly in the realm of research and development, occur within organizations. Financiers contract with mining firms, CEOs manage R&D teams, voters manage politicians, etc. 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, joint venture contracts between investors and mining companies often specify fixed or flow transfers, in addition to commissions on findings, see Root (1979). Similarly, university licensing agreements commonly include both fixed fees and royalties, see Jensen and Thursby (2001). When search observations are independent, the optimal contract is generally not amenable to analysis—see our Supplementary Materials for details. Optimal contracts for retrospective searchers, however, are highly tractable.

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 speed. 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 exponential search costs, we show that the agent searches with lower speed when under a commission contract relative to when searching alone. In fact, we show that contractual frictions come at a substantial cost of nearly a quarter, or 24%, of the surplus.

From a technical viewpoint, retrospective search combines a stopping problem with a control problem. Such problems have received limited attention in both the social sciences and the mathematics literature, particularly in the presence of discounting, see our literature review below. The paper offers a methodological contribution by illustrating techniques for solving these types of problems, which we hope will be useful for a variety of other contexts.

2 Related Literature

As already noted, the basic model of search with independent observations was introduced by Stigler (1961), McCall (1970), and Mortensen (1970). The model has been extended in many ways, see, for instance, Olszewski and Weber (2015) and references therein. It 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). There are two important differences between our setting and the canonical search model. First, we consider samples that are correlated over time. Second, we allow the agent to choose her search speed at any point.

Modeling correlation over time using a Brownian path in a search setting is inspired by Callander (2011).[4] However, in Callander's setting, agents are short-lived, and can only 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 far 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. Urgun and Yariv (2021) consider a similar setting to this paper's in which a non-discounting agent is constrained to search within a fixed amount of time. Cetemen, Urgun, and Yariv (2022) consider search over a Brownian path by evolving teams of infinitely-patient agents. Wong (2022) 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, it is shown that exploration hastens after poor outcomes and that search stops when observed outcomes are sufficiently poor.

As already discussed, recent papers provide empirical evidence for Gaussian correlation in consumer product search and oil exploration, see Hodgson and Lewis, 2021 and Hodgson (2021). Hodgson and Lewis (2021) present a theoretical framework for estimating correlations between spatially proximal alternatives. They further characterize some features of behavior when agents use a one-period look ahead heuristic.

The idea that speed, or variance, might be a control variable associated with costs appears in other experimentation models. For instance, Moscarini and Smith (2001) consider

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

a sequential sampling setting in which an agent can control the precision of the signals she receives. 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 search speed highlights a different dimension of search efforts when observations are correlated.

From a technical perspective, our results contribute to the mathematics literature on optimal stopping in which the objective is related to the maximum seen so far, see e.g. Peskir (1998) and Obłój (2007). Most of that literature focuses on agents who experience flow costs and no discounting. In addition, that literature studies optimal stopping absent a control. Pedersen (2000) do consider optimal stopping in the presence of discounting, but without a control. Even then, they provide an Ordinary Differential Equation (ODE) for the stopping boundary, without a general solution. Peskir (2005) considers an agent who controls the drift, but assumes an exogenous stopping rule: search stops whenever observations exit a given interval. Furthermore, there is no discounting. He shows a bang-bang solution. The techniques we develop allow for the analysis of such stopping problems with both discounting and the inclusion of a control—in our case, the costly search speed. We hope the methods we introduce open the door for further studies in the area.

Our analysis of contracts relates to the budding literature on contracts for experimentation, which has thus far focused on the (independent) one- or two-armed bandit setting, see Manso (2011), Halac, Kartik, and Liu (2016), and Guo (2016). As we show in the Supplementary Materials, an analogous agency setting to ours in which the agent searches over a path of independent samples is far less tractable.

3 A Model of Retrospective Search

Consider a risk-neutral agent—an online shopper, a drug company, or a mineral prospecting team—who is searching in continuous time. Time is indexed by t and runs through $[0,\infty)$. When the agent stops searching, she gets benefits from the maximal value found throughout her search, net of her accumulated search costs.

The progress of discoveries follows a Wiener process, where the realized sample path describes the link between new findings and their expected value to the agent. We assume there is a natural progression of exploration. For instance, an online shopper may follow a suggested path of related items, or an ordered sequence of retailers (De los Santos et al., 2012); the adjustment of drug dosages in phase I/II clinical trials often follows a pre-determined sequence of proximal dosages (Thall and Cook, 2004); and mineral prospecting most commonly follows a one-dimensional vein (U.S. Department of Agriculture, 1995). The Wiener process has been axiomatically founded in the innovation context by Jovanovic and Rob (1990). Its use in our model allows us to capture the correlation between expected values of similar findings, and the effects of search speed. Recent empirical work supports the assumption of Gaussian correlation that underlies oil exploration and consumer choice, see Hodgson (2021) and Hodgson and Lewis (2021).[5]

Formally, time proxies for the sequence of ordered samples our agent searches through. For any time t, let B_t denote the standard Brownian motion with $B_0 = 0$. The realized sample path captures the expected value of each (ordered) finding.

The agent can influence the speed at which the path of findings is traversed. In our examples, the investment of resources—financial expenditures, volume of patients recruited to participate in a medical trial, human capital hired, etc.—affects how rapidly search is conducted. Specifically, the agent chooses a continuous and measurable mapping σ_t^2, where $\sigma_t \in [\underline{\sigma}, \bar{\sigma}]$. 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 speed σ_t^2. We assume that c is twice continuously differentiable, increasing, and convex. This notation simplifies our presentation. However, throughout, we slightly abuse terminology and refer to σ_t, rather than σ_t^2, as the agent's search speed. If $\underline{\sigma} = \bar{\sigma}$, our setting boils down to one in which the agent has fixed search speed she cannot control and only chooses when to terminate search.[6]

Speeding up search is tantamount to the "scaling" of time. We use the fact that such scaling is equivalent to a change in the standard deviation of the original Wiener process: see, for instance, Section 8.5 in Øksendal (2003). That is, when the agent's speed is σ^2, we can describe the generated values observed at time t—the expected value of the discovery—that we denote by X_t, using the following law of motion:

$$dX_t = \hat{\sigma}_t dB_t,$$

with $X_0 = 0$.[7]

Search rewards correspond to the best discovery the agent has: an online shopper ultimately buys from the cheapest store, a drug company promotes the most efficacious dosage, and a mineral prospecting team pursues the most promising plot on its track. As

[5]Our assumption that search occurs in continuous time helps with tractability. It is consistent with the observation that in many search settings, investigations are inherently incremental. Agents are not aware of the full set of alternatives at the outset. As search progresses, the menu of possibilities expands.

[6]For details on this special case, see our Supplementary Materials.

[7]We assume no drift since, in many of our applications, the mere passage of time does not provide search improvements. Nonetheless, in Section 6.2 we discuss the impacts of drift. One 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.

search unfolds, the best discovery may improve. Its value is given by M_t, where

$$M_t = \max_{0 \le r \le t} X_r.$$

If search ends at time t, the best discovery, the one yielding M_t, gets implemented.

When the search speed is constant over time, $\sigma_t = \sigma^*$ for all t, we can use the Reflection Principle (see, for instance, Rogers and Williams, 2000) to infer that M_t follows the same distribution as $|\sigma^* B_t|$. Therefore, for any t,

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

That is, the speed 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 *substantial 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 speed of search captures a rather different aspect of search efforts. Nonetheless, the impact on the expected returns to search are similar, with higher search speeds yielding higher expected outcomes.

Ultimately, the agent's problem can be written as:

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

The agent's risk neutrality simplifies our analysis. We discuss the impacts of risk aversion when describing our results, and in Section 6.1. We provide the detailed analysis of retrospective search by a risk-averse agent in the Supplementary Materials.

When observations are independent, as in Stigler (1961), McCall (1970), or Mortensen (1970), recall plays no role. The option value of search continuation is fixed over time. Thus, an outcome rejected in the past is never more appealing in the future. The agent optimally *satisfices* and stops search as soon as she observes a sufficiently high outcome, with no regard to prior realizations. With correlation, recall plays an important role. In

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

$$\Pr(M_t \ge a) = \Pr(T_a \le t) = 2\Pr(X_t \ge a) = \Pr(|X_t| \ge a),$$

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

fact, 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 not worthwhile.[9] The use of recall is present in applications. For example, De los Santos et al. (2012) report online customers returning to sales platforms visited earlier in their search.

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 speed 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 speed, the agent's control. That dependence is non-trivial: the optimal stopping boundary would change were search speed very low or very high. For very low search speeds, the agent would want to stop quickly as there is not much to gain from costly search. For very high search speeds, continuation is prohibitively costly and the agent would terminate search rapidly as well. For intermediate levels of search speeds, the agent may benefit from non-trivial search.

Let $\mathbb{T}_{[a,b]}$ denote the first time a standard Brownian motion (with instantaneous variance of 1) escapes an interval $[a,b]$. Naturally, this time depends on the initial observation in $[a,b]$. We denote by $\mathbf{E}\left(e^{-r\mathbb{T}_{[a,b]}}|X\right)$ the corresponding expected discounted time with an initial observation $X \in [a,b]$. We soon describe it explicitly. Proposition 1 fully characterizes the optimal search policy.

[9] Risk 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.

Proposition 1 (Optimal Search Speed). *For any continuous stopping boundary $g(M)$, the optimal search speed $\sigma^r(M,X)$ 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 speed, 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.[10] 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 proposition suggests that the speed of search changes throughout the search path: optimal search speeds are adjusted as new discoveries are made. This observation bears empirical implications. A large literature in operation management focuses on product development speed and its link to outcomes (see the meta-studies of Chen et al., 2010 and Cankurtaran et al., 2013). Our analysis suggests that naively considering correlations between outcomes and a snapshot of development speeds may yield misleading results.

We now describe the precise dependence of the optimal search speed on the status of search. The expected discounted time a standard Brownian takes to escape the interval $[g(M),M]$ starting from observation X can also be described explicitly:[11]

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

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 Speeds). *For any continuous stopping boundary $g(M)$ and any $r>0$, the optimal search speed $\sigma^r(M,X)$ is 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*

[10] The value function can then be used to identify the optimal search speed among those solutions, see the Appendix for details.

[11] The hyperbolic functions sinh and cosh are defined as follows:

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

optimal search speed 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-speed is U-shaped. The agent responds to details of her environment, searching more rapidly when nearing a breakthrough—namely, when a new maximum is likely to be achieved in a short horizon—or when recent results are discouraging and search breakdown 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 large body of experimental literature, going back Hull (1932), provides evidence for the hypothesis. A recent explosion of work followed Kivetz et al. (2006). The hypothesis has been used to explain effort patterns in a rich set of environments, from marketing, to labor settings, to athletics. Much of this literature views the effect as arising from deeply-ingrained psychological instincts.[12] 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 speed 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. Certainly, many search units consider long activity horizons and arguably exhibit less pronounced discounting. For example, while mineral exploration teams can choose the speed of mine digging on their path, standard practices since the 1950s dictate constant and pre-prescribed speed and depth of sampled wells, independent of prior observations, see Zhilkin (1961). Similarly, in clinical trials for new drugs, research suggests the use of a constant sample size for each consecutive dosage, see Brock et al. (2017).

The intuition for Proposition 1 follows several steps. The first has to do with the impacts of small changes in search speed on the attained maximum. Consider panel (a) of Figure 1, depicting a situation where, at time τ, $X_\tau < M_\tau$. A small perturbation to the search speed at τ does not impact the observed maximum M_τ. Thus, the chosen search speed should depend only on the local process features, namely X_τ. What happens at times τ in which $X_\tau = M_\tau$, as depicted in panel (b) of Figure 1? Our continuous-time formulation implies that, within any infinitesimal interval of time, with probability 1, the agent reaches a new value different from the current maximum.[13] Formally, in Lemma A1 of

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

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

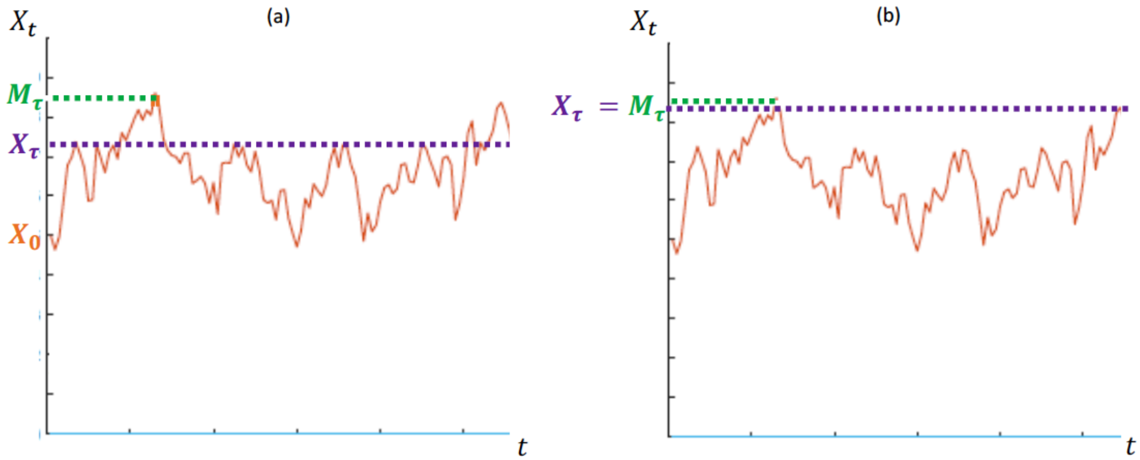

FIGURE 1: Independence of local choices on established maxima

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 speed 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 speed 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 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 speed 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*.

The optimal stopping boundary identified in Proposition 2 is consistent with search strategies used in the field. For instance, it suggests that when prospecting for, say, oil, excavating wells that appear worse than those observed in a region explored earlier should limit additional excavation. Similarly, when seeking the optimal drug dosage in clinical trials, evidence of severe side effects should limit exploration of larger dosages. In fact, Thall and Cook (2004) prescribe an algorithm with a stopping policy resembling the one we suggest.

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 speed 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:

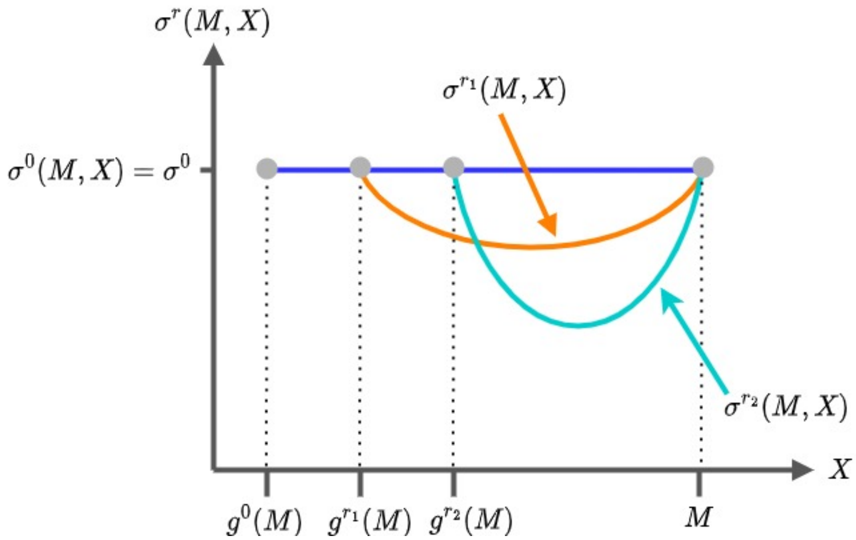

Figure 2: Features of the optimal search policy with $r_2 > r_1 > 0$

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

Figure 2 summarizes our characterization by displaying the shape of the optimal search speed as a function of observed values for different discount rates ($r_2 > r_1 > 0$). 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 speed 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.

Comparison with Independent Search When search is over independent samples, as in the canonical models of Stigler (1961), McCall (1970), and Mortensen (1970), agents cease search when *sufficiently high* values are realized. In contrast, when discoveries are correlated, as in our setting, low realizations indicate that far more research is needed to accomplish a breakthrough. Agents therefore stop when observing *sufficiently low* realized values. Nonetheless, the optimal policy has a similar threshold flavor captured by the drawdown size. In addition, while the characterization of the optimal search speed follows similar lines, its closed-form solution is far less amenable to analysis, a point we return to when discussing agency frictions in search.[14]

Whether or not observations are correlated can affect conclusions from empirical studies of search behavior. For instance, in a classic paper, De los Santos et al. (2012) study customers' search for low book prices online. They assume platforms' book prices are independent. Thus, customers searching sequentially should, in principle, stop their search when observing sufficiently low prices and never recall prior observations. The data are inconsistent with these patterns, leading the authors to conclude that consumers choose their sample size statically. Our analysis suggests that true or perceived correlation across platform prices should lead search to halt when a sufficiently high price is observed and may entail recall, consistent with some of De los Santos et al. (2012)'s evidence.

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,

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

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

[14]Suppose search is conducted over a path governed by independent draws from a standard normal distribution each period. The optimal control, translated into the optimal standard deviation of the normal is given by maximizing $R(\sigma) = \psi^{-1}(c(\sigma))\sigma$, where $\psi(x) = \phi(x) - x(1 - \Phi(x))$, with ϕ and Φ representing the density and cumulative distribution of the standard normal, respectively. See Urgun and Yariv (2021) and our Supplementary Materials for details.

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 $r = 0$. Search duration is inversely proportional to the search speed. 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 speed and the optimal drawdown. Given any drawdown d, the efficient speed is constant and minimizes the cost per speed. Conversely, for any fixed search speed 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 speed and, in turn, is maximized at the optimal search speed. 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 speed is analogous to investment in a reduction of production costs, which impacts quantity.[15]

When $r = 0$, natural comparative statics emerge. Suppose search speed is exogenously fixed at a constant $\hat{\sigma}$ with associated search cost of $\hat{c}$. Keeping the cost fixed, as search speed increases, expected payoffs go up. Keeping the search speed 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.[16] 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

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

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

by σ_i^* and $\mathbf{E}(M_i^*)$ the optimal search speed and expected project value under cost function $c_i(\cdot)$ when $r = 0$. From Proposition 1, $\sigma_1^* > \sigma_2^*$ for any given stopping boundary. Costs affect the stopping boundary as well. The impact on $\mathbf{E}(M_i^*)$ or, equivalently, on the expected search payoff, depends on how the drawdown size changes. If $c_2(\sigma_2^*) > c_1(\sigma_1^*)$, 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 $c_2(\sigma_2^*) < c_1(\sigma_1^*)$, the comparison is generally inconclusive.

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 speed should decline with γ when γ is sufficiently low, namely $\gamma < 2e$. Indeed, the optimal search speed $\sigma^* = \left(\frac{2}{\gamma}\right)^{1/\gamma}$ declines for low values of γ and asymptotes at 1 as γ increases, see panel (a) of Figure 3. 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 3. The expected project value is given by $\mathbf{E}(M_\tau^*) = \frac{1}{2}\left(\frac{2}{e\gamma}\right)^{2/\gamma}$, which is double the expected payoff from search, and is non-monotonic in γ. It is decreasing initially and then increasing indefinitely to an asymptote, as depicted in panel (c) of Figure 3.

In other words, increases in log-convexity as defined above make the agent less ambitious in terms of search speed, 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 speed. 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 speed corresponding to $\bar{c}(\cdot)$ then satisfies:

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

Since the cost function c is convex, $\sigma c'(\sigma)$ is non-decreasing. Hence, the impact of wages on search speed need not be monotonic and depends on the cost function's curvature.

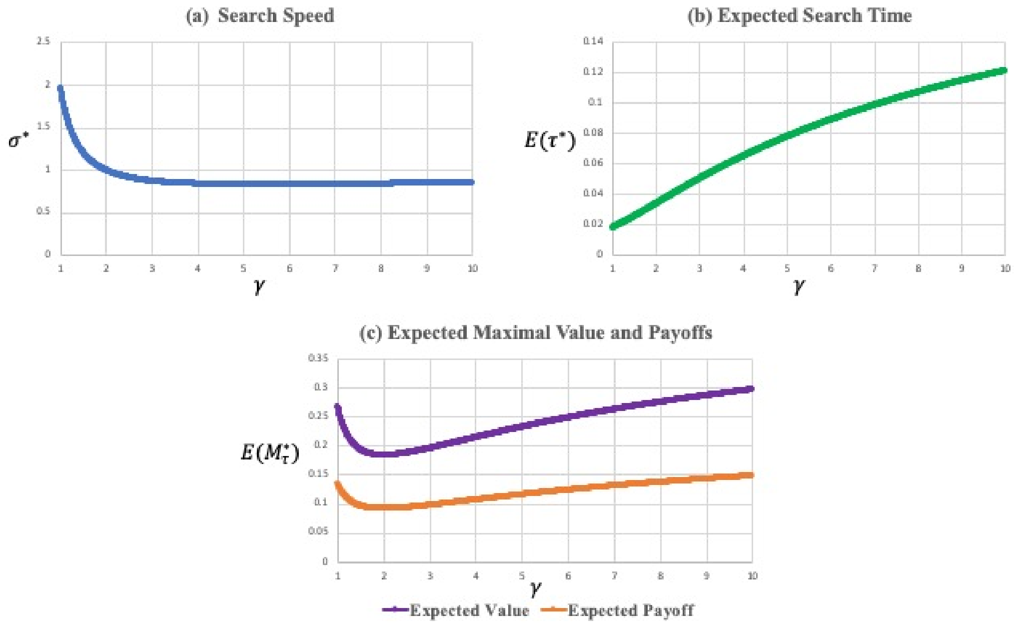

FIGURE 3: Impacts of cost changes when $c(\sigma) = \exp(\sigma^\gamma)$

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 knows whether the agent is on the job or not, but cannot monitor the effort the agent exerts in his search or the resulting observed values. That is, the principal observes neither the search speed σ nor the path of values $\{X_t\}_t$. In this setting, we can think of $\sigma_t = \underline{\sigma}$, with the agent exerting the minimal speed, as the agent shirking. For simplicity, we restrict our analysis to the case in which there is no discounting for both the principal and the agent.

Certainly, if the principal can sell the project to the agent, agency frictions cease to matter. However, in many applications that is not possible: individual researchers infrequently acquire universities or research labs, galleries do not always purchase the full portfolio of art they display, and real-estate agents do not buy the entire stock of houses they represent. We therefore consider contractual relationships that do not entail ownership transfers.

Specifically, 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.[17]

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

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

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

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

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

generalizing Propositions 1 and 2:

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

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

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

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

Intuitively, a wage effectively shifts downward the search costs by a constant amount. That is, the agent effectively considers a cost $\tilde{c}(\cdot) = c(\cdot) - w$. The formula for the optimal search speed 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 speed 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 speed 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 speed. Since costs are convex, $c'(\sigma)\sigma$ is non-decreasing in σ, and 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 speed. 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, speed. We note, however, that, in general, wages may be low, even negative. Indeed, we do not impose any limited-liability constraints. The inclusion of a lower boundary on admissible wages would not alter the methods we present and, if anything, would make interior search intensities easier to sustain optimally as wages would naturally be forced to be higher. We maintain no such constraints both for presentation simplicity and for realism. Indeed, it is not uncommon for commissioned researchers to rent labs or testing equipment, cover various experimental outlays, etc.[18] Such flow expenses would formally translate into negative wages.

We now turn to the returns of commissioned search resulting from the agent's optimal policy. For each commission contract (w, α), we can identify the expected payoff to both the principal, denoted by $\mathbf{E}(\Pi^P_{w,\alpha})$, and the agent, denoted by $\mathbf{E}(\Pi^A_{w,\alpha})$. The precise formulation of these expected payoffs naturally generalizes our characterization in Proposition 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:*

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

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:

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

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

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

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

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 speed the principal optimally targets should satisfy a first-order condition whenever interior. As a consequence, we have the following characterization:[19]

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

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

Contracting with Independent Search When search is over independent samples, the optimal commission contract is far less tractable. To our knowledge, it has not been investigated for general distributions of instantaneous values. As we show in the Supplementary Materials, when (discrete) independent observations are normally distributed, the optimal control σ depends on both the wage and the commission level. This contrasts our setting, where the optimal control depends on flow wages alone. Intuitively, when observations are independent, the agent stops search when the *current search observation* is equal to the option value of continuing. Any instantaneous decision accounts for both the wage, which affects net search costs, and the commission, which affects current and future payoffs. The dependence is not amenable to analysis.[20]

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 exponential costs in order to illustrate simply how the optimal commission contract responds to the environment's features.

5.3 Contracting with Exponential Search Costs

Suppose search costs are exponential, $c(\sigma) = e^{b\sigma}$, where $b > 0$. For the agent searching on his own, the optimal constant search speed would be given by $\sigma^{NC} = \frac{2}{b}$, see Proposition 1.

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

[20] It is given by maximizing $R(\sigma, w, \alpha) = \psi^{-1}(\frac{c(\sigma)-w}{\alpha})\sigma$, where $\psi(x) = \phi(x) - x(1 - \Phi(x))$, with ϕ and Φ representing the density and cumulative distributions of the standard normal, respectively.

Using Proposition 5 above, the optimal commission contract (w^*, α^*) satisfies

$$w^* = e^{\frac{1}{2}(\sqrt{17}-1)} - \frac{1}{4}\left(\sqrt{17}-1\right)e^{\frac{1}{2}(\sqrt{17}-1)} \quad \text{and} \quad \alpha^* = \frac{\sqrt{17}-1}{2\left(\frac{1}{2}\left(\sqrt{17}-1\right)+2\right)} \approx 0.438.$$

with induced search speed of $\sigma^* = \frac{\sqrt{17}-1}{2b}$.

The agent searches with higher speed when under a commission contract relative to when searching on his own. Intuitively, higher wages shift costs down and induce the agent to search at greater speeds, see Section 4.3. The principal therefore trades off the time it takes the agent to search and the wages needed to induce the necessary speeds. Costs being exponential simplifies dramatically the structure of the optimal contract. In particular, wages are set so that the resulting instantaneous costs—and, in the exponential case, the marginal costs as well—are independent of b. As a consequence, the optimal commission and wages are both independent of the cost parameter b.

Proposition 4 suggests that the expected maximal value of the project is $\mathbf{E}[M_\tau^*] \approx \frac{0.126}{b^2}$, while the expected search duration is $\mathbf{E}[\tau^*] \approx \frac{0.007}{b^2}$. Thus, the overall expected maximum and search duration are decreasing in the cost parameter b. The resulting expected payoffs to the agent and the principal are then:

$$\mathbf{E}(\Pi^A_{w^*,\alpha^*}) = \frac{\left(31-7\sqrt{17}\right)^2 e^{\frac{1}{2}-\frac{\sqrt{17}}{2}}}{\left(\sqrt{17}-1\right)^3 b^2} \quad \text{and} \quad \mathbf{E}(\Pi^P_{w^*,\alpha^*}) = \frac{\left(3\sqrt{17}-11\right)e^{\frac{1}{2}-\frac{\sqrt{17}}{2}}}{4b^2}.$$

In other words, with exponential costs, no matter the cost parameter, the agent receives approximately 30% of the overall surplus through the commission contract.

With a single-agent's retrospective search, Proposition 3 suggests an expected payoff of $\mathbf{E}(\Pi^*) = \frac{1}{b^2e^2} \approx \frac{0.135}{b^2}$. In contrast, the overall surplus generated by the agent and the principal is $-\frac{16\left(33\sqrt{17}-137\right)e^{\frac{1}{2}-\frac{\sqrt{17}}{2}}}{\left(\sqrt{17}-1\right)^3 b^2} \approx \frac{0.103}{b^2}$. Thus, with exponential costs, contractual frictions come at a cost just shy of *one quarter*, or about 24%, of the surplus.

6 Conclusions and Discussion

This paper proposes a simple model of retrospective search. Agents—online shoppers, politicians, geological survey teams, or medical drug developers—observe evolving options that are correlated over time. They have two decisions at each point: at what speed to search and whether to stop and collect the maximal observed value. The optimal search policy entails a U-shaped search speed. The optimal search speed 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 speed of search depends only on the wage, while the induced stopping boundary depends on both the wage and the commission. When speed of search entails exponential costs, we show that contractual frictions come at a cost of roughly one quarter of the surplus.

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

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 speed do not depend on the linearity of the agent's utility. The characterization of the optimal speed 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? 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 the body of the paper, it is easily solvable. In general, we show in the Supplementary Materials that it ties to a well-known class of ODEs, *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 policies 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. However, 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 keen 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 suits the search applications we have in mind: when looking for new wells, the right drug dosages, 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 speed, our analysis remains virtually identical. The stopping boundary does need to be adjusted, however, with greater drifts associated with more lenient stopping boundaries: with a substantial drift, the agent has a strong incentive to continue searching.

One could also contemplate an agent controlling the drift, instead of the search speed, as considered by Peskir (2005). That turns out to generate 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.

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 speed 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:

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

Equivalently,

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

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 speed 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 $\tilde{\sigma}_M(M_t, X_t, \sigma_t)$.[21] 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

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

arguments whenever no confusion is caused, we have

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

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

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

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:

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

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:

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

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 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 $V_x(g(M), M) = u_x(M) = 0$.[22] 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

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

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, Henry Jr, 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 unkilled 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 unkilled process X.[23]

2. The (controlled) infinitesimal generator of the process $\hat{Z} = (\hat{M}, \hat{X})$, denoted $\mathcal{A}_{\hat{Z}}$, 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:[24]

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

[24] The 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 unkilled diffusions.

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

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

A.2 Reducing Dimensionality

The following lemma, which we use throughout our analysis, suggests that a marginal change in search speed 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 speed, 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, Shepp, and Shiryaev (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.

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,

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

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 $t\to 0$, 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 $\frac{\partial W(M,M)}{\partial M}\left(\lim_{t\to 0}\frac{\mathbf{E}_{M,M}(M_t-M)}{t}\right)$. Therefore,

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

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

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:

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

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

Let σ_t^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:

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

As described in our background section, at the point of stopping, the agent's utility from the achieved maximum coincides with her continuation value: $u(M_\tau) = V(M_\tau, X_\tau)$. Furthermore, the stopping time has to be of the form $\tau^* = \inf\{t \geq 0 : X_t \leq g^r(M_t)\}$ 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:

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

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

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

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,

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

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

$$P(X_{\tau_{a,b}=a}|M,X) = \frac{S(b)-S(x)}{S(b)-S(a)} = \frac{b-x}{b-a}, \text{ and} \qquad (3)$$
$$P(X_{\tau_{a,b}=b}|M,X) = \frac{S(x)-S(a)}{S(b)-S(a)} = \frac{x-a}{b-a},$$

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:[25]

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

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

$$\mathbf{E}\left(\int_0^{\tau_{a,b}} f(X_t)dt \middle| X\right) = \int_a^b f(y)G^r_{a,b}(X,y)m(dy), \qquad (4)$$

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:

$$E\left(-\int_0^{\tau_{g(M),M}} [c(\sigma(x_t))]d_t \middle| M,X\right) = -\int_{g(M)}^M G^r_{g(M),M}(X,y)(c(\sigma(M,y)))m(dy).$$

Thus,

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

Reorganizing the above,

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

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

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

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

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

From Borodin and Salminen (2012) (Appendix 1, page 105), the Green's function of

[25] Conditioning on values of M or X implies these are the current values of the maximum value or the search observation, respectively.

the exponentially killed Brownian motion is given by:

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

Taking the limit as X approaches $g(M)$ and using the smooth-pasting conditions,

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

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

$$V(M,X) = u(M) + (X-g(M))\int_{g(M)}^{M} \frac{\sinh(\sqrt{2r}(M-y))}{\sinh(\sqrt{2r}(M-g(M)))}\frac{2c(\sigma(M,y))}{\sigma(M,y)^2}dy \tag{7}$$

$$-\int_{g(M)}^{M} G^r_{g(M),M}(X,y)\frac{2c(\sigma(M,y))}{\sigma(M,y)^2}dy.$$

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:

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

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 X:

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

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

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

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

$$\frac{2c(\sigma(M,X))}{\sigma(M,X)^2} - \frac{c'(\sigma(M,X))}{\sigma(M,X)} = \int_{g(M)}^{M} 2rG^r_{g(M),M}(X,y)\frac{2c(\sigma(M,y))}{\sigma(M,y)^2}dy \tag{8}$$

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

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

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

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

By definition, the Green's function is symmetric. That is, $G^{r}_{g(M),M}(X,y)=G^{r}_{g(M),M}(y,X)$. 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:

$$r\int_{g(M)}^{M}G^{r}_{g(M),M}(X,y)2dx=r\mathbf{E}\left(\int_{0}^{\mathbb{T}_{[g(M),M]}}e^{-rt}dt|X\right).$$

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 $r\mathbf{E}\left(\int_{0}^{\mathbb{T}_{[g(M),M]}}e^{-rt}dt|X\right)=\mathbf{E}(1-e^{-r\mathbb{T}_{[g(M),M]}})$, 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, $\lim_{r\to 0}\mathbf{E}(e^{-r\mathbb{T}_{[g(M),M]}}|X)=1$, 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):*

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

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:

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

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

$$(g'(M)-1)\left(\int_{g(M)}^{M}\frac{2c(\sigma(M,y))}{\sigma(M,y)^2}\frac{\sinh\left(\sqrt{2}\sqrt{r}(y-g(M))\right)}{\sinh\left(\sqrt{2}\sqrt{r}(M-g(M))\right)}\sinh\left(\sqrt{2}\sqrt{r}(y-g(M))\right)dy\right) \quad (11)$$

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

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

$$-(M-g(M))\frac{2c(\sigma(M,M))}{\sigma(M,M)^2}+u'(M)=0.$$

As mentioned in the text, analysis following Peskir (1998) shows that the optimal stopping boundary is the maximal solution $g(M) \le M$ 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 $u'(M)$.[26] 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 $g^r(M)$

[26] Peskir (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."

is linear in M (with a derivative of 1), we can simplify equation (11) to:

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

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:

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

Now, recall that the continuation HJB is:

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

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

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

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

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

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

$$\frac{2c(\sigma(M,X))c'(\sigma(M,X))\Upsilon'(X)}{c'(\sigma(M,X))^2(1-\Upsilon(X))+c(\sigma(M,X))c''(\sigma(M,X))\Upsilon(X)} = \sigma_X(M,X).$$

By definition, $\Upsilon(g(M)) = 1$ and $\Upsilon'(g(M)) = -\sqrt{2r}\sinh(\sqrt{r/2}(M-g(M))$. Plugging these back in and simplifying further yields:

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

Plugging this into equation (12),

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

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 $d^r = M - g^r(M)$:

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

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:

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

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:

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

Similarly, from equation (8), we have

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

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 speed. 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:

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

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:

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

subject to

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

The optimal search speed σ^* is pinned down uniquely by the choice of w. It follows that if the principal induces a search speed 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:

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

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

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

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

$$\begin{aligned}
f_{\alpha\alpha} &\equiv -\frac{\sigma c'(\sigma) + 2c(\sigma)}{c'(\sigma)^2},\\
f_{\sigma\sigma} &\equiv \frac{\alpha}{2c'(\sigma)^4}\Bigg(\left(-12\alpha c(\sigma)c''(\sigma)^2 + c'(\sigma)^2\left((\alpha-2)\sigma c'''(\sigma) + (8\alpha-4)c''(\sigma)\right)\right)\\
&\quad + c'(\sigma)\left(4\alpha c(\sigma)c'''(\sigma) - 2(\alpha-2)\sigma c''(\sigma)^2\right)\Bigg),\\
f_{\alpha\sigma} &\equiv \frac{4\alpha c(\sigma)c''(\sigma) + (1-3\alpha)c'(\sigma)^2 + (\alpha-1)\sigma c'(\sigma)c''(\sigma)}{c'(\sigma)^3}.
\end{aligned}$$

Since $f_{\alpha\alpha}$ is clearly negative, the first-order approach is valid whenever $f_{\alpha\alpha}f_{\sigma\sigma} - (f_{\alpha\sigma})^2 \geq 0$. For example when the costs are exponential with coefficient b, that is $c(\sigma) = e^{b\sigma}$,

$$\begin{aligned}
\alpha^* &= \frac{\sqrt{17}-1}{2\left(\frac{1}{2}\left(\sqrt{17}-1\right)+2\right)} \approx 0.438447\\
\sigma^* &= \frac{\sqrt{17}-1}{2b}\\
w^* &= e^{\frac{1}{2}(\sqrt{17}-1)} - \frac{1}{4}\left(\sqrt{17}-1\right)e^{\frac{1}{2}(\sqrt{17}-1)}
\end{aligned}$$

To observe that this is indeed a local maximer we check the hessian at the optimum

$$f_{\alpha\alpha}f_{\sigma\sigma} - (f_{\alpha\sigma})^2|_{(\alpha^*,\sigma^*)} = \frac{\left(7\sqrt{17}+17\right)e^{1-\sqrt{17}}}{\left(\sqrt{17}+3\right)^2 b^2} > 0$$

Thus α^*, σ^* w^* are local maximizers for the principal. We also need to ensure that given b, α and w the σ^* choice is also a maximizer for the agent. First we immediately notice that it must be the case that for any wage that is considered we must have $w < c(\underline{\sigma})$ as otherwise the gains of the agent is unbounded and the search newer concludes. The total payoff of the agent as a function σ given that b, w and α is:

$$\frac{(\alpha)^2\sigma^2}{4\left(e^{bs} - w\right)}$$

For an interior maximizer the second order condition with respect to σ is

$$\frac{\alpha^2\left(we^{b\sigma}\left(b^2\sigma^2 + 4b\sigma - 4\right) + e^{2b\sigma}\left(b^2\sigma^2 - 4b\sigma + 2\right) + 2w^2\right)}{4\left(e^{b\sigma} - w\right)^3} < 0$$

we know we need $w < e^{b\sigma}$ for all $\sigma \in [\underline{\sigma}, \bar{\sigma}]$ thus for $b > 1$ and $\underline{\sigma}$ such that $e^{b\underline{\sigma}} > w^*$ the agent's solution is indeed a maximizer. Whenever the first-order approach is invalid, the principal chooses a boundary solution that, in turn, induces a boundary search speed for the agent. The payoffs to the principal and and the agent are respectively, $\frac{(3\sqrt{17}-11)e^{\frac{1}{2}-\frac{\sqrt{17}}{2}}}{4b^2}$ and $\frac{(31-7\sqrt{17})^2 e^{\frac{1}{2}-\frac{\sqrt{17}}{2}}}{(\sqrt{17}-1)^3 b^2}$ and a total welfare of $-\frac{16(33\sqrt{17}-137)e^{\frac{1}{2}-\frac{\sqrt{17}}{2}}}{(\sqrt{17}-1)^3 b^2}$. Without the contractual frictions again under the assumption that $b > 1$ the expected payoff from the search is $\frac{1}{b^2e^2} \approx \frac{0.135335}{b^2}$. In other words the loss of welfare due to contractual frictions correspond to approximately 24 percent.

■

References

Aghion, P. and J. Tirole (1994). The management of innovation. *The Quarterly Journal of Economics 109*(4), 1185–1209.

Borodin, A. N. and P. Salminen (2012). *Handbook of Brownian motion-facts and formulae*. Birkhäuser.

Brock, K., L. Billingham, M. Copland, S. Siddique, M. Sirovica, and C. Yap (2017). Implementing the efftox dose-finding design in the matchpoint trial. *BMC Medical Research Methodology 17*(1), 1–15.

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.

Cankurtaran, P., F. Langerak, and A. Griffin (2013). Consequences of new product development speed: A meta-analysis. *Journal of Product Innovation Management 30*(3), 465–486.

Cetemen, D., C. Urgun, and L. Yariv (2022). Collective progress: Dynamics of exit waves. *mimeo*.

Chen, J., F. Damanpour, and R. R. Reilly (2010). Understanding antecedents of new product development speed: A meta-analysis. *Journal of Operations Management 28*(1), 17–33.

De los Santos, B., A. Hortaçsu, and M. R. Wildenbeest (2012). Testing models of consumer search using data on web browsing and purchasing behavior. *American Economic Review 102*(6), 2955–2980.

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.

Hodgson, C. (2021). Information externalities, free riding, and optimal exploration in the uk oil industry. *mimeo*.

Hodgson, C. and G. Lewis (2021). You can lead a horse to water: Spatial learning and path dependence in consumer search. *mimeo*.

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, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics*, 113–126.

Miller, R. A. (1984). Job matching and occupational choice. *Journal of Political economy 92*(6), 1086–1120.

Mortensen, D. T. (1970). Job search, the duration of unemployment, and the phillips curve. *American Economic Review 60*(5), 847–862.

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.

Obłój, J. (2007). The maximality principle revisited: On certain optimal stopping problems. In *Séminaire de probabilités XL*, pp. 309–328. Springer.

Øksendal, B. (2003). *Stochastic differential equations*. Springer.
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.
Pedersen, J. L. (2000). Discounted optimal stopping problems for the maximum process. *Journal of applied probability 37*(4), 972–983.
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.
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.
Stigler, G. J. (1961). The economics of information. *Journal of Political Economy 69*(3), 213–225.
Taylor, H. M. et al. (1975). A stopped brownian motion formula. *The Annals of Probability 3*(2), 234–246.
Thall, P. F. and J. D. Cook (2004). Dose-finding based on efficacy–toxicity trade-offs. *Biometrics 60*(3), 684–693.
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.
Wong, Y. F. (2022). Spatial experimentation: Optimal costly exloration of correlated alternatives. *mimeo*.
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.