

Benkert, Jean-Michel; Letina, Igor

Working Paper

Never-ending search for innovation

Discussion Papers, No. 24-10

Provided in Cooperation with:

Department of Economics, University of Bern

Suggested Citation: Benkert, Jean-Michel; Letina, Igor (2024) : Never-ending search for innovation, Discussion Papers, No. 24-10, University of Bern, Department of Economics, Bern

This Version is available at:

<https://hdl.handle.net/10419/308756>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by/4.0/>



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics
and Social Sciences

Department of Economics

Never-ending Search for Innovation

Jean-Michel Benkert, Igor Letina

24-10

December, 2024

DISCUSSION PAPERS

Schanzeneckstrasse 1
CH-3012 Bern, Switzerland
<http://www.vwi.unibe.ch>

Never-ending Search for Innovation

Jean-Michel Benkert and Igor Letina*

December 2024

Abstract

We provide a model of investment in innovation that is dynamic, features multiple heterogeneous research projects of which only one potentially leads to success, and in each period, the researcher chooses the set of projects to invest in. We show that if a search for innovation starts, it optimally does not end until the innovation is found—which will be never with a strictly positive probability.

Keywords: Innovation, Optimal Search, Infinite Horizon

JEL Codes: D83, O31

*Benkert: Department of Economics, University of Bern. Letina: Department of Economics, University of Bern and CEPR. We thank Georg Nöldeke for very valuable comments and Martial Zurbiggen for excellent research assistance.

1 Introduction

In 1880, Alphonse Laveran discovered that malaria was caused by a parasite from the *Plasmodium* genus. This discovery immediately led to the next question: how was *Plasmodium* transmitted from one person to another? To answer such a question, researchers could design and conduct various experiments to test the many potential vectors, including transmission via soil, water, or a living host such as a bird or—which ultimately proved true—a mosquito (Desowitz, 1991). For example, researchers could expose healthy subjects to infected water to determine if the disease could be transmitted that way, or allow mosquitoes to feed on infected individuals and then examine the insects for the presence of the parasite.¹ There are countless experimental variations, such as the species of mosquito used, the time elapsed after infection before examination, the specific parts of the mosquito dissected, and the techniques employed to detect the parasite. Researchers, aiming to find the answer as efficiently as possible, would pursue the most promising and least costly experiments first, and move on to more challenging, expensive, or less likely experimental avenues if success initially proves elusive.

The above example has two important features. First, given a set of potential vectors, discovering that one is not responsible for transmission is informative about the probability that a remaining vector is how the disease actually spreads. For instance, if the parasite is not found in the soil, the probability that it will be found in water increases. Second, the cost of testing different vectors varies greatly. For example, examining soil samples is cheap, while dissecting living hosts is costly.

We propose a model exhibiting both of the above features, informational externalities and heterogeneous projects. In doing so, we go beyond the existing literature and we find that, if the search starts, it does not end until the answer is found. In particular, the search will go on forever with strictly positive probability. This is in sharp contrast with the existing results in the literature, where the search tends to end in finite time even if the optimal search rule prescribes to continue searching forever. This can be seen in McCardle, Tsetlin, and Winkler (2018), who study monopolistic search in an infinite-horizon, continuous-time bandit model. In their model, depending on the arrival rate of successes and failures, the optimal stopping time is finite or infinite. However, even when the optimal stopping time is infinite, the search will end in finite time with a probability going to 1, as eventually either a success or a failure will realize.²

¹These are some of the experiments conducted by Ronald Ross, who in 1897 demonstrated that mosquitoes transmitted malaria.

²A similar observation holds true in the seminal search models (without informational externalities) such as McCall (1965) and Morgan (1983). See Lippman and McCall (1976) and Chade, Eeckhout, and Smith (2017) for surveys, or for a more recent example see Benkert, Letina, and Nöldeke (2018). In this class of models, even if the optimal search rule prescribes searching forever, the probability that search continues forever will converge to zero, thus ending search in finite time.

Instead of the bandit framework, we study a model where the agent is faced with a set of research projects of which only one has the potential to lead to the breakthrough. Like the bandit framework, our model also features learning, as failure from one set of projects is informative about the chances that a success will be found in the remaining set. Our model is closely related to de Roos, Matros, Smirnov, and Wait (2018), who also consider a model where an agent searches in a set of research projects. Importantly, they assume that all projects are homogeneous, and that a successful project exists with certainty. When this is the case, then the optimal search prescribes that the entire set of available projects is examined in finite time, causing the search to also end in finite time.³ We instead assume that projects are sufficiently heterogeneous, so that a rational agent optimally never examines the entire set of available projects. Paired with the possibility that no project could be successful, we obtain the contrasting result that optimal search may never end.

Heterogeneity of research project is not sufficient for the search to potentially remain active forever. This can be best seen in Weitzman (1979), who famously showed that optimal search ends—in finite time—when some reservation value is met. In Weitzman (1979), opening one box is not informative about the value of yet unopened boxes. Thus, after the agent has exhausted all the promising boxes, she stops. In contrast, Bayesian updating leads our agent to become more optimistic about the value of remaining projects, and, as we show, this causes the agent to never willingly abandon an active search.

In what follows, we formally introduce our model (Section 2), present our result on potentially never-ending search (Section 3), and conclude with a discussion of the result and its policy implications (Section 4).

2 Model

There is a set of heterogeneous research projects $J = [0, 1)$ and a single agent searching for an innovation in J . In any period $t = 1, 2, \dots$, the agent can examine an arbitrary (measurable) set of projects $S \subseteq J$ at the cost $C(S) = \int_S c(j) dj$. We assume that $c(j)$ is continuous, strictly increasing, and satisfies $\lim_{j \rightarrow 1} c(j) = \infty$.

With probability $p \in (0, 1)$, the innovation is *feasible*, meaning that there is a single project $\hat{j} \in J$, which, if examined, will lead to a successful innovation of value v . Examining any other project leads to a dead end. We assume that any project $j \in J$ is equally likely to yield success. The future costs and payoffs are discounted by $\delta \in (0, 1)$.⁴

³Using a similar model as de Roos et al. (2018), Matros and Smirnov (2011) and Matros and Smirnov (2016) focus on the strategic interaction between multiple searchers. In extensions, they consider heterogeneous projects (in terms of success probabilities rather than costs) and allow for no successful project to exist. Nevertheless, they always conclude that search ends in finite time, contrasting with our findings.

⁴This is a dynamic version of the model introduced in Letina (2016).

3 Never-ending Search

An agent's search strategy $\sigma = (L_1, L_2, \dots)$ consists of a sequence of (potentially empty) measurable sets $L_t \subseteq [0, 1) \cup \emptyset$, $t = 1, 2, \dots$ determining which projects will be examined in period t if previous periods did not yield success. We say that a search strategy σ is *interval-based* if the subsets L_t are intervals. Moreover, we say that a search strategy is *increasing-interval-based* if the subsets L_t are intervals such that $\sup L_t = \inf L_{t+1}$.

Let $S_t := \cup_{t' < t} L_{t'}$ and $U_t := [0, 1) \setminus S_t$. Thus, S_t represents the set searched up to time $t - 1$, and U_t is the set still unsearched when deciding at time t . Therefore, any optimal search in period t will occur in U_t .

Let $\hat{p}_\sigma(t) := \Pr[\hat{j} \in U_t \mid \hat{j} \notin S_t]$ denote the agent's Bayesian posterior belief that success is feasible at time t when following the strategy σ , i.e., $\hat{p}_\sigma(t) = \frac{\mu(U_t)p}{\mu(U_t)p + (1-p)}$, where μ denotes the Lebesgue measure. Then, we can write the agent's value function at time t , when S_t has been searched so far and when searching according to strategy $\sigma = (L_1, L_2, \dots)$, as

$$V_\sigma(S_t, t) = \hat{p}_\sigma(t) \frac{\mu(L_t)}{\mu(U_t)} v - C(L_t) + \delta \left(1 - \hat{p}_\sigma(t) \frac{\mu(L_t)}{\mu(U_t)} \right) V_\sigma(S_t \cup L_t, t + 1).$$

We can now state our result, the proof of which is relegated to the appendix.

Proposition 1 *The optimal search strategy σ^* is such that either $L_t = \emptyset$ for all t or $\mu(L_t) > 0$ for all t .*

Intuitively, there are two forces leading to the result. First, the informational externality of not finding the innovation induces the agent to continue searching. Second, the increasing costs of searching dampen the incentives to continue the search. The former effect ensures that the search continues, whereas the latter effect (and in particular the assumption that $\lim_{j \rightarrow 1} c(j) = \infty$) prevents the search from being exhaustive.

More specifically, if, contrary to our result, there were a final period T in which the search occurred, the agent would want to search all projects with a positive expected value in that period, implying that the marginal project would have an expected value of zero. Because of the increasing costs of search, that marginal project would have to be interior. However, if there was no success in period T , the belief that the marginal project will succeed would increase, implying a strictly positive expected value in the neighborhood of the marginal project. Thus, stopping the search could not be optimal. Further, as the remaining projects are increasingly costly to search, exhaustively searching all remaining projects and ending the search in finite time is never optimal. Hence, once started, the search only stops if the innovation is discovered. Since the innovation is not feasible with probability $1 - p > 0$, the search could continue forever with a strictly positive probability, leading to the following corollary.

Corollary 1 *If any search is optimal, the probability of active search in any period $t \geq 1$ is strictly above $1 - p > 0$.*

This corollary highlights the distinction of our result relative to the existing literature, where even if the optimal strategy is to search forever, the search ends in finite time. In contrast, our result is not only about the optimal strategy mandating infinite search but that the actual search may last forever.

Finally, for our results to be meaningful, we need to establish conditions under which some search is indeed optimal.

Lemma 1 *If $pv > c(0)$, it is optimal for the agent to engage in the search.*

It is immediate that this condition ensures that the agent would rather search once than not search at all.⁵ Intuitively, the agent compares the marginal benefit of searching the initial interval with the marginal cost of doing so. The condition ensures that the former exceeds the latter, so that search starts. However, by Proposition 1 we know that the agent would not cease searching after one period but continue until a success is found.

4 Conclusion

At first glance, our result may seem counterintuitive: why would a rational agent be willing to keep (forever) searching for a prize of finite value at eventually unbounded marginal costs? To reconcile our result with this (incomplete) intuition, we need to also consider the search *intensity*. Even as marginal costs increase, the agent can adapt the intensity of search so that actual search costs do not increase from period to period. Effectively, while active search may never end, the associated search intensity will converge to zero in the limit because of the unbounded increase in the marginal search costs.

This is not merely a technical point, but rather has an important policy implication for the question of when to optimally abandon a (research) project. McCardle et al. (2018) note that, in practice, search often lasts longer than what is *perceived* as optimal and suggest that one should set an “alarm clock” to ensure a timely end to the search effort. Our results suggest that this might not always be the appropriate policy, and that a research project may optimally need to continue for longer than intuition alone would recommend. According to our results, an outside decision maker (who does not fully account for the informational content of the failures that the agent has experienced) is especially likely to perceive the search as going beyond the optimal threshold. This is particularly relevant in environments with informational spillovers (so that failure today

⁵In the one-period model, searching an interval of length q^* is optimal if and only if $pv = c(q^*)$. Hence, the condition in Lemma 1 yields $q^* > 0$, ensuring that the agent would rather search once than not search at all.

is informative of success chances tomorrow) and adjustable research intensity, which we would argue are more common than not in research. Empirically, our result also highlights the value of collecting data on research intensity, in addition to data on whether research is active or not.

In short, the main message of our paper is that external decision makers should exercise caution when recommending that a research project be abandoned. What might appear to be the stubbornness of a researcher unwilling to admit defeat could, in fact, be a rational response to the lessons learned from past failures.

Appendix: Proof of Proposition 1

Lemma 2 *The optimal search strategy σ^* satisfies: there do not exist intervals L and L'' with $\mu(L), \mu(L'') > 0$ and $\inf L > \sup L''$ such that L is searched and L'' is not.*

PROOF: Suppose not. Then, there exists a set L , searched at some time t and a set L' , with $\inf L > \sup L'$ and $\mu(L') = \mu(L)$, that is never searched. As L is searched at time t , we can write $L_t = L \cup S$ where S may be empty. Then, the value function at time t when searching according to σ^* is

$$\begin{aligned} V_{\sigma^*}(S_t, t) &= \hat{p}_{\sigma^*}(t) \frac{\mu(L \cup S)}{\mu(U_t)} v - C(L \cup S) \\ &\quad + \delta \left(1 - \hat{p}_{\sigma^*}(t) \frac{\mu(L \cup S)}{\mu(U_t)} \right) V_{\sigma^*}(S_t \cup L \cup S, t+1). \end{aligned}$$

Consider the alternative strategy σ' , coinciding with σ^* except that L is replaced by L' at time t . Then

$$\begin{aligned} V_{\sigma'}(S_t, t) &= \hat{p}_{\sigma'}(t) \frac{\mu(L' \cup S)}{\mu(U_t)} v - C(L' \cup S) \\ &\quad + \delta \left(1 - \hat{p}_{\sigma'}(t) \frac{\mu(L' \cup S)}{\mu(U_t)} \right) V_{\sigma'}(S_t \cup L' \cup S, t+1). \end{aligned}$$

The difference in value functions $V_{\sigma'}(S_t, t) - V_{\sigma^*}(S_t, t)$ reads

$$\begin{aligned} &C(L \cup S) - C(L' \cup S) \\ &+ \delta \left(1 - \hat{p}_{\sigma}(t) \frac{\mu(L' \cup S)}{\mu(U_t)} \right) (V_{\sigma'}(S_t \cup L' \cup S, t+1) - V_{\sigma^*}(S_t \cup L \cup S, t+1)) \\ &= C(L \cup S) - C(L' \cup S) > 0. \end{aligned}$$

The equality obtains because the deviation from σ^* to σ' entails no change in the posteriors, i.e., $\hat{p}_{\sigma'}(t') = \hat{p}_{\sigma^*}(t')$ for all t' , as the measures of the searched and unsearched sets are the same with both strategies. Hence, since the two strategies coincide for any $t' > t$, we have $V_{\sigma'}(S_t \cup (L' \cup S), t+1) = V_{\sigma^*}(S_t \cup (L \cup S), t+1)$. Finally, the inequality follows from the fact that $\inf L > \sup L'$ and c is increasing. This contradicts the optimality of σ^* . ■

Lemma 3 *The optimal search strategy σ^* satisfies: there do not exist intervals L and L'' with $\mu(L), \mu(L'') > 0$ and $\sup L < \inf L''$ such that L is searched at time t and L'' at time $t' < t$.*

PROOF: Suppose not. Then, there exist sets L and L' that are searched at times t and t' with $t > t'$, $\sup L < \inf L''$ and $\mu(L') = \mu(L)$. Consider the alternative strategy σ' which

coincides with σ^* except that L and L' are swapped at t and t' , respectively. Analogously to the proof of Lemma 2, the posterior beliefs and success probabilities induced by the two strategies are the same. Hence, the only difference in the value functions is induced by a change in the timing of the search costs. The swap reduces the search costs at t' by K , while increasing them by the same amount at t , but this increase is discounted by at least δ . Thus, the net reduction in costs is at least $K - \delta K$, implying $V_{\sigma'}(S_t, t) - V_{\sigma^*}(S_t, t) > 0$, contradicting the optimality of σ^* . ■ Lemmas 2 and 3 imply the following:

Corollary 2 *The optimal search strategy σ^* is increasing-interval-based.*

With this we can reinterpret the strategy $\sigma = (L_1, L_2, \dots)$ as a collection of increasing intervals, i.e., $[0, l_1), [l_1, l_2)$ and so forth.

Lemma 4 *The optimal search strategy σ^* satisfies: if there exists a period t with $\mu(L_t) = 0$, then $\mu(L_{t'}) = 0$ for all $t' > t$.*

PROOF: Suppose not. Then, there exists a period t with $\mu(L_t) = 0$ and a period $t' > t$ with $\mu(L_{t'}) > 0$. Consider the alternative strategy σ' which coincides with σ^* except that for all $s \geq t$ we replace L_s by L_{s+1} . Then, $V_{\sigma'}(S_t, t) = V_{\sigma^*}(S_{t+1}, t+1)$ and $V_{\sigma^*}(S_t, t) = \delta V_{\sigma^*}(S_{t+1}, t+1)$. Since stopping search yields a payoff of zero, it cannot be that $V_{\sigma^*}(S_{t'}, t') < 0$ for any t' . Additionally, since $\mu(L_{t'}) > 0$, it must be that $V_{\sigma^*}(S_{t'}, t') > 0$, since otherwise removing some subset from L' would strictly increase the payoff in t' . Thus, $V_{\sigma^*}(S_{t+1}, t+1) > 0$ and since $\delta < 1$, then $V_{\sigma'}(S_t, t) > V_{\sigma^*}(S_t, t)$. ■

We can now prove Proposition 1.

PROOF: Lemma 4 implies that only three types of strategies can be optimal: 1) never search, 2) search in every period until time $T > 0$ (unless there's a success) and then not search anymore, 3) search in every period until there is a success. Thus, we only need to show that 2) cannot be optimal. Suppose there is a $T > 0$ such that the agent searches in every period $t \leq T$ and then stops. Then, the agent's value function at time T , the last period of search, is given by

$$V_{\sigma^*}(S_T, T) = \frac{(l_T - l_{T-1})p}{(1 - l_{T-1})p + (1 - p)}v - C([l_{T-1}, l_T)),$$

because $V_{\sigma^*}(S_T \cup [l_T - l_{T-1}), T + 1) = 0$ and by Bayes' rule $\hat{p}_{\sigma^*}(T) = \frac{(1-l_{T-1})p}{(1-l_{T-1})p+(1-p)}$. Further, because it is optimal to search in period T , l_T is implicitly defined by⁶

$$\frac{\partial V_{\sigma^*}(S_T, T)}{\partial l_T} = 0 \Leftrightarrow \frac{pv}{(1 - l_{T-1})p + (1 - p)} = c(l_T). \quad (1)$$

⁶Observe that the value function is strictly concave in l_T as c is strictly increasing.

Since $\lim_{j \rightarrow 1} c(j) = \infty$ and the LHS of equation (1) is finite, it cannot be optimal that $l_T = 1$ and the interval $[0, 1)$ is searched in finite time.

Thus, take $l_T < 1$ and consider the agent's incentives in period $T + 1$. In particular, consider one more period of search before stopping forever. Then, the value function of searching a positive interval, i.e., setting $l_{T+1} > l_T$ is

$$V_{\sigma'}(S_{T+1}, T + 1) = \frac{(l_{T+1} - l_T)pv}{(1 - l_T)p + (1 - p)} - \int_{l_T}^{l_{T+1}} c(j)dj.$$

For σ^* to be optimal, we need to have $\frac{pv}{(1-l_T)p+(1-p)} - c(l_{T+1}) < 0$, i.e., the marginal value of continuing search needs to be negative. Rearranging and using equation (1), we obtain

$$\frac{pv}{c(l_{T+1})} < 1 - l_T p \Leftrightarrow \frac{c(l_T)}{c(l_{T+1})} < \frac{1 - l_T p}{1 - l_{T+1} p},$$

which must hold for any $l_{T+1} \in (l_T, 1)$. This is violated for l_{T+1} close enough to l_T , as the LHS converges to 1 while the RHS is constant and strictly smaller than 1 as l_{T+1} approaches l_T from above. Hence, there is an incentive to continue searching in period $T + 1$; thus, stopping in T cannot be optimal. ■

References

- BENKERT, J.-M., I. LETINA, AND G. NÖLDEKE (2018): “Optimal search from multiple distributions with infinite horizon,” *Economics Letters*, 164, 15–18.
- CHADE, H., J. EECKHOUT, AND L. SMITH (2017): “Sorting through search and matching models in economics,” *Journal of Economic Literature*, 55, 493–544.
- DE ROOS, N., A. MATROS, V. SMIRNOV, AND A. WAIT (2018): “Shipwrecks and treasure hunters,” *Journal of Economic Dynamics and Control*, 90, 259–283.
- DESOWITZ, R. S. (1991): *The Malaria Capers: More Tales of Parasites and People, Research and Reality*, New York: W.W. Norton, chapters 11–13.
- LETINA, I. (2016): “The road not taken: competition and the R&D portfolio,” *RAND Journal of Economics*, 47, 433–460.
- LIPPMAN, S. A. AND J. J. MCCALL (1976): “The Economics of Job Search: A Survey,” *Economic Inquiry*, 14, 155–189.
- MATROS, A. AND V. SMIRNOV (2011): “Treasure game,” Economics Working Paper Series 2011-10, University of Sydney.
- (2016): “Duplicative search,” *Games and Economic Behavior*, 99, 1–22.
- MCCALL, J. J. (1965): “The economics of information and optimal stopping rules,” *Journal of Business*, 38, 300–317.
- MCCARDLE, K. F., I. TSETLIN, AND R. L. WINKLER (2018): “When to abandon a research project and search for a new one,” *Operations Research*, 66, 799–813.
- MORGAN, P. B. (1983): “Search and optimal sample sizes,” *The Review of Economic Studies*, 50, 659–675.
- WEITZMAN, M. L. (1979): “Optimal search for the best alternative,” *Econometrica*, 47, 641–654.