Searching for Serendipity

Isaac Tucker* Martín González Cabello* Olav Sorenson*

Last revised: June 24, 2026

ABSTRACT

Serendipity would seem to preclude purposeful search. To understand the relationship between search and the probability of serendipitous discovery, we build a formal model of the process, an extension of the NK model. Our simulations indicate that searchers guided by theory find unexpected solutions to their problems far more often. Those who search longer, regardless of the reason, more frequently encounter solutions to problems that they had not been trying to solve. But these unanticipated use cases may represent distractions. Our simulations suggest that those searching might have found even better solutions to their original problems had they not been seduced by their serendipitous discoveries.

MANAGERIAL SUMMARY

Disruptive innovations, novel business strategies, and inspired designs are often unexpected and depend on some measure of good fortune. Though seemingly random, our model reveals that the frequency of these fortuitous discoveries depends on choices. Search guided by theory most commonly finds unexpected solutions to the problem at hand. Theory can also help to find solutions to unanticipated problems, though it does so by ensuring that search continues longer. Importantly, however, our results suggest that these unanticipated use cases often represent a trap, distracting the searcher from their original problem.

Keywords: Serendipity; Innovation; Search; Exploration; NK Model

*UCLA Anderson School of Management

Fortune brings in some boats that are not steered - Shakespeare, Cymbeline

INTRODUCTION

Many an important innovation—from aniline purple, the first synthetic dye, to Zoloft, an antidepressant—appears to have been discovered by accident, to have been an act of serendipity (Busch, 2024; Dew, 2009; Merton & Barber, 2004; Yaqub, 2018). The inventors had been searching for something else when they happened upon their fortuitous discoveries. William Henry Perkin, for example, had been trying to synthesize quinine, an antimalarial, when one of his "failed" experiments produced the purple residue that turned out to be a stable dye.

Given the value of these discoveries, the question naturally arises as to whether searchers could engineer serendipity, increasing the frequency with which it arises. At first blush, that question seems almost oxymoronic. By definition, serendipity implies that the searchers had not been seeking the unforeseen outcome. But, consider a simple example. Imagine that you want to reference an article, perhaps one on serendipity. These days, you would probably just search Google Scholar or the Web of Science to find a copy of the article. Yet, in an earlier era, you would have had to find the journal on a shelf and to page through it. Along the way, you would have encountered several other articles. Some of them might have been relevant to answering your question or they might have proved useful to something else that interested you. Physical search though less efficient, seems more likely to lead to a serendipitous discovery.

More broadly, what features of the search process influence the propensity of serendipitous discovery? When will those choices prove most valuable?

To begin to answer these questions, we introduce a formal framework for modeling serendipitous search. Agents explore NK landscapes (Kauffman, 1993; Levinthal, 1997). These landscapes simulate a search space with ridges, peaks, valleys, and plains. We extend the standard NK model in two ways: We allow a given set of N components to generate both a primary and a secondary landscape. Discovery could occur on either one. Agents exploring the space, moreover, have theories that generate expectations of where they should look (cf. Csaszar & Levinthal, 2016;

Gavetti & Levinthal, 2000). Much of the landscape nevertheless remains *terra incognita*. We define serendipity as the discovery of a high value combination on either the primary landscape in *terra incognita* or on the secondary landscape, for an unanticipated use.

Our model yields a somewhat surprising result: Serendipity, the discovery of unexpected peaks, does not arise most frequently from random sampling, meandering, or jumping into the unknown. Rather, the highest rates of serendipitous discoveries accrue to those who pursue a hypothesis, based on their beliefs about valuable locations on the landscape (e.g., Camuffo et al., 2020; Felin & Zenger, 2009; Fleming & Sorenson, 2004). Searchers interested in serendipity should take the scenic route, traveling toward expected destinations by way of uncharted paths.

Theory-based search promotes serendipity for two reasons. First, though incomplete and imprecise, theories have a basis in reality. They therefore lead agents to search regions of the space with higher value combinations. Our example of searching for an article nicely illustrates the intuition behind this result. Choosing a book at random off of the shelf or flipping to random pages rarely yields useful information. But methodically stepping through nearby titles and articles often does. Why? Even in highly interdependent fitness spaces, the peaks—the valuable points in the space—exhibit spatial autocorrelation (Kauffman, 1993). They tend to cluster in the combinatorial space. Moving toward one valuable point therefore tends to put you closer to others.

Second, agents search longer when they have strong beliefs in the existence of some valuable solution. They persist in the face of negative feedback that might otherwise end search (Fleming & Sorenson, 2004; Gavetti & Levinthal, 2000). Incremental moves amplify these benefits by forcing the agent to explore even more of search space. More broadly, almost anything that lengthens the time spent searching increases the probability of a serendipitous discovery, something that we call the *persistence principle*.

Our model also reveals at least two other important insights. First, theories can also stymie serendipitous discovery. Those with more extensive beliefs, with less accurate beliefs, and with positively biased beliefs frequently ignore unexpected discoveries because they believe that they might still find something better. As they experiment, they correct their errant beliefs. But, with finite time, they may never return to the fortuitous discoveries that they made.

Second, serendipity in terms of discovering something that provides a solution to something that the agent had not been trying to solve, may represent a distraction. The model allows us to establish a counterfactual: Would the agents have found better solutions to their original problems if they had not switched to a serendipitous discovery that solves a different problem. We find evidence for a *serendipity trap*. In most cases, agents would have eventually found a higher fitness point as a solution to their primary problem had they not abandoned those searches to pursue their serendipitous discoveries.

SERENDIPITY

In 1856, William Henry Perkin an 18-year-old chemistry student had been trying to synthesize quinine, an anti-malarial medication (Travis, 1990). It was not an easy problem. Quinine requires not just the right chemical ingredients but also a specific three-dimensional structure (Cova et al., 2017). Perkin failed multiple times. However, one of his failures produced a black sludge that became a brilliant purple when combined with ethanol (Travis, 1990). The color caught his attention. Purple was uncommon, reserved for the rich and royalty. Its main source at the time had been a dye derived from a rare sea snail (Cooksey, 2013). His accidental discovery, analine purple, went on to become the first synthetic dye, launching the artificial chemical industry in Germany (Travis, 1990).

Perkin's story captures many of the common elements of what has been labelled as serendipity. It involves both a process and an outcome (Busch, 2024). Serendipitous discoveries emerge from a search process. Perkin, for example, had been trying to create quinine. They involve an element of surprise, finding something unexpected. In this case, he had not even been searching for a dye. That accidental discovery, moreover, turned out to have value.

We define serendipity as finding something valuable while searching for something else. That definition encompasses two somewhat distinct types of serendipity.[1] In one, people searching

[1] Yaqub (2018) splits serendipity into four types, distinguishing also searching for a solution to a specific problem

for a solution to one problem find a solution to a different one (e.g., Dew, 2009; Garud et al., 2018; Hippel & Krogh, 2016). In another, people find a solution to the problem that they had been trying to solve but they do so through an unanticipated path (e.g., Merton & Barber, 2004).

Scholars across a range of areas in management have increasingly been interested in serendipity (for a recent review, see Busch, 2024). In competitive environments, serendipity may allow some firms to discover and exploit valuable resources and strategies (e.g., Denrell et al., 2003). In innovation, serendipity appears to have been a factor in many important inventions and scientific discoveries (e.g., Busch, 2024; Busch & Barkema, 2022; Merton & Barber, 2004). In social relationships, it may play a surprisingly important role in who becomes friends with whom (Feld, 1981; Festinger et al., 1950; Sorenson, 2023)

Despite this interest in the phenomenon, studying serendipity has been difficult for two reasons. First, serendipitous discoveries appear relatively rare (Denrell et al., 2003; de Rond, 2014). People find themselves surprised on a regular basis, encountering something unexpected. But most of these surprises end up being missed or dismissed. They do not seem interesting enough to pursue or they do not result in valuable discoveries.

Second, the fact that serendipity involves both surprise and the identification of something novel and useful means that researchers can only identify it in retrospect. As a consequence, research on serendipity effectively selects on the dependent variable. Researchers document and record examples of serendipity, attempting to infer the common elements from those cases (Busch, 2024; Yaqub, 2018). But we have little sense of all of the activities that did *not* lead to a fortuitous discovery. Perhaps both the successful and unsuccessful had engaged in the same search strategies but only some got lucky.

Given this fundamental difficulty in studying serendipity, we see substantial value in developing a formal model of the processes underlying it. Our model omits many of the nuances that have been discussed in the literature. Yet, we believe that both the process of translating the concept into a model and the analysis of the model have value (Adner et al., 2009). The former

from searching without a clear goal. In the model, we will treat undirected search as a baseline case.

requires us to be precise about our assumptions. The latter reveals some important intuitions about the sources of serendipitous discoveries.

MODELING SERENDIPITOUS SEARCH

To model search and serendipity, we extend the NK model, which depicts a search space as being delineated by N elements with a level of interdependence defined by K. Kauffman (1993) first introduced the model as a means of understanding evolutionary processes in biology. Management scholars, however, quickly adopted and adapted it to study search and innovation processes (e.g., Fleming & Sorenson, 2001; Levinthal, 1997; Rivkin, 2000). Its explicit incorporation of interactions both aligned well with the notion that elements of inventions and strategies might act as complements or substitutes and allowed researchers to assess the success of search strategies in more and less challenging environments.

In adapting the model to management, many modifications have been made. Some, for example, have incorporated the idea of cognition, of having mental representations of the space that allow for sophisticated forms of search (e.g., Csaszar & Levinthal, 2016; Fang et al., 2025; Gavetti & Levinthal, 2000; Singell, 2025). Others have overlaid organizational and technological structures onto the model that decompose search into sub-problems (Ethiraj & Levinthal, 2004; Fang & Kim, 2018; Rivkin & Siggelkow, 2003). Yet others have modified the interaction structure to allow for more and less central components (Rivkin & Siggelkow, 2007).

Because the NK model has been used extensively in the management literature, we provide a condensed description of it. For those interested in an extended treatment, Levinthal (1997) and Rivkin (2000) provide excellent introductions to the baseline model.

The NK model represents combinations of $N \in \mathbb{N}$ binary components, indexed by i, with $K \in \{0, 1, \ldots, N-1\}$ interactions, as a fitness landscape. Agents traverse these landscapes by recombining components in a search for useful combinations.

Each element i corresponds to an important component of an invention or strategy, such as a routine, a resource, or a physical feature. In the case of analine purple, it might include both the antecedent chemicals and processes for combining them. Each state $s \in \mathcal{S}$, or location on the landscape, then corresponds to a specific string of N binary choices: $s = (s_1, \ldots, s_N)$.[2]

The fitness $f(s)$, or value, associated with any combination s comes from averaging across the component values for each element $i \in s$:

$$f(s) = \frac{1}{N} \sum_{i=1}^{N} w_i\big(s_i,\ s_{C_i}\big).$$

Those component values, $w_i\big(s_i,\ s_{C_i}\big)$, come from assigning a random draw from the unit uniform distribution to each possible combination of i, plus the K randomly selected elements with which it interacts (C_i).[3] Because this fitness averages values from 0 to 1, it also ranges from 0 to 1.

At one extreme, when $K = 0$, each element i contributes independently to fitness, leading to a smooth landscape with a single optimal state. At the other extreme, $K = N - 1$, every choice depends on every other choice. The resulting landscape ends up being packed with peaks and valleys, often adjacent to one another.

Secondary landscapes

One type of serendipity involves agents finding a solution to a problem that they had not been trying to solve (Dew, 2009; Garud et al., 2018; Hippel & Krogh, 2016; Yaqub, 2018). That possibility falls outside the realm of the standard NK model because agents have been seen as searching for something valuable but not for a solution to a particular problem. Researchers have typically interpreted the fitness associated with a particular point on the landscape $f(s)$, for a combination of components s, as an overall usefulness—the value associated with the best use of those components (e.g., Fleming & Sorenson, 2001; Levinthal, 1997). In the case of Perkin, the failed attempt to synthesize an antimalarial would have had high fitness because it produced a valuable dye.

[2]The state space includes $\mathcal{S} = 2^N$ possible combinations. In our simulations with $N = 15$, the state space has 32,768 combinations.

[3]For each $i \in \{1, \ldots, N\}$, randomly select $C_i \subseteq \{1, \ldots, N\} \setminus \{i\}$ with $|C_i| = K$.

To address this issue, we recast the concept of fitness as being specific to solving a particular problem or use case u. We therefore index fitness by its use: $f_u(s)$. Each set of N components has a multitude of possible uses. The chemicals combined by Perkin, for example, would have separate fitness values as an antimalarial and as a dye. Although a set of components could have many potential uses, we simplify the model by limiting it to two, a primary landscape ($u = 1$) on which agents search, and a secondary one ($u = 2$) on which they might also find a serendipitous discovery.[4]

Combinations of components that provide good solutions to one problem need not do so for others. We treat the w_i as being specific to a use case, drawing independent values to generate each landscape. Components need not even interact in the same ways in providing potential solutions to different problems. In other words, different landscapes drawing on the same N components could have different values of K.

We assume that the agents can observe the fitness values on the secondary landscape for the states that they visit. That effectively imposes a degree of "openness" as a scope condition for serendipity. The agent must at least recognize that the other potential use case exists (e.g., Busch & Barkema, 2022; Hippel & Krogh, 2016). We, however, do not see this assumption as overly constraining. Openness has both an extensive and an intensive margin. We see awareness of the alternate use as an extensive margin. But agents can still dismiss those discoveries if they do not exceed some threshold of usefulness, an intensive margin. Indeed, variation in this threshold accounts for several of our results.

Beliefs and expectations

The very notion of surprise as an aspect of serendipity suggests that those searching had expected something (different). Perkin believed that the chemicals that he had been combining might produce quinine. Our model therefore endows agents with beliefs that can guide their search.

NK models have most commonly incorporated beliefs by giving the agent a cognitive map

[4] Because we could arbitrarily assign any use u to being the secondary landscape, this simplification does not affect the generality of our results.

or representation of the search space, one less detailed than the actual space (e.g., Csaszar & Levinthal, 2016; Fang et al., 2025; Gavetti & Levinthal, 2000; Winter et al., 2007). Although this approach has generated many important insights about the value of experience in search, for our purposes it has the disadvantage of typically being applied to the entire space. It leaves limited room for the unknown.

We therefore adopt an alternative approach, one inspired by the emerging theory-based view (e.g., Felin & Zenger, 2009, 2017; Sorenson, 2024). Agents have initial beliefs, potentially guided by causal theories, about whether or not specific states, combinations of components, might provide effective solutions to the primary problem (e.g., Fleming & Sorenson, 2004; Rivkin, 2000; Singell, 2025).[5] But they do not have beliefs about the entire space. They may even remain unaware of important elements (e.g., Bryan et al., 2022).

Agents have imperfect beliefs. Their beliefs about states can deviate from the true fitness values of those states in two ways. First, they may include noise, perhaps because agents do not understand or are unaware of all of the factors relevant to determining the fitness of a combination. Second, they may be biased, overestimating the true fitness of states.[6]

We denote the set of beliefs as $B_{u,t} \subseteq \mathcal{S}$, where u designates the use case and t indicates the time period. For each $s \in B_{1,t=0}$, the initial belief about the fitness of the state includes three components:

$$b_{1,t=0}(s) = \underbrace{f_1(s)}_{\text{true fitness}} + \underbrace{\mu}_{\text{optimism}} + \underbrace{\epsilon_s \sim \mathcal{N}(0, \sigma^2)}_{\text{noise}}.$$

We also include in the belief set, $B_{u,t}$, information that the agent acquires about the actual fitness of states, both on the primary and secondary landscapes. Agents initially possess actual information only on the fitness of their starting position and of locations adjacent to it.[7] However, as agents learn the true fitness of a state, by visiting it, that information becomes part of their

[5] We could allow agents to have beliefs about the secondary landscape as well but, because their search decisions depend on trying to find a solution to the primary problem, any such beliefs would not alter our findings.

[6] Theoretically, agents could also be pessimistic about a state, undervaluing it. Practically, however, the agent would then never visit these locations (e.g., Denrell & March, 2001).

[7] In NK modeling, it has been standard to grant agents awareness of the fitness values for all states that deviate from their current position by a single element.

belief set. If they already had a prior belief about its fitness (on the primary landscape), the true value replaces it: $b_{1,t}(s) = f_1(s)$.

With those details in place, we can offer a formal definition of serendipity. Our definition builds on the intuition that serendipity involves the discovery of something valuable for which the agent had not been searching.

Definition 1 *We define a* serendipitous discovery *as any state $s \in \mathcal{S}$ not in the original belief set $s \notin B_{1,t=0}$ that has a fitness value on either landscape equal to the maximum value of the current, updated belief set:* $\max(f_1(s), f_2(s)) = \max B_{u,t}$.

Discoveries can meet this definition in two ways. On the one hand, the agent could find an attractive solution to their primary problem but through an unexpected path, based on a set of elements about which it had not had priors ($s \notin B_{1,t=0}$). On the other hand, an agent might find a solution to a problem that it had not been trying to solve. If the fitness of a state on the secondary landscape exceeds that of any state explored on the primary landscape as well as the beliefs that the agent has about states not yet explored, the agent adopts that serendipitous solution.

Notice that our definition allows for backtracking (Weitzman, 1978). If an agent had believed that some state offered a good solution but then discovers that it does not upon visiting the location, the updating of $B_{u,t}$ effectively allows the agent to return to the best actual solution, the highest fitness value, seen up to that point in time on either landscape.

Although we define the serendipitous discovery as a state having a fitness value equal to the maximum value of the current belief set, search does not stop there. We allow agents to engage in hill climbing from that point. Agents will therefore usually locate an even better solution before discontinuing their search.

Search processes

In the original version of the NK model, agents could search either through a myopic "walk" uphill or by jumping to some location on the landscape, by randomly changing two or more elements simultaneously (Kauffman, 1993). Walks involve the agent observing the fitness of local neighbors, states that involve changing only a single element, and then moving to the adjacent one with the highest fitness. Walks end when none of the adjacent combinations offer higher fitness (i.e. when the agent reaches a local maximum).[8]

With beliefs, search processes become more interesting. Agents can direct their searches in ways that they expect to improve their odds of finding valuable solutions (Csaszar & Levinthal, 2016; Gavetti & Levinthal, 2000; Singell, 2025). We separate the search process into two choices: the first involves selecting a target, the second a method for traveling there.

Target selection

We assume that the agent attempts to find high fitness states by testing their initial beliefs, their hypotheses about valuable states. Each period t, the agent chooses as a target the state s believed to have the highest fitness ($\max b_{1,t-1}(s)$). Agents treat both initial beliefs and beliefs that have been acquired through experience as equal. They therefore only shift to a new target state if one of their initial beliefs turned out to be false.

As a baseline for comparison, some of our simulations also assign a target state s at random. We see that baseline as similar to not having any beliefs about valuable points on the space.[9] Though almost certainly less common than directed search, Yaqub (2018) offers multiple cases of serendipitous discovery that emerged from apparently undirected exploration (see his discussion of "Bushian" and "Stephanian" serendipity). Having chosen a target, agents can move toward it by either walking or jumping.

[8] Standard implementations of the NK model effectively assume that agents have awareness of all N elements, allowing them to experiment with them. Although that assumption seems somewhat in tension with the idea that serendipity might spring from the unknown, versions of our simulations in which we do not allow agents to hill climb produce qualitatively similar results. Hill climbing primarily influences the value of serendipitous discoveries rather than their frequency.

[9] With a random target, even though beliefs do not influence the search direction, they still define a threshold for when the agent considers an unexpected discovery sufficiently high fitness to stop exploring.

Walking

In walking, the agent changes one element at a time. Walking reflects a process of systematic experimentation, isolating the effects of each change. Agents only consider as possible changes elements on which their current state differs from their target state. The agent randomly selects one of the possible one-step moves that would bring it closer to its target. Agents repeat the process until reaching the target state (or a serendipitous discovery on the way there).[10]

Jumping

In jumping, the agent changes multiple elements at a time. Upon arrival, the agent learns the true fitness of the destination state, s, and its one-step neighbors.[11] If none of the states offers a higher fitness than the agent believes the next most valuable state to have, it will jump toward that next belief. We allow jumping agents to mix jumping with walking, so if the target sits adjacent to the current state, the agent can walk to it. The agent continues until it exhausts its beliefs or until it finds a serendipitous discovery.

Some research on NK models has systematically varied the "length" of the jump, the number of elements being changed, to see how that search choice influences success (e.g., Kauffman, 1993). However, similar to Rivkin (2000), we only consider cases in which the agent attempts to jump directly to a target. Jump length varies, but the length depends only on the distance, the number of differing elements, between the current state and the target. We also follow Rivkin (2000) in assuming that these jumps occur with error. Each change succeeds with probability p. The agent therefore may not arrive at the intended destination. Longer jumps include more errors on average: When an agent attempts $J > 1$ simultaneous changes, it arrives at

[10] This procedure deviates somewhat from the standard approach to myopic hill-climbing on NK landscapes (e.g., Kauffman, 1993; Levinthal, 1997). In the standard approach, agents select the neighboring state that offers the highest fitness. That approach, however, essentially assumes that agents can evaluate states before moving to them. Adopting the standard approach does not qualitatively change any of our results but it does increase the rate of serendipitous discovery.

[11] This approach also deviates somewhat from prior research. Levinthal (1997) and Rivkin (2000), for example, require jumpers to search their new neighborhoods after arrival. Our modification ensures that walking and jumping reveal similar amounts of information on each move. The walking agent gains information about no more than $N - 1$ adjacent states while a jumping agent can gain information about $N + 1$ states (the destination plus the adjacent states).

its intended destination with probability p^J, missing by an average distance of $J(1-p)$.

Simulation

Except for edge cases, the NK model eludes analytical evaluation (Rivkin, 2000). We therefore sampled the parameter space and explored the behavior of the model numerically.[12] Our simulations fix N at 15.[13] We explored the parameter space on K, the number of starting beliefs, the noisiness and amount of bias in those beliefs, and the accuracy of jump traversal. We allowed K to vary across its full range. However, because the behavior changes smoothly with small changes in K, our figures often report only four values to simplify the presentation of the results ($K \in \{2, 4, 10, 14\}$).

In the results reported below, the primary landscape and the secondary landscape have been generated using the same K. In unreported simulations, we have relaxed that assumption. Interestingly, the ease or difficulty of searching the secondary landscape does not matter to any of the results. That may come as a surprise to many. However, recall that search itself, where agents explore and whether they persist in their searches, depends only on beliefs about the primary landscape. Fitness values from the secondary landscape appear to the agent as random draws.

Agents could have from 5 to 50 starting beliefs ($|B_{1,t=0}| \in \{5, 10, 20, 30, 40, 50\}$). We varied the noise surrounding those beliefs from 0.1 to 0.5 ($\sigma \in \{0.1, 0.2, 0.3, 0.4, 0.5\}$) and the bias from 0 to 0.5 ($\mu \in \{0.0, 0.1, 0.2, 0.3, 0.4, 0.5\}$).[14] We allowed jump accuracy to range from 80% to 99% ($p \in \{0.8, 0.9, 0.95, 0.99\}$). In combination, we sample 10,800 points in the parameter space.

We simulate each set of parameters 1,000 times. For each individual iteration, we first generate a landscape with randomly-drawn interactions and fitness values. We then draw a set of

[12] LLM Disclosure: We used two LLMs (ChatGPT 5-Thinking and Claude Opus 4.1) when coding our simulations. We first coded the model without LLM assistance. We then used LLMs to speed up the simulation, to generate richer logs, and to test a range of alternative search behaviors. We validated the output of these enhanced simulations by comparing them to our initial hand-coded results on a portion of the parameter space.

[13] Although researchers sometimes treat N as a variable of interest, the research that has done so has found that nearly all of the interesting variation in the behavior of the NK model stems from K (e.g., Kauffman, 1993).

[14] High levels of noise and optimism can push beliefs to exceed one.

starting beliefs and a starting position. For each search method, we simulate one agent per landscape and record their per-period and final fitness, number of serendipitous discoveries, number of periods before termination, and whether or not the agent discovered the global maximum. We average the results across the 1,000 agent-landscape iterations.

We adopt a short, fixed horizon T for our headline comparisons ($T = 40$). In this model, most interesting differences occur in the early periods. To reflect the fact that fully random agents do not have awareness of the impending end of the simulation, we measure fitness performance using the best visited state up to the time horizon.

SERENDIPITOUS OUTCOMES

What does the model tell us about serendipity? In most of the results, we focus on the share of simulation runs that result in a serendipitous discovery. Our measure counts the number of serendipitous discoveries, dividing them by the total number of simulation runs. In other words, it reports the proportion of simulated searches from a set of parameters that resulted in a serendipitous discovery.

But we also sometimes assess the fitness associated with serendipitous discoveries. Even in the absence of serendipitous discovery, some search behaviors yield better outcomes on NK landscapes. We therefore calculate these values as being in excess of the next best non-serendipitous discovery, the marginal value of serendipity.

Figure 1 first reports the frequency of serendipitous discoveries as a function of K. We split serendipitous discoveries into two types: those on the primary landscape reflect cases in which the agent finds a valuable solution to the problem that it had been trying to solve but where that solution comes from a state of the search space about which the agent did not have any prior beliefs (i.e. from a place where it had not been expecting a solution). Those on the secondary landscape, meanwhile, capture instances of finding a valuable solution to a problem that the agent had not been trying to solve. The figure depicts two types of search. Agents can walk or jump toward the state that they believe offers the highest fitness.

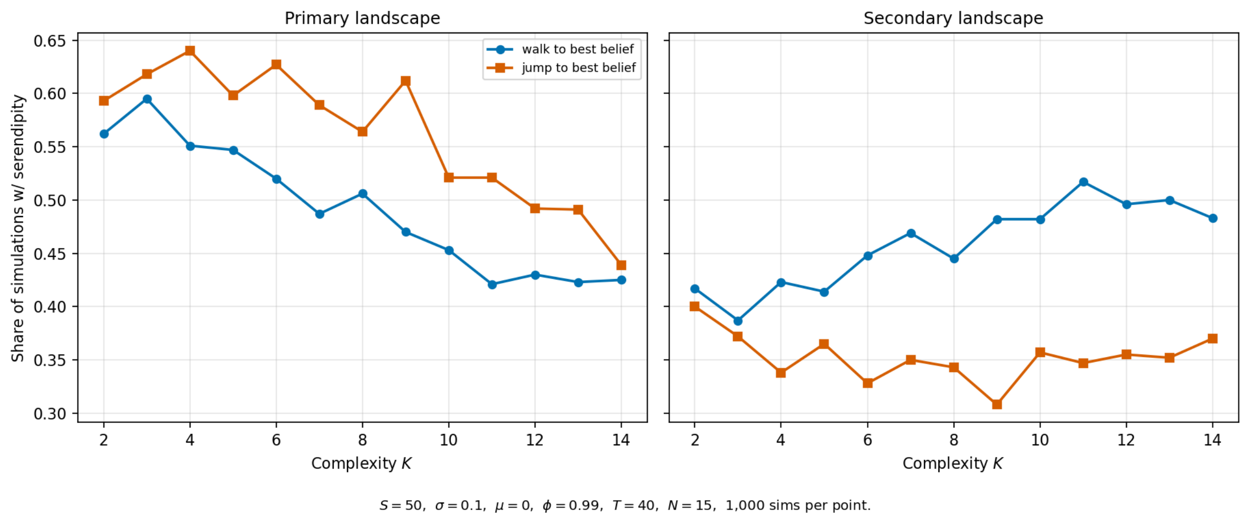

Figure 1: Serendipitous discoveries as a function of K.

Many may find the rates of serendipity reported here quite high. In some cases, more than 60% of simulation runs end in a serendipitous discovery on the primary landscape. Most of them stem from hill climbing. The agent travels to the state that it believes has the highest fitness but then, upon arriving, it discovers that a neighboring state, one that requires changing only one element, has even higher fitness. So, according to our definition, that incremental step represents a serendipitous discovery because it had not been expected.

Should we consider these improvements from hill climbing to be serendipitous discoveries? Most of them represent minor improvements in fitness. By contrast, the majority of examples of serendipity in the literature have been major innovations. The discovery of aniline purple launched an entire industry. Do these minor improvements and major discoveries represent the same phenomenon?

In this respect, we believe that the model surfaces an interesting issue. Unexpected discoveries actually appear quite common. In large combinatorial search spaces, even agents with deep expertise will rarely anticipate all of the factors that matter to solving a problem. They will often stumble upon an ingredient or a step that could improve the outcome. Indeed, almost every invention or innovation involves a measure of serendipity. But the modal discovery ends up being incremental.

But perhaps we want to understand the sources of major discoveries rather than minor improvements. Then the question becomes: Do the search processes that generate incremental improvements in fitness differ from those that discover major ones? In most cases, our model suggests that the answer is no. Raising the threshold for what we consider a serendipitous discovery reduces their frequency but it does not qualitatively change the relationships between search behaviors and the relative rates of serendipitous outcomes. So, whether we wish to consider all unexpected discoveries as serendipitous or only those that exceed some value threshold, the conclusions end up being largely the same. However, in a few cases that we will highlight below, the search process does influence the expected fitness of the serendipitous discovery.

Perhaps the most notable difference between the two types of serendipity comes from their differing relationships to K. When searching based on beliefs, discoveries on the primary landscape decline as a function of K while they remain constant or rise on the secondary landscape. Lower K landscapes have higher levels of spatial autocorrelation (Kauffman, 1993). In other words, high fitness locations tend to cluster in certain parts of the state space. Because beliefs (at least noisily) capture the true value of states, moving toward promising states leads agents to more fertile regions of the combinatorial space. As K increases, however, this autocorrelation declines, meaning that information about the locations of some valuable states provides less of an advantage in finding others.

On the secondary landscape, however, a different principle operates: We have assumed that fitness on the secondary landscape is uncorrelated with fitness on the primary one. Information on the location of valuable points on the primary landscape therefore tells the agent nothing about where valuable points might exist on the second one. But search on the primary landscape becomes harder as K rises. Agents search longer. As they explore more of the space, particularly when they walk between points, their odds of finding a high-value solution to an unexpected problem rises.

Those patterns have an interesting implication: Serendipitous discoveries that involve unexpected solutions to the primary problem should occur relatively more frequently for those trying to solve "easy" problems, problems that involve less interdependence between the elements. By contrast, researchers and others engaged in trying to solve really difficult problems—those with high K—end up being relatively more likely to stumble upon a solution to something else.

Expertise

Many of the examples of serendipitous discoveries in the literature involve scientists (e.g., Merton & Barber, 2004; Yaqub, 2018). In attempting to solve some problem, many a scientist has either accidentally uncovered an unexpected solution to that problem or stumbled upon something that solves another problem. Aniline purple, aspartame, and penicillin, for example, all follow similar story lines.

That common theme suggests that something about the way that scientists approach search may increase their odds of serendipitous discoveries. We believe that at least three features of our model capture aspects of how scientists differ from others in their search processes: number of beliefs, belief noise, and search method.

Most directly, scientists usually have theories guiding them (Kuhn, 1962; Popper, 1959). They have hypotheses (beliefs) about what might solve the problem. Those theories potentially influence search in at least three ways. First, they direct the scientist to specific regions of the search space (Dosi, 1982; Fleming & Sorenson, 2004; Kuhn, 1962). In our model, beliefs have at least some basis in truth (subject to noise and optimism). Pursuing beliefs therefore will tend to direct agents toward better basins of attraction.

Second, theories allow the scientist to persist in the face of negative feedback (Fleming & Sorenson, 2004; Lakatos, 1970). Most experiments fail, or at least fail to produce something of significant value. Without a theory, a myopic searcher might shift directions, assuming that a path represents a dead end. But those with confidence that the solution might still reside in that direction often persist in their pursuit of it. Even though our agents update their beliefs when they reach a target, beliefs in our model still allow agents to traverse across valleys that might otherwise trap them in local optima on their paths to these targets (cf. Winter et al., 2007).[15]

Third, theories establish a benchmark. When explorers encounter something unexpected, how can they decide whether that discovery seems promising enough to pursue? Scientists again benefit from having theories (beliefs). Their expectations about the value of a set of possible solutions provides a basis for deciding whether they should shift their search to an unexpected region of the state space or to another application area. In the model, agents only shift to the secondary landscape or to *terra incognita* on the primary landscape when their discovery exceeds the value of what they expect in other places.

Scientists differ from other explorers not only in being guided by theories but also in their methods of searching the space. In their training, scientists learn to isolate the effects of individual elements. When experimenting, they will typically try to vary only one factor, holding everything else constant (e.g., Camuffo et al., 2020). In the model, that method feels analogous to walking, to changing one element at a time. By contrast, those without that training may attempt to jump directly to what they see as the solution.

To explore these mechanisms, we begin by examining the frequency and value of serendipity as a function of the depth of the agent knowledge (number of beliefs) and the precision of those beliefs (noise).

Belief depth and precision

Scientists potentially differ in their theories in two ways from non-scientists: (i) the depth of their understanding, and (ii) the predictive accuracy of their theories.[16] The model implements these two characteristics in separate ways. We interpret the number of initial beliefs held by the agent as being an indicator of depth of understanding. Noise, or the lack of it, provides variation in the predictive accuracy of those beliefs.

[15] Although we have chosen to model theories as beliefs about the value of specific states, if agents instead had lower dimensional representations (maps) of the space (e.g., Gavetti & Levinthal, 2000), an interesting extention of the model might involve allowing agents to update their representations but with varying strengths of prior beliefs in those maps.

[16] Theories may also differ on a number of dimensions not included in our model, such as their logical consistency and the extent of their elaboration (Sorenson, 2024).

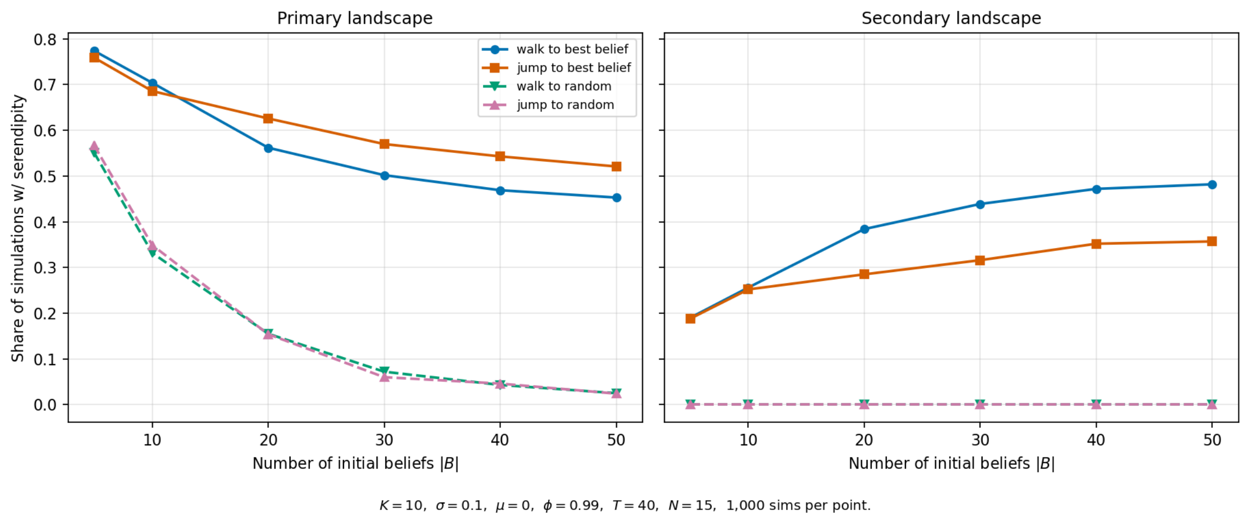

Figure 2: Serendipitous discoveries as a function of $|B|$.

Figure 2 reports serendipity frequency as a function of the number of beliefs held by the agent. As a benchmark, these figures include runs where the agent has beliefs but does not use them to guide its search. Beliefs in those runs serve only to establish a benchmark for the agent to determine when it should discontinue its search.

Being guided by beliefs, a theory, dramatically increases the odds of a serendipitous discovery. Consider first the primary landscape. Beliefs actually have two competing effects: On the one hand, because of the spatial autocorrelation in the landscape, they direct agents to regions of the combinatorial space that have more high-value states. Compare the solid lines to the dotted ones. Those pursuing targets they believe have value encounter serendipitous solutions at much higher rates than those visiting random locations.

On the other hand, beliefs raise the bar for discontinuing search (Weitzman, 1978).[17] Beyond a small number of beliefs, the discovery frequency *declines* with further increases in the numbers of beliefs. That effect stems from two closely-related factors. First, those with a deeper understanding of the state space more frequently have awareness of a possible solution to the primary problem. Many valuable discoveries, in other words, end up not being serendipitous

[17] Weitzman (1978) studies optimal stopping rules given independent draws but the structure of the NK model creates spatial autocorrelation in fitness values. Our definition of serendipity, moreover, represents a behavioral rule rather than an optimal policy.

(because they had been expected). Moreover, when those with a deep understanding do come across valuable states of which they had not been aware, they more frequently dismiss them, believing that they could find something better if they kept searching.

On the secondary landscape, this raising of the bar has an unexpected benefit. It keeps the agents searching longer (Weitzman, 1978).[18] They explore more of the space, meaning that they more frequently find serendipitous solutions to problems that they had not been trying to solve.

Although our model depicts the agent as an individual, we might also think of it as a team. In that case, the number of initial beliefs would probably increase as a function of team size. Larger teams then would less frequently find serendipitous discoveries to their primary problems. As more researchers with varied beliefs join the group, it becomes harder to build consensus around the idea that some unexpected state or solution might be the best option. However, their inability to agree might lead them to more serendipitous solutions to unexpected problems.

We next explore the effects of predictive accuracy. Figure 3 reveals that serendipitous discovery on the primary landscape declines monotonically with noise (inaccuracy). But noise increases discovery on the secondary landscape.

Noise in beliefs about the value of states has countervailing effects. On the one hand, noise extends search. It encourages the agent to explore regions of the state space that it might not have otherwise. The state with the highest true fitness may appear to have limited value. One with moderate fitness, on the other hand, may appear attractive. That circuitous search gives the agent more opportunities to encounter unexpected high-fitness states.

But noise has two other effects that reduce the probability of serendipitous discovery. First, noise, by definition, reduces the correlation between the agent's beliefs about the value of a state and its true fitness. Much of the value of these beliefs in serendipitous discovery on the primary landscape comes from spatial autocorrelation, from attracting the agent to regions of the combinatorial space with more high-value states. Noise compromises this effect.

Second, noise raises the bar for a serendipitous discovery (Weitzman, 1978). If an agent

[18] Stopping on the secondary landscape ends up being a version of Pandora's Box

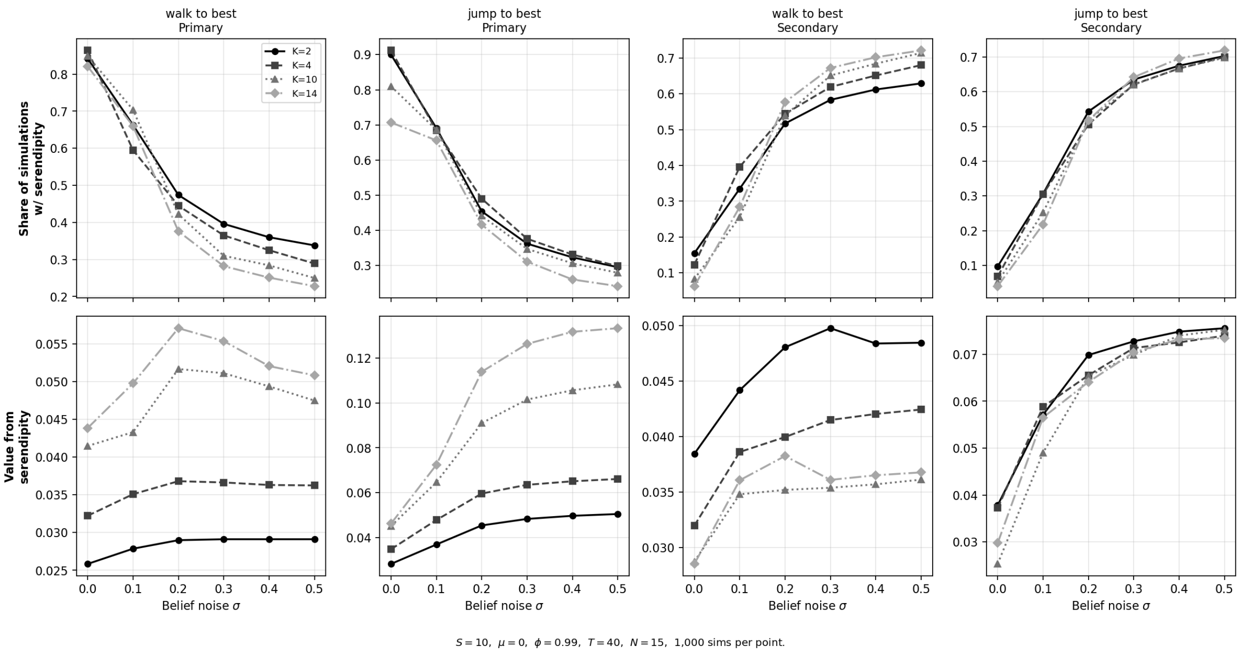

Figure 3: Serendipitous discoveries as a function of noise.

believes that some high value solution remains untested, it will pass over promising discoveries. Experience helps to mitigate this second effect. When an agent visits a state to which it has erroneously assigned a high value, it revises its belief to the true fitness. It may then return to earlier discoveries, potentially recognizing their value. But that process requires time.

On the primary landscape, the reduction in spatial autocorrelation, in the value of being pulled to valuable regions, dominates. The probability of serendipitous discovery declines rapidly with increasing noise. But the spatial autocorrelation on the primary landscape has no value on the secondary landscape. Without this negative effect, the value of searching longer leads those with noisier beliefs to more frequently arrive at high-value solutions to unexpected use cases.

Walking versus jumping

Scientists not only have deeper and more predictive theories but they also explore the space differently. They run controlled experiments, often trying to vary just a single element. If they believe that factors interact, they might implement a crossed design, testing all of the combinatorial combinations, to explore the interaction fully. In the mechanics of our model, they may walk rather than jump.

Revisiting Figure 1, we can see that jumping tends to generate more serendipitous solutions to the primary problem but that walking produces more discoveries on the secondary landscape. Walking differs from jumping in two important ways. First, it tends to keep the agent searching in the same region of the combinatorial space, for good and for bad. Second, it keeps the agent searching longer, on average. Because of the spatial autocorrelation, the first effect often puts the walking agent at a disadvantage relative to the jumper in finding serendipitous discoveries on the primary landscape. But the extended search time benefits the walking agent in terms of serendipity on the secondary landscape.

Jumping also interacts in an interesting way with belief noise. Notice that in Figure 3, on both the primary and the secondary landscapes, the value of the discoveries found *increase* with noise when the agent jumps to their target beliefs, particularly with higher levels of interdependence (K). That effect stems from the longer time spent searching. Because belief noise raises the bar, jumping agents often encounter many high-value states but they keep searching because they believe that they may still find something better. Eventually, however, the agent visits all of the states that it believed had high values. After it corrects its errant beliefs, it then returns to the highest value point that it had visited along the way, usually a serendipitous solution. With walking, that return does not happen because the agent has not had enough time to update all of its errant beliefs. Agents run out of time.

Of course, even when searchers attempt to move directly to states that they see as promising, they may fail to do so. They may have incorrect or incomplete beliefs about the elements required or they may make mistakes in trying to implement them. These appear as jump errors in the model (Rivkin, 2000). How does jump accuracy influence search outcomes?

Figure 4 plots the proportion of searches that result in a serendipitous discovery as a function of jump accuracy. On the primary landscape, the probability of a serendipitous discovery usually first rises and then falls with decreasing jump accuracy. On the one hand, jump errors lengthen search because the agent often needs multiple periods to reach its intended target. On the

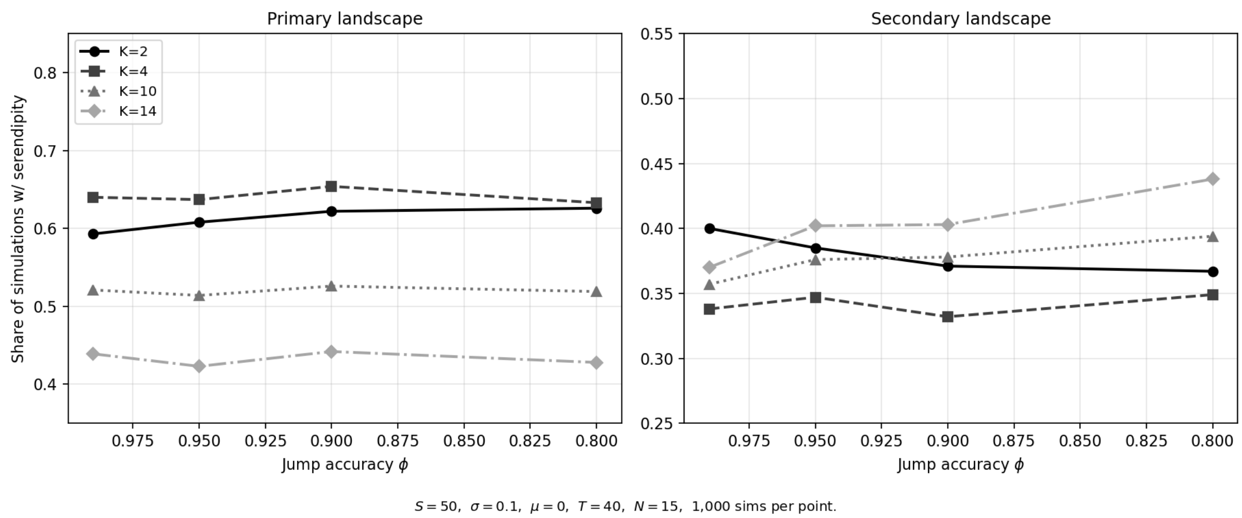

Figure 4: Serendipitous discoveries as a function of jump errors

other hand, errors have an effect similar to noise in the sense that they pull the agents into less productive regions of the combinatorial space (i.e. regions with fewer high-value states).

On the secondary landscape, the results depend on K. At low levels of interdependence, agents can easily correct jump errors by hill climbing. They therefore often discover either an expected or unexpected high point on the primary landscape, ending search. But as K increases, it becomes more difficult for the agent to correct their errors through hill climbing (Rivkin, 2000). They end up searching longer. As a result, they also find more interesting solutions to problems that they had not been trying to solve.

The persistence principle

Nearly all of our results on the secondary landscape follow what we like to think of as the *persistence principle*. Those who explore the largest number of states, the largest number of combinations of elements, enjoy the highest odds of serendipitous finds.

Figure 5 plots the probability of a serendipitous discovery on the secondary landscape as a function of the length of the search (pooling over all of the model parameters, except for jumping versus walking). Given the large number of runs, we bin the results into deciles of the search length distribution. The squares represent jumping searches and the circles walks. For both forms

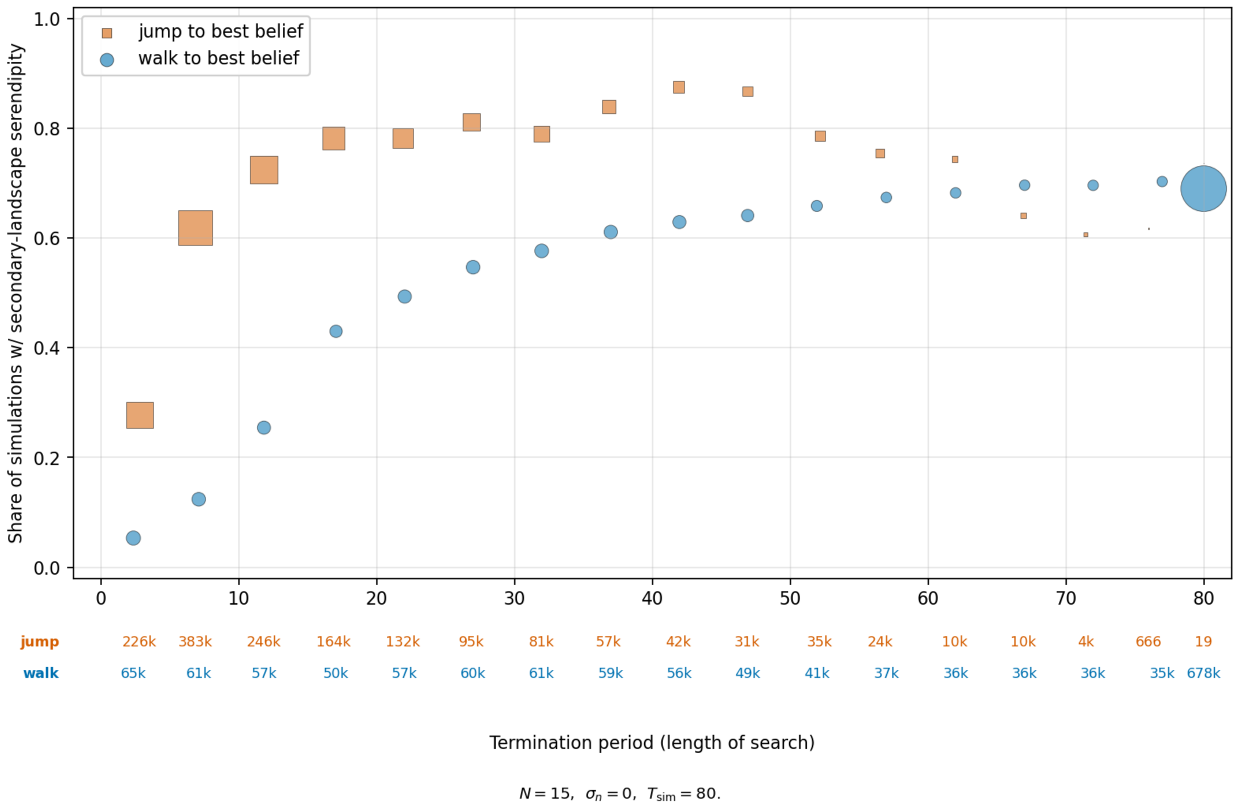

Figure 5: Serendipitous discoveries on the secondary landscape as a function of search length

of search, the length of the search correlates highly with the probability of a serendipitous discovery.

Search length correlates even more strongly with the probability of serendipitous discovery for jumping than it does for walking. That difference stems from the fact that jumping explores the space quickly. Recall that belief noise can stymie serendipitous discovery because the agent believes that better states may remain unexplored. Jumping allows the agent to test its beliefs about these states more rapidly, allowing it to update its errant beliefs before the simulation ends. Once the agent recognizes that these states do not offer the promise that had been expected, it becomes open to other discoveries.

Those familiar with the literature on NK models may feel that this result follows a well known characteristic of search on these rugged landscapes: Anything that pushes agents out of the catchment area for a local optima, whether noise, cognitive limits, or the decomposition of the problem into parts, improves overall search outcomes (Gavetti & Levinthal, 2000; Levinthal, 1997; Rivkin & Siggelkow, 2003).

However, that similarity is superficial. Notice that our definition of serendipity depends on the fitness being higher than that of any expected outcome on the primary landscape. That means that any search characteristics that improve outcomes on the primary landscape actually have a countervailing effect on the odds of discovery on the secondary landscape because they effectively raise the bar for being considered an interesting discovery. Factors that increase the time spent in search, such as walking relative to jumping, moreover often have no advantages for fitness on the primary landscape.

Perils of perfection

In Figure 3, the probability of serendipitous discovery on the primary landscape declines by more than half at high levels of noise. Part of the problem comes from the rising bar for considering a discovery good enough to stop search. The threshold, the max value of the belief set, represents an order statistic. As such, it increases with the standard deviation of the distribution.[19]

Both entrepreneurs and scientists have often been seen as dogmatic, relentlessly pursuing their beliefs about the right solution to a problem. The absence of noise does not quite capture that effect. But we can explore this confidence in the model by varying the level of optimism, the extent to which the agent overestimates the value of its preferred solutions (beliefs).

Figure 6 depicts the probability of finding a serendipitous solution on the primary landscape as a function of the degree of optimism. The panels vary the time horizon from 20 periods up to 80 periods. The effects of optimism interestingly interact with the time horizon. Because agents believe that something better remains out there, optimism keeps them searching. It raises their bar.[20] They correct these beliefs with experimentation. But with shorter time horizons, they often run out of time.

[19] $E[\max B] \approx \mu + \sigma\sqrt{2 \ln |B|}$. The max belief therefore rises roughly as a function of the square root of the log of the number of beliefs.

[20] Optimism has a similar, though even stronger, effects than noise in terms of raising the bar. On average, half of the noisy beliefs end up being overestimates whereas all of the optimistic ones are.

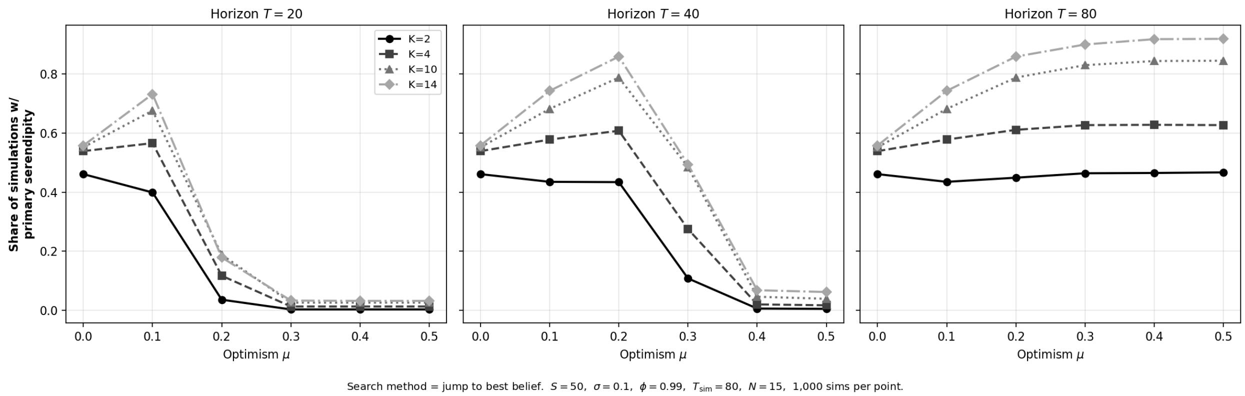

Figure 6: Given time, optimism improves the probability of a serendipitous discovery.

However, when they have enough time to correct their errant beliefs, having had a higher bar ends up being functional (Weitzman, 1978). As they search, agents often visit multiple high fitness states. They can then return to the best of them.

This raising the bar effect also interacts with the depth of the agent's understanding of the landscape (the number of beliefs). If the agent has errant beliefs about a large number of states, they need more time to explore these beliefs and to correct them. The agent may not have time. Deeper knowledge therefore raises the perils associated with errant beliefs, whether due to noise or optimism.

Seductions of serendipity

One of the advantages of having a model comes from being able to explore counterfactual scenarios. When a serendipitous discovery is the solution to a different problem, an interesting question arises about whether the agent did the right thing by switching over to the secondary landscape.[21] In other words, would the agent have been better off persisting in searching for a solution to its original problem? Should Perkin have continued his search for quinine?

We investigate this question by comparing the fitness reached by an agent on the secondary landscape, following a serendipitous discovery, to the fitness that same agent would have reached

[21]This question ends up being similar to a multi-armed bandit problem (Gittins, 1979). Should the agent keep pulling the arm for the primary problem or switch to the arm for the secondary one? However, our model differs from the canonical bandit set up in two critical ways. First, interdependence means that the main bandit delivers correlates draws. Second, the agent seeks to maximize its highest value rather than its average return.

had it continued with its primary search. Interestingly, and perhaps surprisingly, agents that shift to the secondary landscape almost always would have been better off had they continued to search on the primary landscape.[22] We call this tendency to shift to good but not great solutions, the *serendipity trap*.

To illustrate the issues, let us consider two cases. “Focused” agents have beliefs about a small number of states and they have relatively accurate beliefs about their value. “Diffuse” agents, by contrast, have a large number of beliefs but high levels of uncertainty about them. We impose a cost to pivoting (c). In order to “pivot” to the secondary landscape, the value of the discovery must exceed by c the value of any state on the primary landscape that the agent has either visited or believes may exist. Figure 7 reports the differences between the value of the serendipitous states that agents discovered on the secondary landscape and the value of the solution that they would have eventually found on the primary landscape had they not pivoted. Not only does the serendipitous discovery often represent an inferior outcome but the penalty actually increases initially in the costs of switching to the secondary landscape.

Why? Two factors combine to produce this effect. First, when agents find a promising point on the secondary landscape, that point has been selected at random. It might represent a region of the secondary landscape with many high-valued states, but often it does not. By contrast, the fact that beliefs (theories) reflect the true value of points on the primary landscape means that the states that the agent believes have the highest value reside close in the combinatorial space to other high-value states. If the agent keeps searching the space around those points, it will often find an even higher-valued solution to the primary problem.

But the value of the information in the autocorrelation does not explain the relationship between the switching cost and the magnitude of the serendipity penalty. That stems from a subtle effect of raising the bar. As the agent becomes more selective in switching, it more commonly ends up choosing a local optimum. When they impose a lower bar for shifting to the secondary

[22] Walking agents with diffuse beliefs do not suffer a penalty because they never shift to the secondary landscape. With walking, they run out of time before they can correct their errant beliefs. They therefore always select a state on the primary landscape as the highest fitness point in their final belief set.

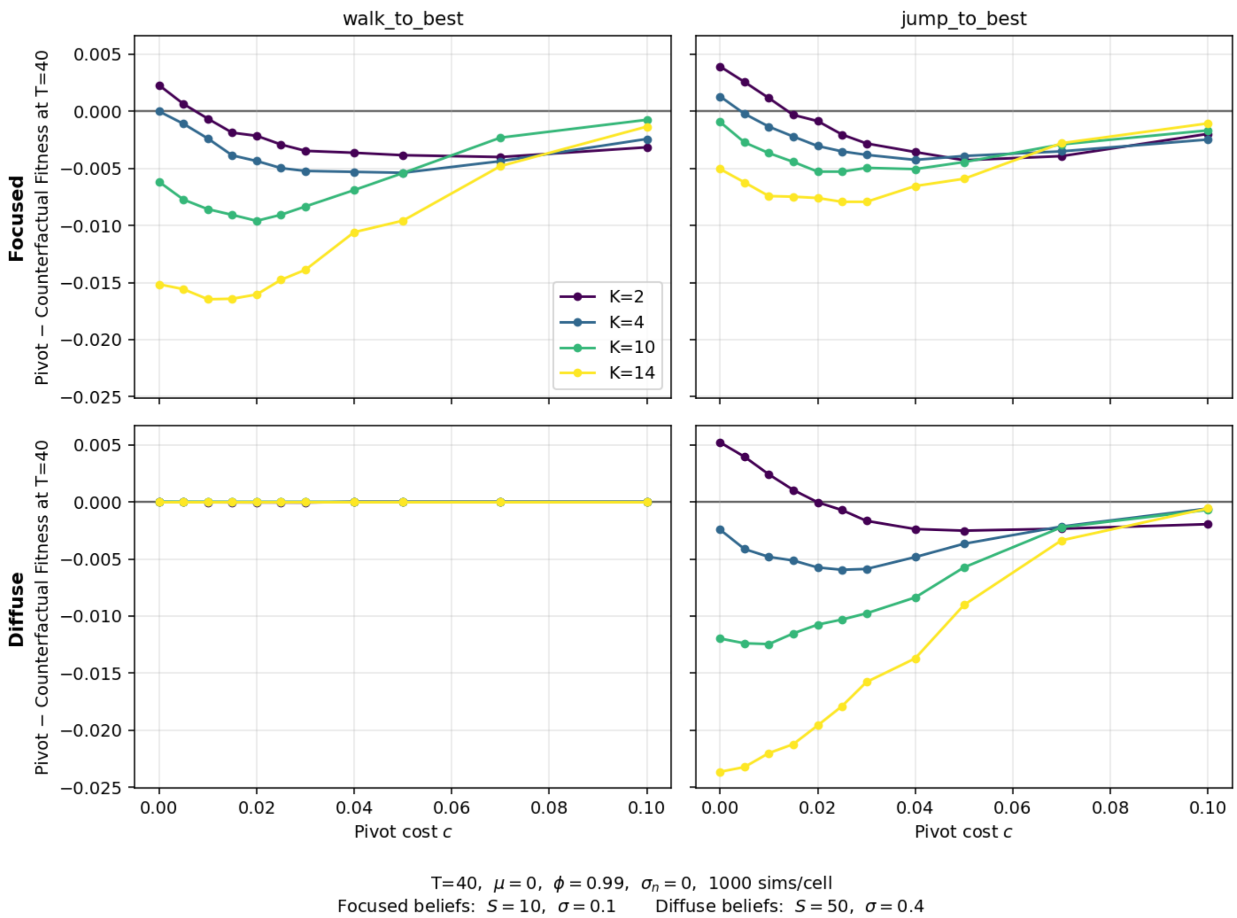

Figure 7: The value of the serendipitous discovery relative to foregone future discoveries

landscape, they often end up on a slope, a state that allows them to hill climb. Because higher fitness peaks have larger catchment areas (Kauffman, 1993), entering at a lower point counter-intuitively often leads them to better solutions to the alternate problem.

Would Perkin had been better off had he ignored his unexpected discovery of aniline purple? Probably not. In this case, the example calls attention to an implicit assumption in our model. By assuming that all use cases have the same scaling in terms of their fitness, our model effectively assumes away differences in factors such as market size and the intensity of competition. Antimalarial medications never became as large of a market as synthetic dyes. In practice, whether pivoting ends up being good or bad depends also on the potential profits associated with the alternative use case relative to the original one.

DISCUSSION

We develop a formal model of serendipitous discovery. We extend the standard NK model to associate a set of N components with both a primary and secondary landscape, representing different use cases. We allow agents to have theories that generate beliefs about high potential points on the primary landscape (e.g., Fleming & Sorenson, 2004; Singell, 2025). These beliefs, however, are incomplete and uncertain. Even on the primary landscape, they may be unaware of some components or may not understand how they might combine. Agents can traverse these landscapes either through systematic steps or by attempting to jump directly to promising locations.

Agents can find two types of serendipitous discoveries. On the one hand, they may find high value solutions to their primary problem of interest but using unexpected combinations of components, ones about which they had no expectations (Merton & Barber, 2004). On the other hand, they may stumble upon a high value solution to a secondary problem while searching for a solution to the primary problem (Dew, 2009; Garud et al., 2018; Hippel & Krogh, 2016).

Our simulation suggests several interesting propositions about who experiences serendipitous discoveries and why they enjoy those advantages. On the primary landscape, having a theory, having beliefs about the high value states increases the odds of serendipitous discovery. Because of the spatial autocorrelation in the fitness landscape, even when agents have incomplete and imperfect beliefs, their beliefs direct them to regions of the space with more high-value combinations.

Though our primary interest here has been serendipitous discovery, our results raise interesting questions also about abduction. In our current set up, agents only update their beliefs to the correct values. With abduction, these empirical surprises might lead agents to revise their theories, which would in turn lead them to develop beliefs about states not in their initial belief set. Those changing expectations might mean that many of the states identified as serendipitous on the primary landscape might end up being expected by the time the agent found them. Extending the model in that direction might also raise interesting questions about the relative value of theories with different structures (Sorenson, 2024).

Many of the findings regarding the probability of serendipitous discovery on the secondary landscape—for finding the solution to a problem for which the agent had not been trying to solve—meanwhile, fall under what we think of as the principle of persistence. Quickly arriving at a high value solution to the problem of interest ends the search process. Any features of the agent or environment that extend the search process provide the agent with more information about possible solutions to alternative problems (March, 1991).

Our opening anecdote about searching for a source illustrates this principle well. An electronic search might bring the person immediately to the title of interest. Search ends with little possibility of serendipity along the way.

But now imagine a physical search, visiting a library to find a book in its stacks. In searching for the book, the person can typically only get close without having to peruse titles. Dewey decimal numbers at the end of an aisle might narrow the range being searched. But eventually the person must start reading titles. Methodically stepping through these nearby titles often yields serendipitous discoveries. Why? Even in highly interdependent fitness spaces characterized by lots of peaks, the high value points exhibit spatial autocorrelation. In the case of the book search, the closer the books are shelved to the one of interest, the more closely related their subjects as well.

In the age of Large Language Models (LLMs), these serendipitous discoveries to unanticipated problems may become even less common. LLMs further reduce the time spent searching. Even with electronic search, the searcher might encounter some sources with similar titles or authors (depending on the field used to search). Even after finding the article or book, the person would typically read it, encountering not just the information of interest but also other ideas. When using an LLM, by contrast, the searcher often receives a direct answer to their question of interest without being exposed to almost any extraneous information.

More broadly, efficiency often ends up being the enemy of serendipity. Those who value discovery may therefore wish to preserve some friction in how they search. Interviewing customers, searching for sources in the library, and reading widely exposing people to the unexpected, raising the odds of serendipitous discovery.

An exception to this persistence principle, however, comes from the open-ended nature of search. In many cases, solutions are not binary. They are not yes or no. Instead, they represent better or worse solutions to a problem. That raises questions about stopping rules. At what point should they consider a solution good enough. Agents value discoveries in relation to what they already know. Novel theories, inventions, and products only supersede the old if they are seen as sufficiently valuable. As agents become too optimistic about the prospects of points not yet visited, they potentially ignore high value finds. They never stop searching.

We finally see evidence of a serendipity trap. Our model allows us to consider a type of counterfactual that would rarely get documented in the real world: What happens when the searcher ignores the serendipitous discovery, continuing in pursuit of their original problem?

Across a wide range of conditions, it appears that searchers would have been better off continuing with their original question. When they continue searching, they find higher fitness states on the primary landscape than they do on the secondary one. That effect, again, stems mostly from the autocorrelation in the primary landscape. Agents' beliefs put them in proximity to high value points. Shifting to the secondary landscape essentially discards the value of the knowledge that they have.

In the entrepreneurship literature, pivoting has typically been seen as a positive (e.g., Grimes, 2018). Entrepreneurs who can quickly turn to something else in the face of negative feedback avoid wasting resources on futile quests. More recent revisions have emphasized the importance of pivoting only in response to evidence that conflicts with their expectations (Camuffo et al., 2020; Kirtley & O'Mahoney, 2023). But even in those cases, our model suggests that entrepreneurs may pivot too early.

Can the manager, entrepreneur or scientist engineer serendipity? They obviously cannot create serendipity directly. But they can create conditions and build teams that may foster it. Perhaps surprisingly, managers should not try to pursue serendipity by injecting randomness into the research process. Wild bets and leaps into the unknown have low odds of success. Instead, the highest rates of serendipitous discovery accrue to those who pursue a plausible hypothesis through deliberate, incremental steps (cf. Camuffo et al., 2020). Even an imperfect theory directs search toward combinations that are more likely to have value. Systematic, one-factor-at-a-time experimentation, moreover, forces the searcher to traverse more of that space, exposing them to more unexpected possibilities along the way.

But managers should also recognize a tension in how aggressively to pursue any serendipitous discoveries. On the one hand, the same confidence and depth of expertise that sustain a search can blind an organization to the value of its own discoveries. The agents in our model frequently pass over valuable finds because they believe something better remains undiscovered. They often run out of time before returning to pursue what they had already found. When solving their problem of interest, managers may wish to try to "lower the bar" for acting on a promising but unanticipated result.

On the other hand, the serendipity trap offers a cautionary tale to managers. When considering serendipitous solutions to problems not being pursued, those searching might have been better off ignoring these discoveries. Because a guiding theory points toward clusters of high-value solutions, the return to staying the course tends to exceed that to chasing an opportunity that arrived from a single random draw. That does not necessarily mean that entrepreneurs, scientists, and managers should always avoid pivoting. It undoubtedly depends on the relative value of the different use cases. Our model assumes that all use cases have the same potential. In reality, applications often vary substantially in their potential demand and in the intensity of competition. In cases where the serendipitous discovery represents a larger opportunity, the pivot may still pay off.

Many of the other features that encourage serendipity point to selection rather that process as promising policies. Searchers typically do not choose the depth of their understanding, the accuracy of their theories, or their levels of optimism. But managers, investors, and others who allocate resources can potentially observe those attributes. Our results suggest that smaller teams and those who hold fewer and looser priors more often stumble onto and act on the unexpected. Inventors and scientists with less experience, for example, may end up being more open to serendipitous discoveries. Optimism can also help to sustain search in the face of disappointment (cf. Hmieleski & Baron, 2009; Malmendier & Tate, 2008).

This project emerged from a conversation about the Encyclopedia Britannica. Despite fond memories of alphabetized and illustrated entries, none of us owned a copy. Why would we? It's an anachronism. Digitization, search engines, and LLMs have made it fast, cheap, and easy to retrieve any piece of trivia. Wikipedia returns information on Shakespeare's Cymbeline without requiring the reader to flip through entries for Canada, cheetahs, and cyanide. We no long run the risk of overshooting and being sidetracked by entries on cytokines or czars. And yet, we all had the feeling that something had been lost.

REFERENCES

Adner, R., Polos, L., Ryall, M., & Sorenson, O. (2009). The case for formal theory. *Academy of Management Review*, *34*(2), 201–208.

Bryan, K. A., Ryall, M. D., & Schipper, B. C. (2022). Value capture in the face of known and unknown unknowns. *Strategy Science*, *7*, 157–189.

Busch, C. (2024). Towards a Theory of Serendipity: A Systematic Review and Conceptualization [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joms.12890]. *Journal of Management Studies*, *61*(3), 1110–1151. https://doi.org/10.1111/joms.12890

Busch, C., & Barkema, H. (2022). Planned Luck: How Incubators Can Facilitate Serendipity for Nascent Entrepreneurs Through Fostering Network Embeddedness [Publisher: SAGE Publications Inc]. *Entrepreneurship Theory and Practice*, *46*(4), 884–919. https://doi.org/10.1177/1042258720915798

Camuffo, A., Cordova, A., Gambardella, A., & Spina, C. (2020). A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial [Publisher: INFORMS]. *Management Science*, *66*(2), 564–586. https://doi.org/10.1287/mnsc.2018.3249

Cooksey, C. (2013). Tyrian purple: The first four thousand years. *Science Progress*, *96*(2), 171–186.

Cova, T. F. G. G., Pais, A. A. C. C., & de Melo, J. S. S. (2017). Reconstructing the historical synthesis of mauveine from perkin and caro: Procedure and details. *Scientific Reports*, *7*, 6806.

Csaszar, F. A., & Levinthal, D. A. (2016). Mental representation and the discovery of new strategies. *Strategic Management Journal*, *37*(10), 2031–2049.

Denrell, J., Fang, C., & Winter, S. G. (2003). The economics of strategic opportunity [_eprint: https://sms.onlinelibrary.wiley.com/doi/pdf/10.1002/smj.341]. *Strategic Management Journal*, *24*(10), 977–990. https://doi.org/10.1002/smj.341

Denrell, J., & March, J. G. (2001). Adaptation as Information Restriction: The Hot Stove Effect. *Organization Science*, *12*(5), 523–538. https://doi.org/10.1287/orsc.12.5.523.10092

de Rond, M. (2014). The structure of serendipity [Publisher: Routledge _eprint: https://doi.org/10.1080/14759551.2014.967451]. *Culture and Organization*, *20*(5), 342–358. https://doi.org/10.1080/14759551.2014.967451

Dew, N. (2009). Serendipity in Entrepreneurship [Publisher: SAGE Publications Ltd]. *Organization Studies*, *30*(7), 735–753. https://doi.org/10.1177/0170840609104815

Dosi, G. (1982). Technological paradigns and technological trajectories: A suggested interpretation of the determinants and direction of technological change. *Research Policy*, *11*(3), 147–162.

Ethiraj, S. K., & Levinthal, D. A. (2004). Modularity and innovation in complex systems. *Management science*, *50*(2), 159–173.

Fang, C., & Kim, J.-h. (2018). The power and limits of modularity: A replication and reconciliation. *Strategic Management Journal*, *39*(9), 2547–2565.

Fang, C., Kim, J.-h. (, & Yang, H. (2025). Learning in Temporally Complex Problems: The Role of External Knowledge [Publisher: INFORMS]. *Organization Science*, *36*(5), 1861–1880. https://doi.org/10.1287/orsc.2022.16469

Feld, S. L. (1981). The focused organization of social ties. *American Journal of Sociology*, *86*, 1015–1035.

Felin, T., & Zenger, T. R. (2009). Entrepreneurs as theorists: On the origins of collective beliefs and novel strategies [Publisher: John Wiley & Sons, Ltd]. *Strategic Entrepreneurship Journal*, *3*(2), 127–146. https://doi.org/10.1002/sej.67

Felin, T., & Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. *Strategy Science*, *2*(4), 258–271.

Festinger, L., Schacter, S., & Back, K. W. (1950). *Social pressure in informal groups*. Harper.

Fleming, L., & Sorenson, O. (2001). Technolgy as a complex adaptive system: Evidence from patent data. *Research Policy*, *30*, 1019–1039.

Fleming, L., & Sorenson, O. (2004). Science as a map in technological search. *Strategic Management Journal*, *25*, 909–928.

Garud, R., Gehman, J., & Giuliani, A. P. (2018). Serendipity Arrangements for Exapting Science-Based Innovations [Publisher: Academy of Management]. *Academy of Management Perspectives*, *32*(1), 125–140. Retrieved October 14, 2025, from https://www.jstor.org/stable/26528741

Gavetti, G., & Levinthal, D. A. (2000). Looking forward and looking backward: Cognitive and experiential search. *Administrative science quarterly*, *45*(1), 113–137.

Gittins, J. C. (1979). Bandit Processes and Dynamic Allocation Indices. *Journal of the Royal Statistical Society: Series B (Methodological)*, *41*(2), 148–164. https://doi.org/10.1111/j.2517-6161.1979.tb01068.x

Grimes, M. G. (2018). The pivot: How founders respond to feedback through idea and identity work. *Academy of Management Journal*, *61*(5).

Hippel, E. V., & Krogh, G. V. (2016). Crossroads–identifying viable 'need-solution' pairs: Problem solving without problem formulation. *Organization Science*, *27*(1), 207–221.

Hmieleski, K. M., & Baron, R. A. (2009). Entrepreneurs' Optimism And New Venture Performance: A Social Cognitive Perspective [Publisher: Academy of Management]. *Academy of Management Journal*, *52*(3), 473–488. https://doi.org/10.5465/amj.2009.41330755

Kauffman, S. A. (1993, June). *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press. https://doi.org/10.1093/oso/9780195079517.001.0001

Kirtley, J., & O'Mahoney, S. (2023). What is a pivot? explaining when and how entrepreneurial firms decide to make strategic change and pivot. *Strategic Management Journal*, *44*(1), 197–230.

Kuhn, T. (1962). *The structure of scientific revolutions*. University of Chicago Press.

Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In I. Lakatos & A. Musgrave (Eds.), *Criticism and the growth of knowledge*. Cambridge University Press.

Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management science*, *43*(7), 934–950.

Malmendier, U., & Tate, G. (2008). Who makes acquisitions? CEO overconfidence and the market's reaction. *Journal of Financial Economics*, *89*(1), 20–43. https://doi.org/10.1016/j.jfineco.2007.07.002

March, J. G. (1991). Exploration and Exploitation in Organizational Learning [Publisher: INFORMS]. *Organization Science*, *2*(1), 71–87. Retrieved March 3, 2025, from https://www.jstor.org/stable/2634940

Merton, R. K., & Barber, E. (2004). *The travels and adventures of serendipity: A study in sociological semantics and the sociology of science*. Princeton University Press.

Popper, K. (1959). *The logic of scientific discovery*. Hutchinson.

Rivkin, J. W. (2000). Imitation of Complex Strategies. *Management Science*, *46*(6), 824–844. https://doi.org/10.1287/mnsc.46.6.824.11940

Rivkin, J. W., & Siggelkow, N. (2003). Balancing search and stability: Interdependencies among elements in organizational design. *Management Science*, *49*(3), 290–311.

Rivkin, J. W., & Siggelkow, N. (2007). Patterned interactions in complex systems: Implications for exploration. *Management Science*, *53*(7), 1068–1085.

Singell, M. (2025). *Any old theory will do: Why cause-and-effect performance links form parsimonious mental models of complex strategic environments* [Working paper, Stanford University].

Sorenson, O. (2023). Does diversity influence innovation and economic growth? it depends on spatial scale. *Research in Organizational Behavior*, *43*, 100190.

Sorenson, O. (2024). Theory, search, and learning. *Strategy Science*, *9*, 372–381.

Travis, A. S. (1990). Perkin's mauve: Ancestor of the organic chemical industry. *Technology and Culture*, *31*(1), 51–82.

Weitzman, M. (1978). *Optimal Search for the Best Alternative* [Google-Books-ID: 5CViWSuPJ1MC]. Department of Energy.

Winter, S., Cattani, G., & Dorsch, A. (2007). The value of moderate obsession: Insights from a new model of organizational search. *Organization Science*.

Yaqub, O. (2018). Serendipity: Towards a taxonomy and a theory. *Research Policy*, *47*(1), 169–179. https://doi.org/10.1016/j.respol.2017.10.007