

Learning in a Correlated World

Arjada Bardhi¹ and Steven Callander²

¹Department of Economics, New York University, New York, NY, USA

²Graduate School of Business, Stanford University, Stanford, California, USA;
email: sjc@stanford.edu

ANNUAL REVIEWS CONNECT

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Annu. Rev. Econ. 2026. 18:391–415

First published as a Review in Advance on
May 6, 2026

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-051624-072515>

Copyright © 2026 by the author(s). This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND 4.0), which permits any noncommercial use, sharing, distribution, and reproduction in any medium or format, provided the original author(s) and source are credited; this license does not permit sharing adapted material derived from this article or parts of it. Images or other third-party material in this article are included in the article's Creative Commons license unless indicated otherwise; see credit lines for license information.

JEL codes: D83, C11, D80, L10, D72



Keywords

correlated learning, Brownian motion, trial-and-error search, Gaussian covariance approach, informational

Abstract

We survey a growing literature on correlated learning—that is, how information from one choice changes beliefs about others when outcomes are correlated. The core modeling innovation is to represent the unknown mapping from choices to outcomes as the realized path of a stochastic process, most commonly the Brownian motion. We show how the framework has been applied to four canonical economic problems in which correlated learning is key but understudied: (a) search and experimentation, (b) communication, (c) innovation and market competition, and (d) attribute problems. We review emerging empirical and experimental evidence on correlated learning and outline new theoretical and methodological directions.

1. INTRODUCTION

No economic observation lives in isolation: Observing one outcome inevitably reshapes expectations about others. Because outcomes are correlated, information gleaned from one context spills over into beliefs about others. These informational spillovers systematically influence decision-making across economic life, as the following scenario illustrates.

Scenario 1 (Learning from experience). A new graduate's first job at Goldman Sachs goes badly: long hours, tedious work, toxic culture. Now they must choose what's next: a similar role at Morgan Stanley, a hedge fund in Chicago, manufacturing in California, or teaching in Ohio.

The graduate's experience at Goldman Sachs naturally informs their next choice. Disliking that job counts strongly against taking a similar role at Morgan Stanley, counts somewhat against joining a hedge fund, but says little about teaching or manufacturing. The outcomes of jobs are correlated, albeit imperfectly and unevenly. The informational spillover from their experience allows the graduate to update their beliefs about other jobs, with its influence fading as the alternatives grow less similar.

We refer to this as correlated learning: Bayesian agents infer an unknown mapping from choices to outcomes across a structured domain, where observing the outcome of one choice informs beliefs about others through their underlying similarity. To take the example from above, Goldman Sachs reveals more about Morgan Stanley than about teaching in Ohio, since both are finance careers. In this sense, correlated learning is a form of Bayesian updating, but one in which the unknown object is an entire choice–outcome mapping, and the correlation structure across choices encodes similarity.

Modeling correlated learning tractably has proved challenging. Many classic frameworks sidestep it by assuming either independence or perfect correlation across choices—extremes that leave important questions out of scope. This survey reviews an approach that directly models imperfect correlation based on choice similarity, is applicable across a wide range of economic settings, and yields sharp theoretical predictions. We synthesize applications to date and highlight open research directions.

1.1. The Broad Economic Relevance of Correlated Learning

Learning from one's own experience is a primary source of correlated learning—whether for a job seeker, a consumer, or a researcher. However, experience is not the only source: Correlated learning can arise from market signals, experts, peers, or experiments, and it matters in decision problems far beyond search, as the following scenarios illustrate.

Scenario 2 (Learning from an expert). A mechanic recommends an expensive set of repairs as necessary. Should the car owner take the expert at their word, or should they infer that minor fixes would likely suffice or that living with a few creaks is acceptable? Or does the mechanic's expertise tip the balance toward following the recommendation despite the cost?

Scenario 3 (Learning from the market). A firm contemplates entry into the airplane manufacturing market. It has observed Boeing's success with the 787, difficulties with the 737 MAX, and Airbus's failure with the giant A380. What do these observations imply about demand for a mid-size plane and features that the market values? Should the firm enter at all?

Correlated learning cuts across many problems and decision contexts. In the first three scenarios, the decision-maker seeks a single choice with a good outcome, learning from their own past

experience, experts, or the market. In others, as in the following scenario, the objective is to learn the entire mapping, choosing the observations that most improve inference.

Scenario 4 (Learning across subpopulations). A researcher has evidence from a randomized controlled trial (RCT) on a drug's efficacy for two demographic subpopulations and funding to test a third. Which subpopulation should they sample next?

The researcher's objective is not to find the subpopulation in which the drug is most effective, but to sample a subpopulation that adds the most information given existing evidence. As drug effects are correlated across subpopulations, their choice depends on the range of demographic variation and the subpopulations studied in prior RCTs. The objective differs from earlier scenarios, but correlated learning is still at the center.

1.2. A New Modeling Approach

The key novelty this approach proposes is to model correlated learning by representing the uncertain outcome mapping as the realized path of a stochastic process—most commonly, the Brownian motion—indexed by available choices or information sources rather than time. This is an unusual but productive use of stochastic processes. The path is drawn once over the space of choices or information sources. Decision-makers then learn the mapping by observing outcomes at selected alternatives and draw inferences about other alternatives.

The new approach contains two conceptual ideas. The first is to represent alternatives spatially, capturing correlation in outcomes via a mapping from alternatives to outcomes. In many decision problems, the set of alternatives possesses some natural ordering—e.g., some jobs are more similar than others. By embedding the alternatives in a space, the correlation in their outcomes can be captured through the set of possible outcome mappings. This set dictates what is possible, and the decision-maker's prior belief over it determines what is likely. With a judicious choice of the admissible set of mappings and prior, the framework disciplines the correlation structure, sidesteps the intractabilities of an unstructured correlation matrix, and captures key features of correlated learning in practice.

The second conceptual idea is that stochastic processes, with Brownian motion as the canonical case, are a particularly tractable and intuitive way to do this. Key properties of the Brownian motion, such as local learnability and proportional invertibility, allow natural extrapolation across nearby alternatives without trivializing learning: The decision-maker learns from a fragment of a larger truth but cannot pin down the entire truth. More generally, in the class of Gaussian processes (GPs), the covariance function is the primitive that encodes similarity among alternatives. The choice of covariance disciplines the nature of correlated learning.

1.3. Applications: Explored and Open

The new approach provides a framework that can be applied broadly. It has already proved productive in four classes of economic problems: search and experimentation, communication, innovation and market competition, and attribute problems. Correlated learning has been introduced to each problem using the same mathematical backbone. Prior to this approach, these literatures largely progressed independently and under starkly different correlation assumptions—some imposed independence and others perfect correlation for tractability—which obscured the importance of correlated learning. The new approach brings these problems together through a common structure.

We review advances that apply the Brownian motion framework across these classes of problems, alongside emerging work that builds related frameworks with other Gaussian and non-Gaussian stochastic processes. We also cover a nascent empirical and experimental literature

that brings correlated learning to data. Many questions remain—richer applications within already explored classes of problems, introduction of correlated learning to new canonical problems, and open methodological questions.

1.4. Why Was Correlated Learning a Gap in the Literature?

The obstacle to modeling imperfect correlation across choices was primarily methodological. In the search and experimentation literatures, for example, the importance of correlation in payoffs was certainly recognized. Yet the classic index results of Gittins (1979) and Weitzman (1979) rely on independence of outcomes, and efforts to extend the analysis to correlated outcomes made little progress except in special cases [such as the negative bandit model by Klein & Rady (2011)]. One could in principle introduce an arbitrary joint distribution over outcomes to capture imperfect correlation, but without further discipline such generality offers little traction. Two restrictions make the approach analytically manageable: (a) Gaussian beliefs, so that covariance becomes the natural primitive for similarity; and (b) an ordered set of choices, which disciplines inference.

Although the use of stochastic processes to model the outcome mapping is surprisingly tractable, that tractability has its limits. Progress was made by initially stepping back in other modeling dimensions. The first models of trial-and-error search assume short-sighted decision-makers, for example, and the communication models identify a tractable subset of equilibria, not a full characterization. These simplifications have allowed the approach to generate concrete insights into how imperfect correlation shapes behavior across a variety of economic problems. With these foundations in place, the literature is now developing tools to relax them.

2. FRAMEWORK

A Bayesian decision-maker faces a world of interdependent uncertainties. In its general form, alternatives $\mathcal{X}$ map to outcomes $\mathcal{O}$ through an unknown mapping $\psi : \mathcal{X} \rightarrow \mathcal{O}$, which is drawn from a family of mappings Ψ according to a prior probability measure Π .¹ The pair (Ψ, Π) encodes similarity across alternatives, and thus the correlation of outcomes. A key challenge in modeling correlated learning is finding a disciplined yet natural and tractable choice of (Ψ, Π) .

The approach represents alternatives spatially and models the unknown mapping as the realized path of a stochastic process indexed by $\mathcal{X}$. In the workhorse formulation, the space is one-dimensional (the real line), and the mapping is the realized path of a Brownian motion that evolves across alternatives rather than through time. That is, Ψ is the family of Brownian sample paths, and Π is the Wiener (canonical) measure. Using stochastic processes in this way disciplines correlation transparently. We begin with Brownian motion, which captures appealing properties of correlated learning, and later we broaden the lens to other processes.

2.1. Brownian Motion Uncertainty

Following the exposition by Callander (2011a), let $\mathcal{X} = \mathcal{O} = \mathbb{R}$ and let ψ be the realized path of a Brownian motion with drift μ and volatility σ^2 , anchored at an observed alternative $(x_0, \psi(x_0))$. Nature draws ψ once and for all. The drift μ captures the expected rate of change across alternatives, whereas volatility σ^2 captures the ruggedness of the process. For this exposition, we take $\mu < 0$ and $\psi(x_0) \geq 0$. **Figure 1a** illustrates a realization of ψ and the expected outcomes at the start.

The agent seeks to learn ψ in order to make better decisions. They know the data-generating process and parameters (μ, σ^2) but observe only a finite data set of alternative–outcome pairs

¹From Section 3 onward, we refer to elements of $\mathcal{X}$ as alternatives, actions, attributes, or information sources, depending on the economic context.




Figure 1

A realization of outcome mapping ψ (solid line) and expected outcomes (dashed line). Panel *a* shows prior expected outcomes. Panel *b* shows expected outcomes given $(x_0, \psi(x_0))$ and $(x_1, \psi(x_1))$.

$\mathcal{D} = \{(x_i, \psi(x_i))\}_{i=0}^n$. This provides two inputs into learning: structural understanding of the process, and data about the mapping itself. Combining them yields posterior beliefs about untried alternatives, which remain uncertain.²

The independent increments property of the Brownian motion implies that the outcome of each untried alternative is distributed normally, with mean and variance that depend only on the nearest observed neighbor(s) in the data set. Let $l := \min_i x_i$ and $r := \max_i x_i$; we refer to these two alternatives as the boundary of knowledge. Alternatives in $[l, r]$ form the body of knowledge, whereas those in $(-\infty, l) \cup (r, \infty)$ lie on the frontier of knowledge.

For any alternative on the frontier, its expected outcome given data set $\mathcal{D}$ is

$$\mathbb{E}[\psi(x) \mid \mathcal{D}] = \begin{cases} \psi(r) + \mu(x - r) & \text{for } x > r \\ \psi(l) - \mu(l - x) & \text{for } x < l \end{cases}, \quad 1.$$

with conditional variance

$$\text{var}[\psi(x) \mid \mathcal{D}] = \begin{cases} \sigma^2(x - r) & \text{for } x > r \\ \sigma^2(l - x) & \text{for } x < l \end{cases}. \quad 2.$$

On the frontier, beliefs are drift driven: The expected outcomes are anchored at the boundary outcome with slope μ . Variance grows linearly with distance from the boundary at rate σ^2 . Thus, novelty, understood as distance from the boundary, governs both the expected departure from the boundary outcomes and uncertainty. Both vanish at the boundary.

Alternatives that are in the body of knowledge sit on a Brownian bridge. For any two neighboring alternatives x_1 and x_2 in the data set, beliefs for alternative $x \in (x_1, x_2)$ are

$$\mathbb{E}[\psi(x) \mid \mathcal{D}] = \frac{x_2 - x}{x_2 - x_1} \psi(x_1) + \frac{x - x_1}{x_2 - x_1} \psi(x_2), \quad 3.$$

$$\text{var}[\psi(x) \mid \mathcal{D}] = \sigma^2 \frac{(x_2 - x)(x - x_1)}{x_2 - x_1}. \quad 4.$$

²Following Callander (2011a), we take noiseless observations of outcomes $(x_i, \psi(x_i))$ as the baseline. With additive noise $y_i = \psi(x_i) + \epsilon_i$, where noise $\epsilon_i \sim \mathcal{N}(0, \tau^2)$ is independent and identically distributed across i , the posterior remains Gaussian and prediction is therefore tractable. However, the nearest-neighbor property fails: Beliefs about an untried alternative depend no longer only on its immediately adjacent observations but also on potentially the entire data set (see also Section 4).

The expected outcome on a bridge is intuitive: It is given by the linear interpolation of the two neighbors' outcomes. Each is weighted by the relative distance from the other neighbor, and the drift term plays no role. Variance is concave on (x_1, x_2) , peaks at the midpoint alternative $(x_1 + x_2)/2$, and naturally vanishes at observed alternatives x_1 and x_2 . **Figure 1b** illustrates the expected outcomes both in the body and in the frontier of knowledge after some alternative x_1 is observed.

The agent's preference depends on the decision problem at hand. We defer details to Section 3, which introduces four classes of problems from the literature.

2.2. Appealing Properties of the Brownian Motion

At the heart of correlated learning is informational spillover: Observing one outcome changes beliefs about the outcomes of nearby alternatives. The extent and shape of the informational spillover determine the nature of correlated learning. The Brownian motion structures informational spillover in a way that captures the following intuitive properties of correlated learning.

- **Ordered expectations:** Due to the complexity of the environment, the agent can order the alternatives according to the expected outcome but not the realized outcome. Thus, the agent knows which alternatives are likely to produce good outcomes and which are not.
- **Partial invertibility:** The outcome mapping is partially invertible. Though knowing the outcome of some alternatives informs beliefs about other alternatives, the informational spillover is imperfect. Beliefs about an alternative are uncertain until it is tried, regardless of how many alternatives the data set contains.
- **Proportional invertibility:** The accuracy of beliefs is decreasing in the distance of the untried alternative from the data set. This proportional invertibility matches the intuitive idea that more novel alternatives carry more uncertainty. Moreover, because the sample path is continuous, uncertainty vanishes arbitrarily close to known alternatives, capturing the idea that similar alternatives produce similar outcomes.
- **Local nonlearnability:** The outcome mapping is not locally learnable under Brownian uncertainty: Learning the outcome of alternatives in one region of the space—even arbitrarily many—reveal only limited information about the global mapping. Local learnability is analogous to an econometrician being able to extrapolate from their data set with an arbitrarily high degree of accuracy.
- **Complexity:** Volatility σ^2 can be thought of as parameterizing the difficulty of extrapolating from the data set, that is, the complexity of the environment. The larger σ^2 is, the less invertible is the mapping and the more complex is the environment. In many contexts, agents will trade off variance against the expected rate of change μ —in such settings, the ratio σ^2/μ captures the relative complexity.
- **Tractability:** Despite the ruggedness of the mapping, prediction is not only tractable but also strikingly simple. The Brownian motion's independent increments imply that beliefs about any untried alternative depend exclusively on its immediately adjacent tried alternatives—what we will refer to going forward as the nearest-neighbor property. Learning reduces to local comparisons. Moreover, as beliefs are normally distributed, the belief moments in Equations 1–4 combine neatly with mean-variance preferences, as commonly used in many applications.

Driving the attractive learning properties of the Brownian motion is that the increments of the process are stationary and independent. Indeed, in an important precursor paper in this literature, Jovanovic & Rob (1990) make the case for the Brownian motion constructively, invoking the well-known axiomatization that the Brownian motion is the unique process with stationary

and independent increments and continuous sample paths (Billingsley 1995). Thus, while it is not necessary to use the Brownian motion to model correlated learning, it is appropriate if we wish to capture the learning properties listed above. It is worth observing that independent increments also imply nondifferentiability of the mapping. In that sense, the ruggedness of the Brownian path is a feature and not a bug in modeling correlated learning.

2.3. Beyond One Dimension

In many applications, the space of alternatives is inherently multidimensional; an alternative $x = (x_1, \dots, x_d)$ is a vector of d features. To allow for this, Bardhi & Callander (2025) extend the Brownian framework to two dimensions through what they refer to as a Brownian staple. This GP retains the underlying logic of independent, stationary increments with a continuous sample path, although the path is now a surface and the increments are rectangular as well as linear.³

Bardhi & Callander (2025) show that many of the learning properties of the Brownian motion carry over to two dimensions. The exception is the nearest-neighbor property. In two dimensions, the space of alternatives is only partially ordered, and hence the Brownian staple is Markov in a weaker sense. The increments within the data set may be overlapping, and hence dependent, rather than independent as in the Brownian motion. Prediction requires disentangling these increments. Consequently, alternatives other than the nearest neighbors matter for prediction. In fact, they show that for generic histories, the entire data set can be relevant for predicting an unknown alternative. Moreover, a generalized form of the nearest-neighbor property—and, thus, predictive simplicity—is restored once the data set is augmented with a finite set of derivative alternatives that recombine alternatives in the original data set in new ways.

3. APPLICATIONS TO FOUR ECONOMIC PROBLEMS

In what follows, we group the literature applying the Brownian framework into four broad domains: search and experimentation, communication, innovation and market competition, and attribute problems. The taxonomy is heuristic, as several papers straddle categories. The common structure across domains is that uncertainty and beliefs follow Section 2; what varies is the strategic and informational environment the agents face.

3.1. Search and Experimentation

In many settings, agents learn by trial and error. They try alternatives, experience outcomes, and learn from them as they go. This raises natural questions: Do the agents eventually find a good alternative? How long does it take? How bold is their search, and in which direction? The Brownian framework offers answers by endowing the alternative space with a metric and a notion of direction, yielding sharp implications for whether, where, and how aggressively to search.

3.1.1. Trial-and-error search. We present the workhorse model of Callander (2011a). Time is discrete: $t = 1, 2, \dots$. In each period t , agent t chooses an alternative $x_t \in \mathbb{R}$ that reveals outcome $\psi(x_t)$. We refer to x_0 as the status quo alternative. All actions and outcomes are publicly observed. The data set that agent t faces is thus $\mathcal{D}_t = \{(x_0, \psi(x_0)), (x_1, \psi(x_1)), \dots, (x_{t-1}, \psi(x_{t-1}))\}$. Agents care only about the realized outcome of their own choice. They share a common ideal outcome of zero and utility $u(x, \psi(x)) = -\psi(x)^2$ from alternative x with realized outcome $\psi(x)$. Because payoffs are myopic and information is public, equilibrium behavior reduces to a sequence of single-period

³The Brownian staple combines Brownian motions on axes (linear increments) with a Brownian sheet on the plane (rectangular increments). We discuss the correlation structure in Section 4.

problems (equivalently, a single myopic agent choosing in each period). Thus, if one agent repeats a known alternative and reveals no information, the next agent faces the same data set and makes the same choice, and so learning and search stop thereafter.⁴

At the heart of the Brownian motion framework is a trade-off between risk and return. An agent can repeat a known alternative and secure its outcome with certainty, or they can explore an untried alternative in hope of improvement but at risk of a worse outcome—a risk that grows with the boldness of their choice. The first agent resolves this trade-off as follows: accept the status quo if its outcome is close to the ideal, and search otherwise. When the agent does search, they search in the direction where better outcomes are more likely to be found, which is to the right of x_0 . The value $\alpha := -\sigma^2/(2\mu)$ provides the critical threshold in their choice.

Proposition 1 (Risk/return trade-off). The first agent's optimal search consists of

- (a) the safe alternative if its outcome is close to the ideal one: $x_1^* = x_0$ if $\psi(x_0) \in [0, \alpha]$;
- (b) a risky alternative otherwise: $x_1^* = x_0 - \frac{\psi(x_0) - \alpha}{\mu}$ with $\mathbb{E}[\psi(x_1^*)] = \alpha$ if $\psi(x_0) > \alpha$.

The agent accepts the status quo, even if it falls short of the ideal, when the expected gain from searching is not worth the added risk. When searching, the agent targets an alternative with expected outcome α rather than the ideal zero. The worse the status quo outcome, the bolder the optimal step. That α is the threshold in both cases is due to quadratic loss, though the underlying trade-off is more general. **Figure 2a** illustrates.

The second agent's problem mirrors the first's, except that two points are now known. This small difference in information can have a substantive impact on search behavior. Suppose the first agent's outcome is positive. If the second agent continues looking to the right, the problem is the same as before but with beliefs anchored at $\psi(x_1^*)$ rather than $\psi(x_0)$. Thus, the agent would repeat x_1^* if its outcome is good enough, and they would search to the right otherwise, choosing the alternative with an expected outcome of α (see **Figure 2b**).

The difference is that when $\psi(x_1^*)$ is far from zero, searching further to the right can dominate repeating x_1^* yet still be inferior to reverting to the status quo. Reaching a good expected outcome would require a step so bold that the added risk would outweigh the gain, and reverting to the status quo would be preferred. In this case, learning becomes stuck.

Proposition 2 (Learning can get stuck). When $\psi(x_1^*) > \alpha$, the second agent's optimal search

- (a) reverts to $x_2^* = x_0$ if $\psi(x_1^*) > \frac{\alpha^2 + \psi(x_0)^2}{2\alpha}$;
- (b) consists of a risky alternative $x_2^* > x_1^*$ such that $\mathbb{E}[\psi(x_2^*)] = \alpha$ otherwise.

When learning is stuck, the agent accepts an outcome previously deemed not good enough, and search halts. Indeed, learning might get stuck even at very poor outcomes, provided that $\psi(x_1^*)$ is sufficiently unfavorable. Learning gets stuck not because the agent believes good alternatives do not exist; the agent's confidence that they do remains unchanged. However, finding them seems now more difficult than before, and they are likely to lie further away from known alternatives. Because outcomes are correlated, the lessons of experience spill over to other alternatives, and in this case, those lessons discourage further search and learning.

⁴The models in Section 3.1 blend elements of search and experimentation. In some applications (including the present one), utility depends on the outcome of the current choice—akin to bandit payoffs—whereas in others it depends only on the best outcome found so far, as in Weitzman's (1979) model. Search is particularly befitting here because agents are myopic and maximize current outcomes; in contrast, experimentation typically emphasizes dynamic trade-offs that value information for future payoffs, as discussed below under forward-looking search (Section 3.1.3). For consistency, we use the terminology of search throughout this section and specify payoffs whenever it matters.




Figure 2

Optimal trial-and-error search for (a) the first agent and (b) the second agent. (a) The first agent's optimal search when the status quo outcome $\psi(x_0)$ exceeds the threshold α , illustrating the risk/return trade-off. (b) The second agent's optimal search after a first-period observation, illustrating the possibility that learning gets stuck.

That learning can become stuck endogenously in this way is a key insight of correlated learning. It can affect any agent after the first, and it can emerge for different kinds of search histories. Notably, with independent outcomes—and no search cost—getting stuck cannot occur: Failures are discarded, and the next agent simply tries an untried alternative.

Callander (2011a) traces the search trajectory over time. Before it stops, search proceeds in two phases. In the monotonic phase, search continues to the right until an outcome overshoots zero. It then reverses course and enters the triangulating phase, iterating back and forth within a narrowing interval, potentially with many changes of direction. Eventually, search stops with probability one.

3.1.2. Search with non-satiated preferences. In many economic contexts, agents seek ever higher outcomes rather than an ideal outcome, with utility increasing in the realized outcome. Callander & Matouschek (2019) study such non-satiated preferences in the Brownian framework. A first challenge is methodological: Which risk preferences generate meaningful search in this environment? On an unbounded space with positive drift, an optimal choice may not exist. For meaningful search to emerge, risk aversion must bound the agent's boldness, without constraining it so much that search is never profitable. Callander & Matouschek (2019) show that standard risk aversion, combined with a crossing condition on the level of absolute risk aversion, addresses both issues.

Search with non-satiated preferences exhibits the workhorse risk/return trade-off and can get stuck after a sufficiently bad outcome. Yet substantive differences emerge. First, search does not begin for a starting outcome below a threshold—a kind of informational poverty trap. Second, the boldness of search increases in the attractiveness of the frontier outcome. Third, triangulation never occurs. Fourth, search has positive probability of continuing forever.

3.1.3. Forward-looking search. A forward-looking agent may value not only the immediate payoff from the realized outcome but also the information spillover for future search. Garfagnini

& Strulovici (2016) study a sequence of two-period-lived agents who can search twice.⁵ The Brownian path is driftless ($\mu = 0$), and preferences are non-satiated and risk neutral. Search is costless within the body of knowledge but incurs an R&D cost on the frontiers that increases and is convex in the boldness of search; the R&D cost ensures the existence of an optimal choice.

There is option value to experimentation: By trying a new alternative in the first period, the agent may find a better outcome, which increases their payoff in the second period when they exploit the best discovered. With $\mu = 0$, the option value is driven solely by variance: A more uncertain alternative is more likely to outperform the current best. In the body of knowledge, attainable variance is disciplined by the distance between any two neighbors. On the frontiers, variance grows with boldness, but it is disciplined by the cost of search.

Forward-looking search reinforces the workhorse logic while revealing new properties. The risk/return trade-off need not apply, but getting stuck remains relevant. If the frontier outcome disappoints, search can return to the body of knowledge. This can generate cycles in search, as if search within the body of knowledge disappoints or the learning opportunities are exhausted, later agents might return to the frontier. Although agents are forward-looking, search ends with probability one. This contrasts with non-satiated search and short-lived agents but aligns with the workhorse formulation. The difference is in the cost of search: In Callander & Matouschek's (2019) model, the cost declines endogenously with risk aversion as outcomes improve, whereas in Garfagnini & Strulovici's (2016) model the R&D cost is constant.

Solving for optimal search with an infinitely lived, forward-looking agent under correlated alternatives remains daunting and largely open. Progress has been made by restricting attention to contiguous search (Urgun & Yariv 2021, 2025; Cetemen et al. 2023; Wong 2025): The agent is constrained to search alternatives continuously in order, as a miner would search for minerals continuously along a vein. What the agent chooses is not which alternatives to take but rather the speed of search, and they pay a cost that is increasing in speed. In Urgun & Yariv's (2025) model the agent stops and claims the best discovered outcome, whereas in Wong's (2025) model the agent collects payoffs along the way from all discovered outcomes (see also Li 2023).⁶

Optimal contiguous search preserves key behaviors from myopic search. In particular, learning can get stuck even with an infinite horizon. Optimal search follows a drawdown rule, that is, it stops when the current outcome falls a certain distance below the peak outcome discovered thus far. In Urgun & Yariv's (2025) model, the optimal search speed is U-shaped—highest when approaching the current peak outcome or when near the lower stopping threshold. In Wong's (2025) model, search increases monotonically the further the current outcome lies below the peak, reflecting the difference in payoff modeling (terminal versus flow payoffs).

3.1.4. Recombinant search. Bardhi & Callander (2025) extend optimal search à la Callander (2011a) to the two-dimensional Brownian staple under the assumption that the history is augmented with derivative alternatives.⁷ They find that recombinant search emerges even without an inherent benefit of combining and despite the multiplicative variance that it brings. Recombinant search emerges when one-dimensional search would otherwise get stuck. Pivoting and switching exclusively to the other dimension means that any progress in the first dimension is abandoned. Recombinant search allows agents to leverage progress that has been made by combining it with elements from the second dimension.

⁵ Readers are referred to Callander & Hummel (2014), discussed in Section 5, for a two-period model with forward-looking agents and ideal outcomes (see also Callander & Matouschek 2019).

⁶ Oliveros (2025) allows firms to reset and restart product development from earlier successes.

⁷ Derivative alternatives are reconfigurations of previously searched alternatives that add no new increments to knowledge. In Callander's (2011a) model, the set of derivative alternatives is empty.

Recombinant search reinforces the workhorse logic. The risk/return trade-off drives choice for recombinant search too, learning can get stuck, and search ends with probability one. Nonetheless, differences emerge. In two dimensions, the risk/return logic generates cycles of search between the multiple frontiers and the body of knowledge, whereas in one dimension no cycles arise under myopic behavior. Another notable difference is that frontier search in two dimensions reveals new increments and expands the body of knowledge but it is not necessarily linked to drift; this gives rise to new forms of search. Relatedly, this raises a new perspective on what it means to be innovative in search. Bardhi & Callander (2025) show that on the frontiers, search expands the body of knowledge but does so one dimension at a time, whereas inside the body of knowledge, search refines rather than expands knowledge but it does so in both dimensions simultaneously.

The early model of Jovanovic & Rob (1990) is a hybrid that merges elements of the above models with classic independent outcomes. The space of alternatives is multidimensional, though each choice reduces to a single dimension. Each dimension is the interval $[0,1]$ with independently drawn mappings. Each agent pays an exogenous cost to sample an alternative, observes its outcome, and then chooses a single alternative. As the cost of sampling is fixed, greater uncertainty is better, as it reveals more information. An existence problem is avoided by the boundedness of each dimension; thus, when a new dimension is sampled, the agent chooses alternative 1 at the right edge.

Critically, there is no correlated learning across dimensions in Jovanovic & Rob's (1990) model. Correlated learning occurs only when an agent samples from a $[0,1]$ interval that has previously been searched, and thus it is confined to the body of knowledge. More importantly, independence across intervals means learning cannot get stuck: No matter how many disappointing outcomes have been observed, the agent can discard them and sample afresh from a new interval.

3.2. Communication

The canonical models of communication take the polar opposite position on correlated learning: Not only does learning across alternatives occur, but it is also perfect. In both the cheap-talk model of Crawford & Sobel (1982) and the verifiable information model of Milgrom (1981), a single underlying state determines the mapping from actions to outcomes, so learning one outcome reveals the state and thus the outcomes of all other actions. Outcomes are therefore perfectly correlated, which is consequential for equilibrium communication. In both models, the striking conclusion is that while communication is possible, it ultimately serves the decision-maker more than the expert. With verifiable information, the unraveling result yields full disclosure to the decision-maker. With cheap talk, communication is coarse, yet the resulting obfuscation harms the expert as well as the decision-maker—indeed, the expert would be better-off simply handing over their information for free.

The key insight from applying the Brownian framework to strategic communication is that the expert can retain power over the decision-maker when their expertise is richer and imperfectly correlated. Whether communication is via cheap talk or via verifiable information, the expert leverages their informational advantage to shift decision-making in their favor. In both cases, they are able to shape decision-making while keeping part of their knowledge private. This balance matches common intuitions about expert advice in practice.

Alternatives are now actions, ordered by distance from a known action. An expert knows the entire Brownian path that maps actions to outcomes—capturing rich expertise—and advises a decision-maker who chooses an action. In cheap talk, the expert sends an action recommendation contingent on the path. With verifiable information, the expert can credibly disclose any set of action–outcome pairs from the path. In all the models discussed below, the decision-maker's utility

matches that in the workhorse model of search: quadratic loss with an ideal outcome of zero. In the cheap-talk models, the expert also has quadratic-loss preferences but with ideal outcome $b > 0$, which is their bias. In the verifiable information models, the expert's utility is independent of outcomes and decreasing in the action taken.

3.2.1. Communication via cheap talk. Callander (2008) shows that if the expert's bias is not too large, they can hold all the power in equilibrium and attain their ideal outcome almost surely. To see why, consider the strategy in which the expert recommends the closest action to the known one that delivers their ideal outcome, that is, x^* such that $\psi(x^*) = b$. This is clearly incentive-compatible for the expert. More surprisingly, it is incentive-compatible for the decision-maker to follow the recommendation. The logic mirrors the risk/return trade-off from search. The decision-maker secures outcome b by following the expert's recommendation; deviating to another action risks a worse outcome in pursuit of a better one. Not surprisingly, an equilibrium exists precisely when the expert's bias is below the familiar threshold α from search.

The expert is able to retain power over the decision-maker in the Brownian framework because outcomes are only imperfectly correlated. Yet that same correlation also limits their power. Once the expert recommends their preferred action, information spills over to nearby actions. For bias above α , the decision-maker may prefer an expert's ideal outcome *ex ante*, but *ex post*—after observing the recommended action and its implied information—they can profitably override the advice. Only for bias less than α is the risk of overriding the expert not worth the return.

Communication in the Brownian framework shapes both the content of advice and who holds power. In equilibrium, decision-making is Pareto efficient—unlike the coarse partition equilibria of Crawford & Sobel (1982). When the expert's informational advantage is a Brownian path, their message can be precise—a single action recommendation—without revealing the entire mapping. By contrast, with perfectly correlated outcomes, precision collapses to full revelation, so advice can be informative only by obfuscating through coarse partitions.

Aybas & Callander (2025a) show that the expert's power can extend to even larger bias levels when the action space is bounded. On a bounded space, there is positive probability that no action delivers the expert's ideal outcome. Surprisingly, by limiting the set of outcomes the expert can obtain, the bound makes it easier for the expert to persuade the decision-maker to accept their recommendation.

3.2.2. Communication via verifiable information. Verifiable information gives credibility to the expert's message. Callander et al. (2021b) show that for $\psi(x_0) > \alpha$, a continuum of equilibria exists in which the expert discloses the smallest action that produces some outcome $q \in [0, \alpha]$. By the risk/return logic, the decision-maker follows this recommendation.⁸ The expert prefers a smaller action, but unlike with cheap talk, if the expert discloses any other action (other than a larger action with outcome q), the decision-maker will know the expert has deviated.

This result is notable, as it does not succumb to the unraveling logic of Milgrom (1981). The difference with Milgrom's (1981) model is that the state space is much richer—a Brownian path versus a single dimension—and the states are not ordered. In the equilibrium, many states pool on each message: There are many paths that potentially go through the disclosed action–outcome pairs. For each state, a better alternative for the decision-maker exists, and they know it is to the right. However, because of the risk/return trade-off, the decision-maker accepts the recommendation, and because of this, no type has an incentive to deviate from the pooling strategy. Incomplete correlated learning and the richness of the Brownian path stop unraveling.

⁸For $\psi(x_0) \leq \alpha$, the expert-optimal equilibrium has no disclosure and decision-maker takes x_0 .

This equilibrium avoids unraveling and minimizes disclosure to a single outcome. Given that unraveling shifts power to the decision-maker, one might expect this equilibrium to be expert optimal. Callander et al. (2021b) show this is not true. They show that the expert is even better off in another equilibrium in which they disclose information beyond the recommended action—what they term referential advice.

A referential equilibrium exploits the same logic that makes learning get stuck in search, in which case the agent backtracks to a known alternative. Through carefully chosen referential advice, the expert deliberately steers the decision-maker toward a smaller action among those that are disclosed. For an expert who prefers the smallest possible action, this engineered backtracking is a feature, not a bug.

3.3. Innovation and Market Competition

Market competition is intense. When products are too similar, price competition can drive profits to zero. To gain market power, firms differentiate, both horizontally and vertically, by innovating under uncertainty. Capturing these forces in one model is hard. The Brownian formulation does so by casting market competition as a search process.

The approach takes the classic model of horizontal differentiation by Hotelling (1929) and layers on a Brownian path to represent product quality. In each period $t = 1, 2, \dots$, a firm t chooses a product $x_t \in \mathbb{R}$ that yields quality $\psi(x_t)$ or stays out. The market begins with an incumbent firm offering product x_0 with quality $\psi(x_0)$. Innovation on the frontier incurs an R&D cost that increases with novelty, which ensures the existence of an optimal product. All products and qualities are publicly observed. Consumers are distributed over the product space with a common willingness to pay and decreasing preference (transportation costs) for products farther from their ideal. Each period, active firms compete in prices. Firms are risk-neutral and their payoff depends on the quality of their own product, the rival products' qualities, and the horizontal differentiation between products.

3.3.1. Does competition spur innovation? The Brownian framework reframes this classic question. Rather than addressing only whether a firm innovates or not, it speaks to how novel the introduced product is relative to existing products. Callander & Matouschek (2022) study this in a single period with posted prices and consumers uniformly distributed on $[0, 1]$. They ask whether an entrant innovates more boldly when independent or when owned by the incumbent.

The benefit of horizontal differentiation in competition spurs the firm to innovate boldly, tempered only by the cost of R&D. An independent entrant seeks novelty to soften price competition, aiming to expand the market rather than compete for existing consumers. By contrast, an incumbent-owned entrant can coordinate prices and instead differentiates to minimize cannibalization of the incumbent's product. Callander & Matouschek (2022) show that the competitive effect dominates: Competition—via an independent entrant—spurs greater novelty in innovation.

This represents a spatial version of Arrow's (1962) replacement effect. Competition is good for innovation because it induces more novel innovation with a higher chance of a breakthrough. Yet, despite greater novelty, Callander & Matouschek (2022) show that the independent entrant is less likely to disrupt the incumbent and drive it from the market. The reason is the interaction of horizontal and vertical differentiation: While the independent entrant's product is more likely to achieve high quality, its greater distance in product space makes rivalry less direct, increasing the incumbent's chance of survival.

3.3.2. The dynamics of innovation. How does innovation in a market evolve over time? Does product novelty increase as the market matures? Do entrants expand the market's boundaries or outcompete incumbents for existing consumers? Does innovation ever stop?

Callander et al. (2025) study a dynamic model to address these questions. In their model, firms engage in third-degree price discrimination; consumers are uniformly distributed on $\mathbb{R}$; products are experience goods; and a competitive fringe copies each innovation after one period, driving the innovator's profit to zero.

The classic Hotelling's (1929) logic favors expanding the market. Innovating on the frontier creates room to escape competition, so serving new customers, all else equal, is more profitable than fighting for existing ones. Nonetheless, Callander et al. (2025) find that firms often enter within the market boundary, in niches between existing products. A niche is especially attractive after a disappointing frontier innovation. Because such disappointments lower the expected quality of frontier products, a niche entry can offer higher expected quality, even after accounting for tougher competition, and thus be more profitable.

Market competition differs from search, and firms react to disappointment differently. In search, a backtracking agent simply repeats a known alternative. In a market, by contrast, a firm learns from predecessors and must compete against them, so imitation is never profitable. Even when it turns away from the frontier, a backtracking firm still innovates with an untried product. This logic generates cycles between frontier and interior innovation. As firms exploit niche products, competition intensifies, making the frontier relatively attractive again. This congestion effect generates innovation cycles that do not arise in the analogous search model.

3.4. Attributes

In an attribute problem, the agent chooses experiments not to select a single alternative for implementation, but rather to learn about the entire outcome mapping. Each alternative is an information source about the underlying mapping: an attribute that, by correlation, reveals part of a larger truth. Payoffs, on the other hand, arrive only through a subsequent policy functional that aggregates across attributes (e.g., a weighted average or a threshold rule).⁹ Informational spillovers are therefore the objective rather than a byproduct. This framing encompasses consumer learning about product characteristics, drug trials measuring efficacy across subpopulations, and courts refining legal doctrine across fact patterns.

The key trade-off is distributional: A sample reduces uncertainty across the attribute space but does so unevenly. Let attributes on $[0, A]$ map to outcomes via a Brownian path anchored at $(0, \psi(0))$. Sampling an attribute farther out—say, a' versus $a < a'$ —reduces uncertainty on the frontier more while reducing it less near the anchor information source. Formally, the posterior variances from Equations 2 and 4 exhibit a single crossing at $\sqrt{aa'}$: Relative to sampling a , sampling a' leaves more residual uncertainty for attributes on $[0, \sqrt{aa'}]$ and less for those on $[\sqrt{aa'}, A]$. Thus, sampling a' trades off precision about more attributes on the Brownian bridge against precision about attributes on the frontier.

The agent then weighs this trade-off according to the details of the attribute problem. We present two natural variants from the literature. In both, the decision-maker samples a finite set of attributes from a continuum, observes outcomes and extrapolates to the rest, and only then chooses a policy—whether to purchase the product, whether to approve the drug, or how to decide on similar legal cases.

3.4.1. Single-dimensional policy. A complex product is a bundle of many attributes. A natural specification for a consumer's value for such a product is a weighted sum of attribute outcomes, that is, $v_i := \int_0^A \psi(a)\omega_i(a)da$, where ω_i captures consumer i 's weighting of attributes and $\omega_i(a) \neq 0$ means that the attribute is relevant. Bardhi (2024) studies a delegated sampling game built on

⁹Attribute problems share with farsighted search that sampling alternatives has informational value.

this idea: An agent chooses which k attributes to sample; once outcomes are observed, a principal sets a scalar policy equal to their estimated value for the product. Both evaluate the product by linear aggregation but place different weights across attributes. The question is which k attributes most effectively shape the principal's policy toward the agent's value. While Bardhi (2024) allows an arbitrary GP for the attribute mapping, we focus here on the Brownian benchmark and return to general processes in Section 4.

In the absence of any conflict over attribute weighting, the optimal sample is the one that most reduces uncertainty about the common value. The value of a sample is summarized by the posterior variance that it induces on the players' estimate of the value—notably, the prior mean plays no role. The most informative sample balances redundancy within the sample with generalizability to the rest of the attributes. In the Brownian benchmark, it is especially simple: For $k = 1$ the agent samples at $\frac{2}{3}A$, for $k = 2$ at $\{\frac{2}{3}A, \frac{4}{3}A\}$, for $k = 3$ at $\{\frac{2}{3}A, \frac{4}{3}A, \frac{6}{3}A\}$, and so on. Intuitively, the most informative sample is tilted toward more uncertain attributes, while the sampled attributes are equally spaced among themselves.

In the case of disagreement ($\omega_A \neq \omega_P$), the value of a sample is pinned down by two statistics: (a) the posterior variance for the principal and (b) the covariance that the sample induces between the principal's policy and the agent's value, which proxies alignment. The agent seeks to align the principal's policy with their own value while keeping the sample as uninformative about the principal as possible. Bardhi (2024) shows that in any optimal k -sample, every attribute is relevant for at least one player, and at most two are relevant to only one player. The optimal sample can be described explicitly under cases of stark disagreement. For example, with nonoverlapping interests,¹⁰ the optimal attribute for $k = 1$ is of relevance to the agent, and as the principal's mass of relevant attributes grows, the sample shifts farther away from the principal's attributes. Intuitively, the agent prioritizes not giving too much information to the principal, even if at the cost of low alignment.

3.4.2. Multidimensional policy. In Callander & Clark's (2017) model, attributes are case facts to be decided upon. All courts, a higher court as well as all lower courts, share a common preference over verdicts: for each case a , convict if $\psi(a) \geq 0$ and acquit otherwise. Each period, lower courts adjudicate a uniform mass of cases. For each case, the lower court can identify the facts of the case but not the correct verdict, that is, it observes a but not $\psi(a)$. For guidance on the correct verdict, the court looks to a higher court that can reveal outcomes of select cases through its verdict and concomitant opinion. This judgment is then a precedent that guides the lower court's decision-making. This is an attribute problem because the higher court chooses which cases to decide not to resolve those disputes alone but to learn and transmit the mapping from cases to correct verdicts across the entire docket.

Two features are worth noting. First, unlike Bardhi (2024), Callander & Clark (2017) posit that the policy is multidimensional: The lower courts choose a binary verdict for each case. Second, the courts' payoff is linear in the distance of the outcome from the zero threshold. Thus, the benefit from a correct verdict and the cost from a wrong one both grow with $|\psi(a)|$, that is, according to how clear-cut the case is. Given a set of precedents, the courts' objective aggregates the case-level payoffs across the full docket.

The optimal case addressed by the higher court is still governed by the key trade-off, but the weighting on where learning matters depends on the current precedent set. Callander & Clark (2017) show that, all else equal, it is more valuable for the courts to learn the outcomes of case facts for which the expected outcome is close to the zero threshold. For these cases, the error

¹⁰That is, the supports of ω_A and ω_P are disjoint.

uncertainty is highest. In contrast to Bardhi (2024), therefore, the authors posit that the optimal sample depends very much on the expected outcome of the case being sampled.

The higher court balances the reduction in posterior variance of hearing a case against a reduction in error uncertainty across all case facts. The optimal case to hear lies between the cases that maximize outcome uncertainty and error uncertainty. Callander & Clark (2017) trace out the evolution of cases the higher court hears—the evolution of judicial precedent and doctrine—capturing the path dependence of the law. This dynamic is akin to a search process albeit within an attribute problem.

Carnehl & Schneider (2025) also study an attribute problem in which each attribute is a research question, and a policy is chosen for each question: either apply existing knowledge to form a conjecture or take an outside option. The payoff from addressing a question falls with its residual uncertainty given current knowledge, and total payoff aggregates across the mass of questions. The differential learning force appears in how adding a new discovery improves the value of knowledge by redistributing precision across conjectures. Resolving a question in the frontier makes it possible to form less precise conjectures about more questions. Addressing questions within the support of existing knowledge increases precision differentially for questions in the vicinity, making questions of intermediate distance most preferred. Frontier discoveries dominate only once knowledge is sufficiently refined, in the sense that the steps taken by past discoveries are small enough.

4. OTHER STOCHASTIC PROCESSES

The success of Brownian motion in modeling correlated learning naturally raises the question of whether other stochastic processes can play a similar role. Brownian motion is in many respects the canonical benchmark, but alternatives are valuable both to confirm the robustness of the insights identified in Brownian environments and to better capture economic settings in which Brownian linearity is less natural.

4.1. The Gaussian Covariance Approach

One way to generalize beyond Brownian motion is to recognize it as a special case of a GP, defined by a prior mean function and a pairwise covariance function. Our exposition here follows Bardhi (2024), who studies correlated learning for a broad class of GPs. A GP prior implies that for any finite set of alternatives $\{x_1, \dots, x_n\}$, the respective outcomes $\{\psi(x_i)\}_{i=1}^n$ are jointly normal. Expected outcomes are given by the mean function $\mu : \mathcal{X} \rightarrow \mathbb{R}$, while dependence across alternatives is governed by the covariance function $\sigma : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, with variance at x naturally given by $\sigma(x, x)$.¹¹

With GP prior $\psi \sim \mathcal{GP}(\mu, \sigma)$ and a data set $\mathcal{D} = \{(x_i, \psi(x_i))\}_{i=1}^n$, beliefs about unobserved alternatives follow Bayes' rule. Three objects drive this updating: (a) the covariance matrix $\Sigma_n = [\sigma(x_i, x_j)]_{i,j=1}^n$ among the observed alternatives (i.e., how informative $\mathcal{D}$ is as a whole); (b) for any target alternative x , the covariance vector between x and each observed alternative (i.e., how similar x is to what we see in $\mathcal{D}$); and (c) the precision matrix Σ_n^{-1} , which down-weights redundant information within $\mathcal{D}$. First, the expected outcome $\mathbb{E}[\psi(x) | \mathcal{D}]$ starts from $\mu(x)$ and shifts toward observed outcomes by a similarity-weighted average of their deviations from their own prior means. Alternatives more similar to x that are also less correlated with the rest of the data set get more weight. Second, the posterior covariance equals the prior covariance minus what

¹¹The covariance function must be positive semidefinite, that is, for any finite selection of alternatives, the corresponding covariance matrix is positive semidefinite.




Figure 3

True outcome mapping (gold line), posterior mean (blue line), and 95% credible interval (shaded area) under three covariance functions and three observed alternatives.

can already be explained by the data set, so uncertainty shrinks most near heavily sampled regions, and correlations vanish when the observations fully mediate the relationship between two alternatives. These mechanics generate the prediction patterns we showed under Brownian motion. **Figure 3** contrasts patterns for three different processes.

The Brownian motion is a special case of this formulation. For the Brownian motion anchored at $(x_0, \psi(x_0))$, the covariance function is $\sigma(x, x') = \sigma^2 \min\{|x - x_0|, |x' - x_0|\}$ if x, x' are on the same side of x_0 and $\sigma(x, x') = 0$ otherwise. This reflects the independent increments property. The Brownian staple of Bardhi & Callander (2025) is another special case. For any two alternatives $(x, y), (x', y') \in \mathbb{R}_+^2$, the covariance between their outcomes is $\sigma((x, y), (x', y')) = \sigma_x^2 \min(x, x') + \sigma_y^2 \min(y, y') + \sigma_{xy}^2 \min(x, x') \min(y, y')$, where σ_x^2, σ_y^2 capture dimension-specific volatilities and σ_{xy}^2 captures cross-dimension volatility.

The appeal of GP priors rests on the flexibility of the covariance function to encode a wide range of economic dependencies while still disciplining learning in a structured way. The choice of covariance is thus a concise statement about how observing one alternative reshapes beliefs about others. We highlight two classes of processes.

4.1.1. Nearest-neighbor property. A distinctive feature of Brownian motion is that increments are independent, so the prediction of any alternative depends only on its nearest observed neighbors. In attribute sampling, Bardhi (2024) refers to this as the nearest-attribute property and proves that this property holds beyond the Brownian motion and is equivalent to a requirement that the process is Markov, which in turn is equivalent to the covariance satisfying, for any three ordered alternatives $x < x' < x''$,

$$\sigma(x, x'') = \frac{\sigma(x, x')\sigma(x', x'')}{\sigma(x', x')}. \quad 5.$$

One can interpret this property as modularity in the landscape: A set of observed alternatives partitions the domain into independent regions, so that prediction at an alternative depends only on the nearest observed alternative on each side.

4.1.2. Distance-based correlation. In many applications, it is natural to assume that the correlation between two alternatives depends only on how far apart they are, not on their absolute location. Brownian motion only partially reflects this idea: For instance, under standard Brownian motion the correlation between $x = 1$ and $x' = 2$ is about 0.71, whereas the correlation between $x = 99$ and $x' = 100$ is about 0.99, even though the distance between alternatives is the same. This is a feature, not a bug, of the Brownian motion: The dependence on location captures the anchoring effect of a status quo alternative. Yet in many settings it is more natural to model similarity as a pure function of distance $|x - x'|$, so that all alternatives are equally uncertain ex ante

and correlations decay with distance. Most of the alternatives to Brownian motion that have been studied in the literature are of this form.

Bardhi (2024) builds a bridge between the nearest-neighbor property and distance-based correlation by introducing a general family of powered-exponential covariance functions:

$$\sigma_p(x, x') = \exp\left(-\left(\frac{|x-x'|}{\ell}\right)^p\right), \quad \ell > 0, p \in (0, 2]. \quad 6.$$

Here the parameter ℓ controls how quickly correlation decays with distance: As $\ell \rightarrow 0$, attributes become independent, whereas as $\ell \rightarrow +\infty$, attributes become perfectly correlated. This class nests important special cases. When $p = 1$, the covariance is Ornstein–Uhlenbeck (OU), which preserves the nearest-neighbor property and is highly tractable as a result (Figure 3b). More generally, conditional on $\psi(x')$, the nonadjacent outcomes $\psi(x)$ and $\psi(x'')$ covary positively for $p < 1$, are independent for $p = 1$, and covary negatively for $p > 1$. When $p = 2$, the covariance is squared-exponential: It produces smooth learning landscapes, but in doing so it breaks the nearest-neighbor property and the tractability that comes with it (Figure 3c). These special cases recur throughout the literature.

4.2. Revisiting the Four Economic Problems

The four economic problems have each been studied with GPs beyond the Brownian motion. For brevity, we focus our discussion on the correlation structures here and, where appropriate, discuss the papers in greater detail in Section 5.

4.2.1. Search and experimentation. Both Bardhi & Callander (2025) and Hodgson & Lewis (2025) study sequential myopic search on multidimensional domains under a GP prior, but they adopt different covariance structures and, as a result, pursue different routes to tractability. Bardhi & Callander (2025) introduce the Brownian staple and aim for closed-form characterization of prediction and optimal search. Hodgson & Lewis (2025) build a structural model of consumer search based on the squared-exponential covariance and take it to data. Malmendier & Veldkamp (2022) analyze the dynamic choice between a risky and a safe action when risky payoffs are correlated across finitely many agents through a flexible covariance matrix.

4.2.2. Communication. Aybas & Callander (2025a) show that efficient cheap talk is not confined to Brownian environments, extending the result on an unbounded space to the class of Itô diffusions. Dong & Mayskaya (2024) study optimal selection of two experts when signals are correlated across an expert–issue space and cross-expert communication is costly. They characterize equilibrium diversity for a broad class of distance-based covariances—with the sharpest, symmetric characterization under OU correlation. Bardhi & Bobkova (2023) study production of verifiable information within minipublics and extend the analysis to cheap-talk communication between the minipublic and the policymaker.

4.2.3. Innovation and market competition. Covert (2015) and Hodgson (2024) build empirical models based on the squared-exponential covariance to model the effectiveness of drilling across locations and fracking inputs. Also using a squared-exponential covariance, Ilut et al. (2020) study, both analytically and numerically, the problem of a monopolist that nonparametrically learns the unknown demand mapping from its price-quantity history, thus providing a novel explanation for price rigidity.

4.2.4. Attributes. As discussed above, Bardhi (2024) introduces the covariance approach in an attribute problem and studies how the nearest-attribute property shapes strategic sampling of payoff-irrelevant attributes. Also in an attribute setting, Liang et al. (2022) characterize optimal forward-looking learning under flexible Gaussian attribute correlation among finitely many

attributes. Bardhi & Bobkova (2023) bring the covariance approach to committee design by casting minipublic selection as an attribute-sampling problem with endogenous disclosure: Each minipublic citizen chooses whether to reveal their own outcome to a biased policymaker. They show that the policymaker's problem reduces to maximizing a minipublic informativeness statistic for any GP prior and characterize precisely the optimal minipublic composition for OU correlation. Related questions of optimal team design with an attribute structure and correlation are also studied by Dong & Mayskaya (2024), who also rely on distance-based Gaussian covariances, and Ichihashi et al. (2025).

4.3. Non-Gaussian and Non-Bayesian Approaches

While GP priors have been effective in many settings, they can be ill-suited for at least two reasons. First, Gaussian likelihoods imply continuous, unbounded outcome support and are unnatural when outcomes are discrete or bounded. Second, unconstrained GPs generally fail to satisfy economically relevant shape restrictions on all sample paths—e.g., monotonicity—without bespoke constraints. Recent complementary approaches discipline the mapping parametrically or adopt non-Gaussian correlation structures.

Aybas & Callander (2025b) parameterize an environment that is the most minimal possible relaxation of perfectly correlated outcomes. In this minimally complex environment, the mapping is known to be linear but the decision-maker knows neither the intercept nor the slope (either $+1$ or -1). Thus, with one point in the mapping observed, informational spillover substantially reduces the decision-maker's uncertainty for each other alternative to two possible outcomes, yet it is not complete until a second point is observed.¹²

In the other direction, several papers have expanded analysis to non-Gaussian stochastic processes, most notably jump processes. Aybas & Callander (2025a) study the class of Itô diffusions. Ichihashi et al. (2025) study contracting in a spatial learning setting in which outcomes across a unit mass of problems evolve via a Markov jump process driven by Poisson shocks. The binary-outcome Markov structure makes the optimal task allocation and contract tractable to characterize.

Malladi (2025) takes a non-Bayesian route to analyzing forward-looking search with spatial learning. He studies an environment in which the outcome mapping lies in a known family—that of L -Lipschitz continuous functions, where L measures how rugged the landscape is. No probabilistic prior is imposed on the family; the agent instead takes a worst-case (robust) approach to search. Remarkably, Malladi (2025) is able to characterize the optimal farsighted search procedure. Relative to myopic search, the optimal strategy explores closer to past discoveries yet is more willing to test extreme alternatives. Relatedly, Banchio & Malladi (2025) extend the worst-case approach to costly sequential search when the agent knows that some alternative attains a target payoff but does not know which one—thus formalizing rediscovery. The family of mappings is the class of L -Lipschitz functions that pass through the target payoff.¹³

4.4. Beyond Economics

The idea that prediction is governed by the structure imposed on dependence has appeared in other fields, although with a different focus from the one on incentives and equilibrium play that is central to economics. In geostatistics, the method of kriging originated with Matheron (1971)

¹²An early Bayesian approach that disciplines the family of mappings directly is proposed by Aghion et al. (1991), who assume that the unknown payoff function is analytic, smooth, and/or quasi-concave and provide limit results for when experimentation leads to the global maximum being identified.

¹³Rugged landscape models in management also adopt a non-Bayesian approach, studying dynamics of search with prespecified learning heuristics, most commonly hill-climbing algorithms (Levinthal 1997).

as a way to construct the best linear unbiased predictor of an unobserved location based on nearby observations. It is built on variograms, which capture how variability grows with distance, which need not be Gaussian. When the underlying landscape is Gaussian, kriging coincides with what in machine learning is known as GP regression, which places a GP prior on the function to be explored and uses covariance functions to capture smoothness, periodicity, or other structured forms of dependence (Cressie 1993, Williams 1998, Rasmussen & Williams 2006, Schulz et al. 2018).

A distinct but related line of work studies Gaussian Markov random fields (Rue & Held 2005). This is a graphical approach in which the key object is the precision matrix rather than the covariance function. The sparsity pattern of the precision matrix defines a graph in which each variable is conditionally correlated only with its neighbors. Prediction of a variable therefore uses only its immediate neighbors, thus conceptually generalizing the nearest-neighbor property to less structured domains. Techniques from this literature are invoked by Bardhi & Callander (2025).

5. WHAT NEXT?

Ongoing work has advanced the framework in several directions, but much remains to be done. We emphasize opportunities that remain open.

5.1. Applications

There is broad scope to delve further into applications with the new approach. Often this requires reinterpreting institutions and practices through the lens of correlated learning. The most progress has come from questions the framework is built to address: informational spillovers across players and time, whether learning is effective or gets stuck, and a sense of distance and direction in choice.

5.1.1. Deepening. The abstract models of the previous sections have broad applicability but lack institutional context. How do markets, politics, and organizations shape correlated learning, and how should they be designed to improve decisions? Embedding the framework in concrete institutions is a natural next step, with work already under way.

5.1.1.1. Political institutions. In politics, policy choice is structured by institutions, whether within legislatures, by the bureaucracy, or in elections. This adds context to correlated learning in two ways: (a) Decision-making power (the policymaking) is separated from authority (the voters), and (b) decisions are made by a group (e.g., a legislature) rather than an individual (i.e., collective choice). Papers in this vein have focused on how these features affect decision-making in complex environments when learning is correlated.

Callander (2011b) explores the first of the two questions above by embedding the workhorse framework into a model of elections. Competition between parties leads in the short run to policies that are less attractive to voters; in the long run, however, it leads to better outcomes for voters, as power switching induces more search that finds better policies. In particular, competitive elections allow search to keep going where it would otherwise be stuck. Callander & Hummel (2014) address the same question by shifting attention to parties' incentives: How might a party choose policy when it knows power will switch to the other party in the future? They show that this creates a purely informational intertemporal linkage: The incumbent preemptively experiments with policy to shape the informational environment of its successor.

Callander & McCarty (2024) turn to the second of the two questions above in the context of legislative policymaking. They embed the workhorse Brownian framework into the canonical model of agenda setting. Correlated learning overturns the classic result that the agenda setter is

better off with an extreme status quo. The take-it-or-leave-it agenda setting model is only one of a large variety of legislative bargaining procedures. A general analysis of legislative bargaining with correlated learning remains wide open. Moreover, the rich uncertainty of the Brownian motion also opens up the question of search in legislative bargaining. Callander & McCarty (2024) offer a preliminary analysis of agenda setting in a two-period model. Howell & Jackman (2013) offer an interesting variant of legislative bargaining in which a policy choice produces two outcomes, one federal and one local. Such a dual-outcome bargaining process has many applications inside and outside politics.

Asymmetric information and expertise play a big role in policymaking. Callander (2008) applies his cheap-talk model to institutional design. Expert power provides a rationale for costly investment in expertise that is lacking in simpler environments, but the incentive is greater if a committee or bureaucrat is misaligned with the legislature. This exacerbates the agency friction, as then the ally principle fails and a legislature prefers to empower nonrepresentative committees and bureaucrats. Moreover, McCarty (2013, 2017) points out that as regulators themselves get their information from the firms they regulate, the firms have an incentive to adjust their behavior to protect their informational advantage. This leads not only to an informational form of regulatory capture but also to a distorted market outcome. Bardhi & Bobkova (2023) show that even with costless, transparent evidence and citizens aligned with the policymaker, the prospect of future disagreement depresses citizens' incentives to discover local evidence within a minipublic. To restore incentives, the policymaker distorts the minipublic's composition, reducing diversity and representativeness. The case of disagreement over attribute weighting discussed by Bardhi (2024) speaks to the same issue.

5.1.1.2. Organizations. The two questions presented in the previous section pertain more broadly to organizations, as do the flow and use of information. Dong & Mayskaya (2024) study communication barriers and expert selection in organizations. Ichihashi et al. (2025) model the optimal assignment of problems to workers in firms (see also Gans 2025a). Clark (2016) examines communication barriers from higher to lower courts through judicial decisions. He shows how courts trade off greater precision about the case facts at hand against applicability across the full judicial docket. Ganz (2020) layers two learning problems on top of each other—a manager learning about best practices and CEOs learning about a manager's ability—and shows that they are fundamentally incompatible.

In firms, even when authority is centralized, decision-making power is often dispersed among managers, and this creates a need to coordinate actions. Dall'Ara (2025) and Callander & Matouschek (2017) study this when learning is correlated. Related is learning across organizations. The infrequency, or often absence, of imitation is a puzzle in firms, politics, and the diffusion of innovations broadly. Callander & Matouschek (2017) show that the distance between firms in the product space matters and provides a barrier to imitation. Glick (2014) shows that imitation in politics can matter even when imperfect, as states sometimes mimic but often learn from and modify successful policies in other states.

5.1.1.3. Markets. In the papers discussed in Section 3, firms learn about product quality in the market. There is much more that firms need to learn to successfully compete. Ilut et al. (2020) study the classic monopolist's pricing problem as it learns about demand. They show that correlated learning in complex environments provides a new explanation for price rigidity. Callander & Matouschek (2022) examine how a regulator's antitrust policy affects innovation when learning is correlated. Hoxha (2024) adds correlated products into a model of market screening. Ideas from market design will surely find a fertile ground of application in this area.

5.1.2. Broadening. Another opportunity is to broaden the set of problems studied with the framework. The differences among search, industrial organization, and attribute problems can be seen as different payoff weightings on attribute–outcome pairs. A continuum of opportunity exists among these cases, many with substantive relevance.

There is also opportunity in the structure of the decision-making environment. For instance, strategic information production and communication under correlated learning with sender commitment to the information structure remains underexplored; Bardhi & Bobkova (2026) examine how statistics versus anecdotes perform in such settings. Another opportunity is in what exactly is unknown. The literature has focused on what might be thought of as data uncertainty—incomplete knowledge about the mapping. In many contexts, however, agents also face structural uncertainty about the parameters (μ, σ^2) . Callander (2011a) provides an example of search when agents lack knowledge of the drift parameter. Much remains to be learned.

5.1.3. Hybrid applications. The four types of problems are not mutually exclusive. In some contexts, more than one type of problem is relevant, and it is natural to combine them into a single model. Aybas & Callander (2025b) study the combination of search and communication and establish a strong interaction effect between the two sources of information. Bardhi & Bobkova (2023) study an extension in which the minipublic members discover outcomes privately and strategically communicate with the policymaker. Callander & Clark’s (2017) work can be seen as an example of search in an attribute problem. Many other hybrid models are possible. For example, it is natural to combine communication with market competition and innovation—for instance, in a setting in which a management consulting firm like McKinsey advises firms competing in a market where learning is correlated.

5.1.4. Behavioral applications. Inference and decision-making when learning is correlated can be cognitively taxing. Individuals may employ shortcuts and heuristics to aid decision-making. Using the Gaussian covariance approach, Ilut & Valchev (2023, 2025) model boundedly rational learning with two modes: an associative mode (System 1) that generalizes from experience and a deliberative mode (System 2) that acquires a costly targeted signal at the current state when the experience-based estimate is too uncertain. The opportunities here are many. The rich behavioral economics literature has clear relevance for correlated learning. Embedding realistic cognitive constraints into applications remains a largely open question.

5.2. Open Methodological Questions

As tractable and fruitful as the Brownian framework has proved to be, many questions remain out of reach and call for methodological advances. Brownian motion’s appeal is its clean geometry (in which nearby options move together), transparent updating, and comparative statics that travel well across applications. Even so, the optimal unconstrained, farsighted search policy over a Brownian path is not yet well understood in full generality, once we drop restrictions on search direction or the agents’ horizon. A natural next step is to seek priors and families of mappings that preserve Brownian-style local correlation while yielding lower-dimensional beliefs. One promising route is to discipline correlation through a Markov random field on a finite set of alternatives: With suitable neighborhood structure, belief updating remains local and tractable. Another route is to identify GP priors with simpler forms of uncertainty than the Brownian motion, mapping out a middle ground between the minimal complexity in Aybas & Callander’s (2025b) model and the maximal complexity in Callander’s (2011a) model and subsequent literature.

A complementary approach is to test the reach of distance-based covariances—such as OU or powered-exponential covariances—more widely on the four canonical problems. The properties of these processes are not yet fully understood, and it is important to explore which processes fit

better in which settings. Yet another direction is to move beyond GPs. Many environments feature breakthroughs or regime shifts, for which jump processes are a better fit. Ichihashi et al. (2025) provide a first framework in this direction, and more remains to be done with richer jump processes.

5.3. Taking Correlated Learning to Data

The growing availability of large administrative and industry data sets, together with new experimental designs, makes it increasingly feasible to bring correlated learning models to the field, a possibility that has only just begun to be explored.

Covert (2015) studies firms' choices of fracking inputs in the Bakken and finds that search follows correlated learning yet is short-sighted. Hodgson (2024) develops a complementary perspective in offshore oil exploration, modeling drilling outcomes as spatially correlated and emphasizing strategic behavior. Firms face a dynamic choice between drilling immediately and waiting to learn from rivals' well exploration. The observability of nearby outcomes creates information externalities that generate delays, underexploration, and clustering of wells. Hodgson & Lewis (2025) analyze consumer search over differentially correlated products, where product j is indexed by a vector of observable characteristics and the squared-exponential prior takes Euclidean distance between two products as the relevant distance. The estimated model generates search paths consistent with spatial learning: (a) Following rarely purchased products, consumers jump farther in the product space; (b) as search progresses, it moves closer to the eventual purchase; and (c) spatial learning delivers markedly higher welfare compared to a no-learning benchmark. Glick (2014) brings the Brownian motion framework to policy diffusion across organizations, providing qualitative evidence from university legal offices on how institutions mimic or adapt each other's responses to ambiguous laws.

Glick & Myers (2015) design an experimental search environment with Brownian motion uncertainty and study whether subjects optimally exploit others' discoveries or engage in excessive exploration. Subjects do not behave as if facing canonical uncertainty, indicating an understanding that Brownian uncertainty is different, yet they still struggle to act optimally in highly complex Brownian environments. More recently, Malladi et al. (2025) show that participants internalize spatial correlation in dynamic search: They move away from low prizes, stay near high prizes, and increasingly concentrate search around the best discoveries. Because such observed behavior is consistent with both Bayesian search under Brownian uncertainty (Callander 2011a) and robust search (Malladi 2025), more experimental work is needed to disentangle the underlying mechanisms—belief formation and miscalibration, the correlation structure and other forms of complexity, and the degree of forward-looking behavior.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

We thank Niko Matouschek, Nicolas Lambert, Yunus Aybas, Nina Bobkova, Joyee Deb, and Christoph Carnehl for many helpful conversations.

LITERATURE CITED

- Aghion P, Bolton P, Harris C, Jullien B. 1991. Optimal learning by experimentation. *Rev. Econ. Stud.* 58:621–54
- Arrow KJ. 1962. Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity*, ed. RR Nelson, pp. 609–26. Princeton University Press

- Aybas YC, Callander S. 2025a. Cheap talk in complex environments. Work. Pap., Stanford University and Texas A&M University
- Aybas YC, Callander S. 2025b. *Experts & experience*. Work. Pap., Stanford University and Texas A&M University
- Banchio M, Malladi S. 2025. Rediscovery. Preprint, arXiv:2504.19761 [econ.TH]
- Bardhi A. 2024. Attributes: selective learning and influence. *Econometrica* 92:311–53
- Bardhi A, Bobkova N. 2023. Local evidence and diversity in minipublics. *J. Political Econ.* 131:2451–508
- Bardhi A, Bobkova N. 2026. The power of anecdotes. *AEA Pap. Proc.* 116:640–45
- Bardhi A, Callander S. 2025. Recombinant search. CEPR Discuss. Pap. DP2135
- Billingsley P. 1995. *Probability and Measure*. Wiley. 3rd ed.
- Callander S. 2008. A theory of policy expertise. *Q. J. Political Sci.* 3:123–40
- Callander S. 2011a. Searching and learning by trial and error. *Am. Econ. Rev.* 101:2277–308
- Callander S. 2011b. Searching for good policies. *Am. Political Sci. Rev.* 105:643–62
- Callander S, Clark TS. 2017. Precedent and doctrine in a complicated world. *Am. Political Sci. Rev.* 111:184–203
- Callander S, Hummel P. 2014. Preemptive policy experimentation. *Econometrica* 82:1509–28
- Callander S, Lambert N, Matouschek N. 2021a. Arrow meets Hotelling: modeling spatial innovation. *AEA Pap. Proc.* 111:538–43
- Callander S, Lambert N, Matouschek N. 2021b. The power of referential advice. *J. Political Econ.* 129:3073–140
- Callander S, Lambert N, Matouschek N. 2025. Innovation and competition on a rugged technological landscape. *Am. Econ. J. Microecon.* In press
- Callander S, Matouschek N. 2017. *Imitation and coordination in firms*. Work. Pap., Stanford University and Northwestern University
- Callander S, Matouschek N. 2019. The risk of failure: trial and error learning and long-run performance. *Am. Econ. J. Microecon.* 11:44–78
- Callander S, Matouschek N. 2022. The novelty of innovation: competition, disruption, and antitrust policy. *Manag. Sci.* 68:37–51
- Callander S, McCarty N. 2024. Agenda control under policy uncertainty. *Am. J. Political Sci.* 68:210–26
- Carnehl C, Schneider J. 2025. A quest for knowledge. *Econometrica* 93:623–59
- Cetemen D, Urgan C, Yavir L. 2023. Collective progress: dynamics of exit waves. *J. Political Econ.* 131:2402–50
- Clark TS. 2016. Scope and precedent: judicial rule-making under uncertainty. *J. Theor. Politics* 28:353–84
- Covert TR. 2015. *Experiential and social learning in firms: the case of hydraulic fracturing in the Bakken Shale*. Unpublished manuscript
- Crawford VP, Sobel J. 1982. Strategic information transmission. *Econometrica* 50:1431–51
- Cressie N. 1993. *Statistics for Spatial Data*. Wiley. Rev. ed.
- Dall'Ara P. 2025. *Coordination in complex environments*. Unpublished manuscript
- Dong M, Mayskaya T. 2024. Does reducing communication barriers promote diversity? *J. Econ. Theory* 222:105932
- Gans JS. 2025a. A model of the Babbage firm. NBER Work. Pap. 34145
- Gans JS. 2025b. A simplified quest for knowledge. NBER Work. Pap. 33815
- Ganz SC. 2020. Hyperopic search: organizations learning about managers learning about strategies. *Organ. Sci.* 31:821–38
- Garfagnini U, Strulovici B. 2016. Social experimentation with interdependent and expanding technologies. *Rev. Econ. Stud.* 83:1579–613
- Gittins JC. 1979. Bandit processes and dynamic allocation indices. *J. R. Stat. Soc. B* 41:148–77
- Glick D, Myers CD. 2015. Learning from others: an experimental test of Brownian motion uncertainty models. *J. Theor. Politics* 27:588–612
- Glick DM. 2014. Learning by mimicking and modifying: a model of policy knowledge diffusion with evidence from legal implementation. *J. Law Econ. Organ.* 30(2):339–70
- Hodgson C. 2024. Information externalities, free riding, and optimal exploration in the UK oil industry. NBER Work. Pap. 33067
- Hodgson C, Lewis G. 2025. You can lead a horse to water: spatial learning and path dependence in consumer search. *Econometrica* 93:1299–332

- Hotelling H. 1929. Stability in competition. *Econ. J.* 39:41–57
- Howell WG, Jackman SP. 2013. Interbranch negotiations over policies with multiple outcomes. *Am. J. Political Sci.* 57:956–70
- Hoxha K. 2024. Selling correlated information products. Preprint, arXiv:2405.11142 [econ.TH]
- Ichihashi S, Li F, Zou D. 2025. *Organizational design in the knowledge economy*. Unpublished manuscript
- Ilut C, Valchev R. 2023. Economic agents as imperfect problem solvers. *Q. J. Econ.* 138:313–62
- Ilut C, Valchev R. 2025. Learning optimal behavior through reasoning and experiences. Preprint, arXiv:2403.18185 [econ.TH]
- Ilut C, Valchev R, Vincent N. 2020. Paralyzed by fear: rigid and discrete pricing under demand uncertainty. *Econometrica* 88:1899–938
- Jovanovic B, Rob R. 1990. Long waves and short waves: growth through intensive and extensive search. *Econometrica* 58:1391–409
- Klein N, Rady S. 2011. Negatively correlated bandits. *Rev. Econ. Stud.* 78:693–732
- Levinthal DA. 1997. Adaptation on rugged landscapes. *Manag. Sci.* 43:934–50
- Li S. 2023. Strategic exploration for innovation. Preprint, arXiv:2108.07218 [econ.TH]
- Liang A, Mu X, Syrgkanis V. 2022. Dynamically aggregating diverse information. *Econometrica* 90:47–80
- Malladi S. 2025. Searching in the dark and learning where to look. Discuss. Pap. 1612, Center for Mathematical Studies in Economics and Management Sciences
- Malladi S, Martinez-Marquina A, Morozov I. 2025. *Space exploration*. Unpublished manuscript
- Malmendier U, Veldkamp L. 2022. *Information resonance*. Unpublished manuscript
- Matheron G. 1971. *The Theory of Regionalized Variables and Its Applications*. Cahiers du Centre de Morphologie Mathématique de Fontainebleau
- McCarty N. 2013. Complexity, capacity, and capture. In *Preventing Regulatory Capture: Special Interest Influence and How to Limit It*, ed. D Carpenter, DA Moss, pp. 99–123. Cambridge University Press
- McCarty N. 2017. The regulation and self-regulation of a complex industry. *J. Politics* 79:1220–36
- Milgrom P. 1981. Good news and bad news: representation theorems and applications. *Bell J. Econ.* 12:380–91
- Oliveros S. 2025. Active product development. Preprint, arXiv:2507.15438 [econ.TH]
- Rasmussen CE, Williams CKI. 2006. *Gaussian Processes for Machine Learning*. MIT Press
- Rue H, Held L. 2005. *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall/CRC
- Schulz E, Speekenbrink M, Krause A. 2018. A tutorial on Gaussian process regression: modelling, exploring, and exploiting functions. *J. Math. Psychol.* 85:1–16
- Urgun C, Yariv L. 2021. Constrained retrospective search. *AEA Pap. Proc.* 111:549–53
- Urgun C, Yariv L. 2025. Contiguous search: exploration and ambition on uncharted terrain. *J. Political Econ.* 133:522–67
- Weitzman ML. 1979. Optimal search for the best alternative. *Econometrica* 47:641–54
- Williams CKI. 1998. Prediction with Gaussian processes: from linear regression to linear prediction and beyond. In *Learning in Graphical Models*, ed. Jordan MI, pp. 599–621. Kluwer Academic Publishers
- Wong YF. 2025. Forward-looking experimentation of correlated alternatives. *Theor. Econ.* 20:883–909