The Simple Economics of Discovery

Martino Banchio* Christoph Carnehl[†]

Abstract

We propose a formal theory to capture the tension between problem-driven and solution-driven search in innovation and science. We model discovery as a two-dimensional match between novel ideas and technical methods. Our framework disciplines empirical evidence by generating sharp predictions about the optimal "research mix" and its response to the economic environment. We characterize how funding shocks, collaboration, and new technologies alter the balance between exploring breadth of questions and mastering depth of solution methods. Our results rationalize heterogeneous findings on the direction of innovation and provide a structural basis for empirical analysis.

1 Introduction

Innovation is often characterized as the result of a successful match between a problem and a solution. This duality is frequently characterized in the qualitative literature, framed as the tension between "technology-push and market-pull" (Di Stefano et al. (2012)) or "solution-driven and problem-driven" innovation (Kruger and Cross (2006)). Sometimes innovators begin with a clear vision of the product they want to market, and persist until the right technique is found. For example, the idea of a durable, safe light source was obvious, but Edison's lightbulb took 6000 trials with different materials and structures.¹ Conversely, other innovations begin with a novel method, and the question is how to apply

^{*}Università Bocconi, IGIER, and Google Research. Email: martino.banchio@unibocconi.it

[†]Università Bocconi and IGIER. Email: christoph.carnehl@unibocconi.it

¹https://www.edisonmuseum.org/practical-incandescent-light-bulb

this technology. Theodore Maiman described his invention of the LASER as "a solution looking for a problem." Only later did researchers match his technology to the most disparate applications in medicine, telecommunications, and consumer electronics.

While the necessity of this match is well-understood conceptually, formal economic models of search have largely abstracted away from this two-dimensional structure. Standard models typically view discovery as a unidimensional process. The canonical model of search in Weitzman (1979) entails a researcher sampling from a distribution to find a value above a certain threshold. In these frameworks, discovery is about finding the "best" option. However, as the examples above as well as the large innovation literature illustrate, discovery is often not about intensity of search along one dimension, but about alignment across two: ideas (problems) and methods (solutions). An idea without a method remains unrealized; a method without an application remains unused. Researchers therefore face a fundamental resource allocation trade-off: should they devote effort to exploring the breadth of potential ideas, or to mastering the depth of technical methods?

In this paper, we bridge the gap between the qualitative literature on problem-solution fit and formal search theory. We develop a tractable model that recasts discovery as a two-dimensional search problem with a matching structure. We represent the fit between ideas and methods, known as the *epistemic landscape* in philosophy of science, as a continuous function. We impose a Brownian prior over this landscape to capture the correlation between similar ideas and their solution methods. A researcher, constrained by her available resources, chooses a "research mix": how broadly to explore ideas and how many methods to master. Discovery occurs only if at least one explored idea can be successfully matched with a mastered method.

Our primary contribution is to provide a simple theoretical framework that generates sharp, testable empirical predictions about the matching approach to search. We first show that the probability of discovery induces convex preferences over ideas and methods. This allows us to map the complex problem of innovation search into the canonical framework of consumer theory. Building on this representation, we derive several predictions regarding how the optimal research mix responds to the economic environment and the epistemic landscape. First, we characterize when idea-driven or method-driven search is optimal. Second, we

find that ideas and methods are substitutes: an abundance of ideas reduces the marginal value of mastering additional methods, and vice versa. Third, we characterize the income expansion path of a researcher: methods behave as "normal goods" while idea exploration can behave as an "inferior good". This implies that expanding a researcher's capabilities or budget may lead them to shift their focus away from generating ideas while deepening their technical mastery.

By leveraging these connections to standard consumer theory, our model offers insights into current debates in the innovation literature. We show how grants and funding structures can alter the direction of research. For instance, our results suggest that unconditional grant increases may shift researchers toward method-heavy, narrower idea portfolios, speaking to heterogeneous empirical findings on the impact of funding on novelty (Myers and Tham, 2023). Furthermore, we find that if a researcher can flexibly invest in improving their skill at generating ideas versus mastering methods, it is optimal to specialize in the skill they already employ more intensively.

Building on these individual-level insights, we study how research collaborations shape discovery, and how such collaborations form endogenously.² As in trade theory, researchers allocate tasks within a team according to comparative advantage. We find that despite synergies from specialization, the production function of teams is submodular in researcher ability. This identifies a specific market failure: when teams form endogenously, stable matches are positively assortative, leading to the "superstar teams" documented in the literature (Ahmadpoor and Jones, 2019). However, the would-be welfare-optimal matching is negatively assortative. Finally, we analyze the impact of directed technical change. Technologies that lower the cost of mastering methods (e.g., rapid prototyping, simulation software) unambiguously increase method mastery. However, technologies that lower the cost of exploring ideas (e.g., Generative AI and Large Language Models) have ambiguous effects, potentially causing method-focused researchers to explore more ideas while inducing idea-focused researchers to rely increasingly on established methods.

A key appeal of our framework is its ability to guide empirical work. The optimal research mix we characterize has straightforward empirical analogues. The

²The increasing role of teams in research and innovation is widely documented and discussed, for example, in Azoulay (2019) and Jones (2021).

"distance" of questions and methods from the existing stock of knowledge can be quantified using semantic similarity or reference network measures (see for example (Luo et al., 2022) for a recent contribution on the measurement of questionand method-novelty). Further, our comparative statics allow for the identification of field-specific parameters through the variation in research mixes across disciplines, and the identification of researcher skills through variation within fields. With these quantifications, we offer a principled approach to answering many important innovation policy quesitons. By formalizing the search for a match, our model provides a unified lens through which to view individual research strategy, team formation, and the impact of new technologies on the rate and direction of innovation.

1.1 Related Literature

Our work relates to the literature on search following Weitzman (1979). While Weitzman (1979) studies the search among independent options, Callander (2011) introduces a tractable search model with correlated options by introducing a spatial search dimension and correlation through an unknown and continuous mapping from locations to outcomes. In particular, Callander (2011) assumes that this mapping is determined by the realization of a Brownian path. Several papers build on this model, extending it in different directions. Garfagnini and Strulovici (2016) consider forward-looking agents, Callander and Matouschek (2019) study the role of risk aversion, and Callander et al. (2025) add a dimension of horizontal differentiation. Urgun and Yariv (2025) study a model in which the agent chooses the speed at which a Brownian path is discovered. Malladi (2025) departs from the Brownian assumption and instead assumes that the outcome-mapping is a Lipschitz-continuous function. All of these models share that a searcher seeks to find a one-dimensional location with a high realization of the mapping.

We depart from the literature by considering a different economic problem. Rather than searching for a peak, our researcher searches for a match in two dimensions.³ Our departure changes the economic problem the researcher faces.

³Other papers drawing on related modeling tools, consider different search targets and different processes; Bardhi and Callander (2026) provide an excellent survey of the literature. For example, Bardhi (2024) and Bardhi and Bobkova (2023) consider sampling points to learn a summary statistic of the process. While the former employs general Gaussian processes, the latter

Any match that is discovered delivers the same value, but the search might fail to generate a match. To maximize the probability of discovering a match, the researcher chooses a two-dimensional search region for the idea-method mapping to pass through.

In this sense, our model is closer to Banchio and Malladi (2025) and Carnehl and Schneider (2025). In the former, a researcher faces a fixed maximal value of the process and searches for a location generating this realization. In the latter, a researcher chooses a fixed location and searches for the realization of the process at the given location. Our searcher is more flexible, attempting to identify an arbitrary match between location and realization.

Scholarship in innovation management and industrial design has long posited that search occurs simultaneously across two distinct dimensions, the problem and the solution method. Maher and Poon (1996) and Maher and Tang (2003) characterize this as "co-evolutionary design," analyzing the iterative heuristics researchers use to align the problem space with the solution space. We depart from such cognitive descriptions by modeling this multidimensional search within a rational choice framework, delivering sharp, testable implications regarding how incentives and resource constraints shape the direction of innovation. Furthermore, our framework unifies disparate strands of the literature, from the solution-first approach to innovation of Gruber et al. (2013) and Von Hippel and Von Krogh (2016), to the contrasting view of Cyert and March (1963) which introduces problem-first innovation as the concept of "problemistic search".

In the economics of innovation and science, several papers address the direction of innovation based on researchers' question-choice incentives.⁴ Our paper instead distinguishes between two alternative but related innovation search margins, and thus, directions of innovation. The tradeoff faced by our researcher adds a novel perspective to the studies on the role of grants in shaping discovery. Our framework can rationalize frequently observed patters in the literature (see Azoulay and Li, 2020; Carnehl et al., 2025, for an overview) and provide novel testable implications.

assumes an Ornstein-Uhlenbeck process. Callander and Clark (2017) studies search for roots of a Brownian motion. However, none of these papers captures the notion of a two-dimensional search for a match.

⁴See, for example, Bryan and Lemus (2017); Bobtcheff et al. (2017); Hopenhayn and Squintani (2021); Hill and Stein (2025); Hill et al. (2025).

2 Model

We develop a model of scientific discovery in which a researcher chooses how to allocate limited time between exploring new ideas and mastering technical methods to produce a publishable finding. The core of the model is an uncertain *epistemic landscape* that connects ideas to the methods required to solve them.

2.1 The Epistemic Landscape

The space of potential research ideas is represented by the positive real line, $x \in \mathbb{R}_+$. Each idea x is associated with a unique solution method, $y(x) \in \mathbb{R}$, that is required for its successful implementation. We normalize the current frontier of knowledge to lie at the origin, x = 0 with solution method y(0) = 0.

The link between ideas and solution methods is inherently uncertain. We model their relationship as a driftless Brownian motion with volatility parameter $\sigma^2 > 0$, $(y(x))_{x \ge 0}$, mapping each idea x to its solution method y(x). Given the knowledge frontier (0,0), for any idea x, the method y(x) is a random variable following a normal distribution with mean $\mathbb{E}[y(x)] = 0$ and variance $\text{var}(y(x)) = \sigma^2 x$.

The volatility parameter of the Brownian motion $\sigma^2 > 0$ quantifies the fundamental ruggedness of the epistemic landscape. A higher volatility corresponds to a more complex, unpredictable field, where a small change in ideas being pursued is more likely to lead to a drastic change in the required solution method.

2.2 The Researcher's Problem

A researcher seeks to make a discovery on the epistemic landscape. Her search strategy defines a set of ideas to investigate and a set of methods to master. Formally, she chooses an interval of ideas, $L \subset \mathbb{R}_+$, and an interval of methods, $H \subset \mathbb{R}$. A discovery occurs if she investigates an idea $x \in L$ for which she masters the solution method y(x), i.e., $y(x) \in H$.

The researcher's payoff is binary: she receives a utility of one if a discovery is made and zero otherwise. Her objective is to maximize the probability of discovery subject to two key constraints, a resource constraint and a novelty constraint.

⁵Jovanovic and Rob (1990) motivate the use of the driftless Brownian motion in the context of innovation through a natural axiomatization of the search process.

First, the researcher faces a resource constraint, which could represent time, funding, or cognitive effort. Let $\ell = |L|$ and h = |H| denote the breadths of the chosen idea and method intervals, respectively. We refer to ℓ as the researcher's investment in *ideas* and h as her investment in *methods*. These investments are costly. The per-unit cost of ideas is $c_{\ell} > 0$ and the per-unit cost of methods is $c_{h} > 0$. Given a total budget B > 0, the researcher's choice (L, H) must satisfy the budget constraint:

$$c_{\ell}\ell + c_h h \le B. \tag{1}$$

Second, the researcher faces a novelty constraint, as the market for discoveries typically does not reward incremental discoveries.⁶ To implement this constraint, we assume that the researcher obtains a reward only for discoveries on ideas with distance of at least $\Delta > 0$ to the current knowledge frontier.

The problem that the researcher solves is

$$\max_{L \subseteq \mathbb{R}_+, H \subseteq \mathbb{R}} \quad \mathbb{P}(\exists x \in L \text{ s.t. } x \ge \Delta \text{ and } y(x) \in H)$$
 (2)

s.t.
$$c_{\ell}\ell + c_h h \le B$$
. (3)

We can simplify the researcher's optimization problem using two general properties of the epistemic landscape. As uncertainty about methods increases in the distance to the knowledge frontier, it is optimal to explore those ideas closest to the frontier subject to the novelty requirement. Therefore, the researcher chooses an idea interval of the form $L = [\Delta, \Delta + \ell]$. Furthermore, for any given idea, the associated method is distributed unimodally and symmetrically around zero. The researcher optimally chooses a methods interval that is symmetric around the benchmark method y(0) = 0, that is, H = [-h/2, h/2].

Lemma 1. The researcher's optimal choice has the following properties:

- 1. The optimal H is of the form [-h/2, h/2].
- 2. The optimal L is of the form $[\Delta, \Delta + \ell]$.

Omitted proofs can be found in the Appendix. These arguments reduce the researcher's problem from choosing the location and size of two intervals to simply

⁶For example, new patents have to be sufficiently different from existing patents. Academic papers have to be sufficiently distinct from the prior literature, etc.

choosing their optimal breadths, $(\ell, h) \in \mathbb{R}^2_+$:

$$\max_{\ell \in \mathbb{R}_+, h \in \mathbb{R}_+} \mathbb{P}\left(\exists x \in [\Delta, \Delta + \ell] \text{ s.t. } y(x) \in [-h/2, h/2]\right)$$
(4)

s.t.
$$c_{\ell}\ell + c_h h \le B$$
. (5)

3 Convex Tastes for Discovery

The key technical result of this paper is that the preferences of the researcher over ideas and methods are convex.

Theorem 1. The preference relation \geq over the set \mathbb{R}^2_+ of methods and ideas defined by

$$(\ell, h) \ge (\ell', h') \iff F(\ell, h) \ge F(\ell', h'),$$

is strictly convex.

Denote by

$$A(\ell,h) := \left\{ \exists x \in [\Delta, \Delta + \ell] \text{ s.t. } y(x) \in [-h/2, h/2] \right\}$$

the event that the realization of the epistemic landscape y(x) "passes through" the researcher's choice of ideas and methods. The probability of such event, $F(\ell,h)$: = $\mathbb{P}(A(\ell,h))$, is the researcher's utility function.

Theorem 1 allows us to represent the researcher's problem as a canonical consumer problem with two "goods", ideas and methods, subject to a budget constraint. Additionally, simple properties of the epistemic landscape guarantee that the preferences are strictly monotonic.

Proposition 1. The marginal utility of h and ℓ is strictly positive, i.e., $F_h, F_\ell > 0$.

Proof. By definition, $A(\ell, h + \delta) \supset A(\ell, h)$ for $\delta > 0$. Since the Brownian motion has positive mass on any open set of continuous functions, the strict inclusion implies that $\mathbb{P}(A(\ell, h + \delta)) > \mathbb{P}(A(\ell, h))$, thus $F(\ell, h + \delta) - F(\ell, h) > 0$. Dividing by δ and taking limits proves that $F_h(\ell, h) > 0$. The same logic applies to the derivative with respect to ℓ .

The representation as a well-behaved canonical consumer search problem allows provides us with several immediate yet useful corollaries and analogies that we will employ throughout.

Corollary 1. For any budget B and costs c_{ℓ} , $c_h > 0$, there exists a unique solution to the researcher's problem.

With this representation at hand, we can characterize the optimal choice of ideas and methods by studying the properties of the marginal rate of substitution (MRS) between the two goods. This will be the subject of our analysis in Section 4 and serve as the building block for our more applied insights in Section 5. The rest of this section instead outlines the proof of Theorem 1 and the corresponding intuition.

3.1 Sketch of Theorem 1's Proof

We break down the proof in steps. First, instead of proving the convexity of preferences directly, we will work with the associated utility function, $F(\ell, h)$.

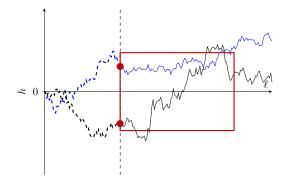
Lemma 2. If $F(\ell,h)$ is jointly strictly concave, the induced preference relation > is strictly convex.

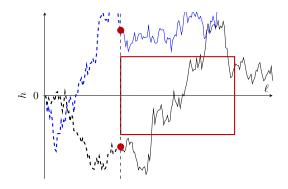
Proof. Joint strict concavity implies strict quasi-concavity of the utility function, which is equivalent to convexity of the induced preference relation. \Box

Second, we break down the utility function $F(\ell,h)$ in the two components that contribute to its value. By the law of total probability, $F(\ell,h)$ is equivalently written as the probability of the event $A(\ell,h)$ conditional on the value $y(\Delta)$ of the landscape at Δ , integrated with respect to the measure μ_{Δ} of $y(\Delta)$:

$$F(\ell,h) = \int_{\mathbb{R}} \mathbb{P}(A(\ell,h) \mid y(\Delta) = z) d\mu_{\Delta}(z).$$

The event $A(\ell,h)$ has conditional probability one when the value $y(\Delta)$ falls within the boundary [-h/2,h/2] of the rectangle $L \times H$. This is represented in the left panel of Figure 1. The event $A(\ell,h)$ instead has conditional probability less than one when the value $y(\Delta)$ falls outside of the boundary H. Simply put, the landscape behaves as the path of a stochastic process with the same law as the original stochastic process but started at value $y(\Delta)$. The conditional probability that the original process hits the rectangle is therefore the probability that the





- (a) The two landscapes represented in this panel have a value $y(\Delta)$ between -h/2 and h/2, so they always hit the red rectangle.
- (b) The landscape with $y(\Delta) > h/2$ does not hit the rectangle. Instead, the landscape with $y(\Delta) < -h/2$ returns to a value $y(x) \in [-h/2, h/2]$ for an $x < \Delta + \ell$.

Figure 1: The two components of the conditional success probability

restarted process enters the region H at any point before ℓ . This is represented in the right panel of Figure 1.

Our Brownian assumption implies symmetry and that $\mu_{\Delta}(z) = \Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right)$. Thus, we can rewrite the utility function as the sum of integrals

$$F(\ell,h) = 2 \int_{-h/2}^{0} d\Phi \left(\frac{z}{\sigma\sqrt{\Delta}}\right) + 2 \int_{-\infty}^{-h/2} \mathbb{P}\left(A(\ell,h) \mid y(\Delta) = z\right) d\Phi \left(\frac{z}{\sigma\sqrt{\Delta}}\right). \tag{6}$$

Third, we compute the probability inside the integral. Consider a value of $y(\Delta)$ below the lower edge of the rectangle, which is -h/2. The probability of ever entering the height band [-h/2, h/2] during the window of length ℓ is exactly the probability that, starting from $y(\Delta)$, the path's supremum over an interval of length ℓ clears the distance to the nearest boundary of the band, i.e., $-h/2-y(\Delta)>0$. Denoting the supremum of a Brownian motion initialized at 0 over the range $[\Delta, \Delta + \ell]$ by M_{ℓ} and by $G_{\ell}(x)$ its distribution function, we can write

$$\mathbb{P}(A(\ell,h) | y(\Delta) = z) = 1 - G_{\ell}(-h/2 - z).$$

With basic calculus, we prove in the appendix that the integrals above are repre-

sented by the simpler integral

$$F(\ell,h) = 2 \int_0^\infty \Phi\left(\frac{h/2 + z}{\sigma\sqrt{\Delta}}\right) dG_{\ell}(z) - 1, \tag{7}$$

where this integral is averaging over "future headroom" (the value of M_{ℓ}) the probability that the value $y(\Delta)$ is within headroom of the barrier -h/2.

Lastly, showing concavity of this integral directly remains challenging because it depends on h and ℓ through a product. Instead, we use a simple quantile-integration trick to represent the last integral as

$$F(\ell,h) = 2 \int_0^1 \Phi\left(\frac{h/2 + Q_{\ell}(p)}{\sigma\sqrt{\Delta}}\right) dp - 1, \tag{8}$$

where $Q_{\ell}(p)$ is the generalized left-inverse of the CDF of M_{ℓ} (or its quantile function). Since the distribution Φ is symmetric, unimodal, and uniformly integrable, a sufficient condition for joint concavity of $F(\ell,h)$ is the joint concavity of the argument $h/2 + Q_{\ell}(p)$ for all $p \in [0,1]$. The advantage of this representation is that the argument is separable and linear in h, so all we are left to check for concavity is that $\frac{\partial^2 Q_{\ell}(p)}{\partial \ell^2} < 0$, or that the quantile function of the supremum is concave in ℓ . The supremum M_{ℓ} is identical in distribution to $\sqrt{\ell} M_1$, so $Q_{\ell}(p) = \sqrt{\ell} Q_1(p)$ which allows us to conclude.

4 Optimal Research Mix

The previous section's representation of our model as a consumer problem allows us to characterize properties of the optimal research mix by leveraging classical results from consumer theory. We will make explicit use of this connection, often referring to the optimal research mix as the *Marshallian demand*. As anticipated by Corollary 1, our Marshallian demand is unique.

As a first step, we obtain a convenient representation of the marginal rate of substitution between methods and ideas. As per canonical consumer theory, the marginal rate of substitution is instrumental in the characterization of the optimal

⁷For all *p* ∈ [0,1], we are composing a strictly concave function $\Phi(\cdot)$ with a positive and concave function $g(\ell,h) = \frac{h/2 + Q_{\ell}(p)}{\sigma \sqrt{\Lambda}}$, which ensures that concavity is preserved.

research mix.

Proposition 2. Let $X = \frac{h/2}{\sigma\sqrt{\Delta}}\sqrt{\frac{\ell}{\ell+\Delta}}$. The marginal rate of substitution (MRS) between h and ℓ is given by:

$$M = MRS_{h\ell} = \frac{F_{\ell}}{F_{h}} = \frac{\sigma\sqrt{\Delta}}{\sqrt{\ell}\sqrt{\ell + \Delta}} \Big(H(X) - X\Big)$$

where H(x) is the hazard rate of a standard normal distribution, $H(x) = \frac{\phi(x)}{1 - \Phi(x)}$.

The MRS relates the marginal effect of an increase in the methods mastered h to an increase in the amount of search ℓ . Additional methods increase the chance that a path with $y(\Delta) > |h/2|$ hits the box by reducing the distance to the methods that the path has to travel over the search length ℓ . Additional ideas increase the search length and thereby raise the dispersion of methods over the ideas pursued. The final expression of the success probability in (8) highlights these respective effects. Both ideas and methods have diminishing returns due to the strict concavity of the standard normal distribution on the positive domain. However, there is a fundamental difference in methods and ideas expansions. Expanding the methods mastered always reduces the distance of realized methods to methods mastered linearly, h/2. Expanding the ideas pursued has intrinsic diminishing returns: Methods disperse over the search length only at a square-root rate, $\sqrt{\ell} Q_1(p)$.

The characterization in Proposition 2 allows us to determine whether the Marshallian demand is a corner solution or is interior.

Proposition 3. For any cost vector (c_{ℓ}, c_h) there exists an income level \overline{B} such that

- 1. If $B \leq \overline{B}$, the Marshallian demand is a corner solution with $h^* = 0$.
- 2. If $B > \overline{B}$, the Marshallian demand is interior, i.e. $\ell^*, h^* > 0$.

Proof. Denote by $\ell(p)$ and h(p) the bundles that belong to the same isoprobability curve of level p. Consider the limit $\lim_{h(p)\to 0} M$ on any isoprobability curve. The MRS in this limit is $\sigma\sqrt{\frac{\Delta}{\ell(p)(\ell(p)+\Delta)}\frac{2}{\pi}}$ where $(\ell(p),0)$ belongs to the isoprobability

⁸Note that the marginal increase in the methods at Δ does not affect the success probability at the margin because paths with $|y(\Delta) - h/2| < \varepsilon$ would hit the rectangle almost surely for any $\ell > 0$ and $\varepsilon \downarrow 0$.

curve. From Proposition 1 we know that $p = F(0,\ell)$ is increasing in ℓ , thus its inverse $\ell(p)$ is increasing in p with $\ell(0) = 0$. The above expression for the MRS is then decreasing in p. Hence, there is always a value \overline{p} such that for all $p < \overline{p}$ the MRS at $(\ell(p),0)$ is larger than c_h/c_ℓ . The expenditure function $e(c_\ell,c_h,p)$ is increasing in p, thus for all $B \le e(c_\ell,c_h,\overline{p})$ the Marshallian demand is $(\ell^*,h^*) = (\ell(p),0)$. Conversely, for all $B > e(c_h,c_\ell,\overline{p})$ the Marshallian demand will have $h^* > 0$.

Finally, note that the limit $\lim_{\ell(p)\to 0} M$ is infinite. The indifference curves are infinitely steep at the intersection with the h axis, thus there can be no corner solution with $\ell^* = 0$.

Leveraging Proposition 2, we can write down a system of equations that fully characterizes the Marshallian demand if the optimal choice is interior, i.e., if $B > \overline{B}$:

$$\begin{cases}
M(\ell^*, h^*) = \frac{c_{\ell}}{c_h} \\
c_{\ell} \cdot \ell^* + c_h \cdot h^* = B.
\end{cases}$$
(9)

Whenever the choice is a corner solution, we immediately have $(h^* = 0, \ell^* = B/c_\ell)$. An analytical characterization of the Marshallian demand beyond this implicit representation is unavailable, but all our insights will follow from borrowing methods of consumer theory and applying them to our setting.

Remark. The characterization of the optimal research in (9) delivers immediate but interesting comparative statics. These comparative statics can guide empirical work in identifying structural parameters. For a given field, represented by the underlying epistemic landscape with volatility σ and novelty constraint Δ , ceteris paribus variations in researchers' relative skills, c_ℓ/c_h , lead to different search approaches that are relatively more idea- or method-driven. Moreover, for given researcher characteristics, the optimal research mix responds to the properties of the research field. In particular, a more rugged epistemic landscape induces a research mix that is more idea-intensive.

4.1 Income Effects

To better understand the optimal research mix, we study how the researcher's choice varies with her budget.

Proposition 4. Fix (c_{ℓ}, c_h) .

- 1. Methods are a normal good; that is, a researcher with a higher budget will employ more methods.
- 2. Ideas are a normal good when the budget is weakly lower than \overline{B} from Proposition 3. Instead, when $B > \overline{B}$, ideas are an inferior good; that is, a researcher with a higher budget will try less ideas.

Proof. First, when $B < \overline{B}$ the Marshallian demand is in a corner solution with $h^* = 0$. Naturally then consumption of both goods weakly increases as long as the Marshallian remains in a corner solution.

Let us then focus on the case $B > \overline{B}$. Since the utility function is twice differentiable, we only need to sign the derivatives $\partial_B h^* = \frac{\partial h^*}{\partial B}$ and $\partial_B \ell^* = \frac{\partial \ell^*}{\partial B}$. By taking the total derivative of the system (9), we get the system of linear equations

$$\underbrace{\begin{pmatrix} M_h & M_\ell \\ c_h & c_\ell \end{pmatrix}}_{A} \begin{pmatrix} \partial_B h^* \\ \partial_B \ell^* \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

The solutions to this system are given by Cramer's rule as

$$\begin{cases} \partial_B h^* = -\frac{M_\ell}{\det(A)} \\ \partial_B \ell^* = \frac{M_h}{\det(A)}. \end{cases}$$

Consider the sign of $\partial_B \ell^*$. First, note that Theorem 1 and Proposition 1 together imply that the indifference curves are downward sloping, i.e.

$$\left. \frac{dM}{d\ell} \right|_{F(h,\ell)=p} = M_{\ell} + M_h \frac{dh}{d\ell} = M_{\ell} - M_h \frac{c_{\ell}}{c_h} < 0.$$

Since $\det(A) = -c_h \cdot \frac{dM}{d\ell}\Big|_{F(h,\ell)=p}$, the determinant is positive. Thus, $\operatorname{sign}(\partial_B \ell^*) = \operatorname{sign}(M_h)$. Proposition 2 gives us that $M_h = \frac{1}{2(\ell+\Delta)}(H'(X)-1)$.

By item 3 of the technical Lemma 3 in the appendix, $M_h < 0$, so $\partial_B \ell^* < 0$, and ideas are an *inferior* good. If one good is inferior, Proposition 1 implies that the other must be normal, so $\partial_B h^* > 0$.

The Marshallian demand for varying budgets is represented by the incomeallocation curve depicted in Figure 2.

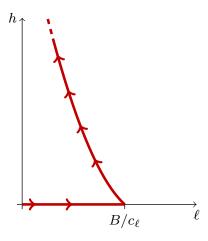


Figure 2: The income-allocation curve. The Marshallian demand moves along the curve in the direction of the arrows as income increases.

While the inferiority of ideas will be at the core of many results in the remainder of the paper, this comparative static is interesting in its own right. The model's prediction that ideas are an inferior good when researchers have sufficiently large budgets implies that reductions in effective "research budgets", such as temporary increases in time constraints, might shift output toward more idea-intensive work. In practice, this suggests, for example, that new parents, who experience a negative shock to available research time, may produce work that is more distant from the current stock of knowledge and produced with methodology closer to the state-of-the-art. Intuitively, when time is scarce, researchers opt for exploration of a broader set of ideas with familiar methods, instead of investing in new methodologies. Thus, life shocks that mimic income effects in our framework can be expected to generate measurable shifts in the novelty of produced knowledge. Such prediction could be tested by measuring the semantic or citation distance of new parents' subsequent papers from the existing frontier, exploiting variation in

⁹Analogous predictions apply to researchers winning a grant, allowing them to buy out teaching hours, for example. Our findings can reconcile the conflicting empirical evidence regarding the effect of winning a grant on the novelty of grant winners' research through variation in the pre-grant budget, or, empirically easier to observe, the revealed pre-grant research mix.

a field's epistemic volatility σ .

Remark. The inferiority of ideas is related to the rate of dispersion of methods. As further ideas grow distant from the current stock of knowledge, methods disperse at a square-root rate. A process whose methods do not disperse, for example the mean-reverting Ornstein-Uhlenbeck process that remains concentrated around the 0-method axis, intrinsically increases the returns to ideas.¹⁰

4.2 Price Effects

A second natural comparative static to perform is the one with respect to costs, that is, with respect to the researcher's costs that relate to her research skills. How the optimal research mix varies with the cost of methods is immediate from the so-called "law of demand":

Proposition 5. An increase in the cost of methods c_h leads to

- 1. Weakly fewer methods, $\partial_{c_h} h^* \leq 0$
- 2. Weakly more ideas, $\partial_{c_h} \ell^* \geq 0$.

Proof. If the Marshallian is in a corner solution, both items are obvious. Consider instead the case that the Marshallian is interior. The first item follows from the fact that the own-price effect of a Hicksian demand is always negative and the Slutsky equation. The second item follows again from the Slutsky equation: the substitution effect is positive because two goods are always net substitutes, and the income effect is negative because ideas are inferior at interior solutions.

In fact, we can prove that ideas and methods are *gross substitutes* by showing the analogous result for increases in the cost of ideas c_{ℓ} .

Proposition 6. An increase in the cost of ideas c_{ℓ} leads to

1. Weakly fewer ideas, $\partial_{c_{\ell}} \ell^* \leq 0$

¹⁰In fact, there is a space-time transform that reduces an Ornstein-Uhlenbeck process to a Brownian, and under that space-time transform the rectangle of ideas and methods is no longer a rectangle, but a strip growing at a square-root rate with ideas. Thus, the rate of dispersion seems to be the correct object to govern the substitution patterns, but a formal connection is beyond the scope of this paper.

2. Weakly more methods, $\partial_{c_{\varrho}} h^* \geq 0$.

Proof. The proof is again trivial if the Marshallian is in a corner solution. Instead, if it is interior, we can take the total derivative of the system (9) with respect to c_{ℓ} and use Cramer's rule to get that $\operatorname{sign}(\partial_{c_{\ell}}\ell^*) = \operatorname{sign}(-\ell M_h - 1)$. Plugging in the formula for M_h computed earlier, we find that the sign is

$$\frac{\ell}{2(\ell+\Delta)}(1-H'(X))-1<(1-H'(X))-1<0$$

where we used again item 3 in Lemma 3. Thus, $\partial_{c_{\ell}} \ell^* < 0$. The sign of $\partial_{c_h} h^*$ is equivalent to the sign of $M + \ell M_{\ell}$. Algebra that we relegate to the Appendix shows that

$$M + \ell M_{\ell} = \frac{K\Delta}{2(\ell + \Delta)^2} \frac{d}{dX} [X(H(X) - X)].$$

Thus, the sign of $\partial_{c_{\ell}}h^*$ is positive if and only if X(H(X)-X) is increasing, which is guaranteed by item 4 of Lemma 3.

The gross substitute condition tells us that there are strong substitution effects in research efforts. A decrease in the cost of ideas (methods) will invariably reduce methods (ideas) in the optimal research mix.

4.3 Indirect Utility

Before turning to applications, we leverage properties of the indirect utility function to provide additional insights on the researcher's optimal decision that we can leverage for various applications later on. Recall that the indirect utility function is defined as

$$v(c_{\ell}, c_h, B) = F(\ell^*(c_{\ell}, c_h, B), h^*(c_{\ell}, c_h, B)).$$

In the following we work out as an example the problem of a researcher with an initial skill vector (c_{ℓ}, c_h) who chooses how to invest in education. Should she try to increase their analytical skills (i.e., lower c_h), or should she invest in her idea generation skills (i.e., lower c_{ℓ})? More formally, suppose that the researcher can invest an amount $I \in (0, \overline{I})$ in her skills.¹¹ In particular, she can freely allocate

¹¹We set $\overline{I} := \min\{c_h, c_\ell\}$ such that both ideas and methods remain strictly costly.

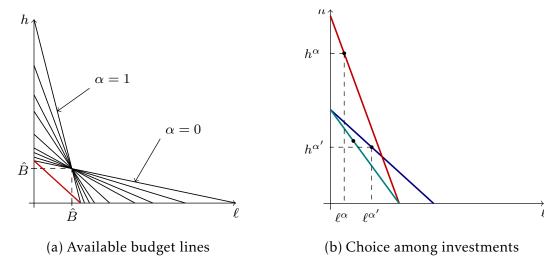


Figure 3: *Skill investments*.

the investment I between idea-generation and method-learning skills, that is, she chooses $\alpha \in [0,1]$ such that

$$c_h^{\alpha} = c_h - \alpha I, \qquad c_{\ell}^{\alpha} = c_{\ell} - (1 - \alpha)I.$$

Each choice of α generates a new budget line with a different slope, but any two such budget lines cross in the same interior point on the 45-degree line, $(h^I = \hat{B}, \ell^I = \hat{B})$ with $\hat{B} := \frac{B}{c_h + c_\ell - I}$.

The optimal choice of the researcher is "bang-bang": she will either invest exclusively in idea-generation or exclusively in method-generation.

Proposition 7. Fix (c_{ℓ}, c_h) and B.

- 1. The researcher either chooses $\alpha = 1$ or $\alpha = 0$.
- 2. If $h^* > \ell^*$, then the researcher will choose $\alpha = 1$. Otherwise, the researcher will choose $\alpha = 0$.

Proof. To see why the first item holds, observe that the researcher chooses α to maximize her indirect utility function:

$$\max_{\alpha} v(c_{\ell} - (1 - \alpha)I, c_h - \alpha I, B).$$

We know from consumer theory that the indirect utility function is quasi-convex in its arguments, and hence its maximizers must be extreme points.

The second part of the proposition tells us that if the researcher's initial skills were sufficiently in favor of methods (ideas), she will invest in methods (ideas) alone. To prove this result, recall Roy's identity relating the indirect utility gain of a cost reduction to the Marshallian demand:

$$-\frac{\partial v}{\partial c_{\ell}} = \ell^* \frac{\partial v}{\partial B}.$$

Then it is clear that the largest gain comes from the skill that is currently the most utilized. \Box

To gain some intuition, it is useful to recall the reason for quasi-convexity of the indirect utility function. Take the two "extreme" budget sets, with $\alpha=1$ and $\alpha=0$ respectively. Any other budget line with $\tilde{\alpha}$ can be expressed as a convex combination of the two extreme budget lines, and any bundle that is feasible under the budget line $\tilde{\alpha}$ must be feasible under either of the extreme budget lines, as can be seen in Figure 3. Thus, it is immediate that any optimal research mix that is feasible under a non-extreme budget set is feasible under an extreme one, hence the result.

Which of the two corner solutions is optimal depends on the initial relative cost c_ℓ/c_h , the parameters of the epistemic landscape σ , and the novelty constraint Δ . The more favorable for methods the relative cost is initially, the more valuable are investments in methods-skills. However, the more uncertain the underlying epistemic landscape, or the stronger the novelty constraint (higher σ or higher Δ), the more valuable are investments in idea-generation skills.

These insights from Proposition 7 are again not only useful for more elaborate applications but interesting in their own right. For example, we may expect that researchers in fields with more complex epistemic landscapes have greater incentives to invest in idea-generation skills than researchers in less complex, more predictable fields.

An alternative interpretation of these results is the optimal design of curricula. Consider a school designs a new graduate program. Should it focus the curriculum more on technical method skills or rather on more creative question-asking

skills? It turns out that the answer is field-dependent and dependent on the students' background. In fields with more volatile epistemic landscapes, the optimal curriculum should focus relatively more on asking questions and generating ideas rather than on teaching more methods.

5 Collaboration

In this section, we leverage our characterization of the optimal research mix and its comparative statics to glean some insights into one fundamental aspect of research, collaboration. Collaboration is an increasingly prevalent feature of innovation and research; see, for example, the survey in Jones (2021). The nature of the team problem allows us to draw parallels between researchers allocating their limited resources in collaborative settings and households making consumption and leisure decisions. We first study how tasks are allocated within a research team and then move on to investigate how and when such teams form. Building on these insights, we show how an organization may structure research collaboration optimally. Finally, we add to the recent discussions on the effects of Large Language Models on cognitive labor by analyzing collaboration with technological tools such as AI assistants.

5.1 Task Allocation in Teams

Consider a team of two researchers, ordered by their comparative advantage on idea-generation. In particular, suppose that researcher i=1,2 has a budget B^i and costs c_ℓ^i, c_h^i such that $c_\ell^1/c_h^1 < c_\ell^2/c_h^2$. The team faces a joint optimization problem reminiscent of the intra-household allocation problem in labor economics. Denote by $\ell^T = \ell^1 + \ell^2$ and $h^T = h^1 + h^2$ the team's aggregate choice of ideas and methods. The team's production function is simply the indirect utility with respect to the joint budget set

$$\mathcal{X} = \left\{ (\ell^T, h^T) \in \mathbb{R}_+^2 : c_\ell^i \ell^i + c_h^i h^i \le B^i \text{ for } i \in \{1, 2\} \right\}. \tag{10}$$

The next proposition shows that the optimal division of tasks within the team is based on comparative advantages and leads to specialization.

Proposition 8. The team-optimal research mix (ℓ^T, h^T) has the following properties.

- (i) One researcher specializes in the task for which she has a comparative advantage; that is, the optimal task allocation is such that either $h^1 = 0$ or $\ell^2 = 0$.
- (ii) If researcher 1's autarky solution is interior, then researcher 2 will provide only methods; that is, $\ell^2 = 0$ and $h^2 = B^2/c_h^2$.
- (iii) The team's budget line is

$$h = \begin{cases} \frac{B^2}{c_h^2} + \frac{B^1}{c_h^1} - \frac{c_\ell^1}{c_h^1} \ell, & \text{if } \ell \leq \frac{B^1}{c_\ell^1} \\ \frac{B^2}{c_h^2} - \frac{c_\ell^2}{c_h^2} (\ell - \frac{B^1}{c_\ell^1}), & \text{if } \ell > \frac{B^1}{c_\ell^1}. \end{cases}$$

The proof of the first and the third item is a direct application of the well-known comparative advantage logic, as in Ricardo (1821). The second item follows as a consequence of the comparative statics in the previous section, in particular of income shifts. If the researcher with comparative advantage on generating ideas was already mixing in some methods (this is case (ii) above), any feasible bundle in the team's budget is also feasible under the budget line of that researcher when shifted by B^2/c_h^2 . By inferiority, the optimal choice on this shifted budget line involves fewer overall ideas, hence researcher 2, who has a comparative *dis* advantage on idea generation, will avoid generating ideas altogether. This is true irrespective of absolute advantages, hence even a highly skilled researcher (who has low c_ℓ and c_h) may focus on methods if collaborating with a teammate with comparatively better idea generation skills.

Proposition 8 highlights how the formation of teams endogenously unlocks synergies through specialization. While the sum of the researchers' autarky choices is feasible for the team, it will generally be suboptimal, as researchers can reallocate tasks within the team and thereby achieve a higher probability of discovery.

5.2 Sorting in Teams

Collaboration being such a fundamental component of research, a natural question is whether private incentives to collaborate implement a socially optimal team

composition. In this section, we investigate this question by considering the following scenario. There is a continuum of researchers who face the same relative costs c_ℓ/c_h , but they vary in their budget B_i . Equivalently, researchers vary in their absolute advantage along both dimensions of research inputs at the same rate. We identify individuals in this collection by their budget and index it with $\alpha \in [0,1]$. Hence, we represent the collection of researchers by their budgets as $\mathcal{B} = \{B_\alpha\}_{\alpha \in [0,1]}$ with higher values of α corresponding to higher budgets (higher ability). Similarly, there is a second collection that has relative cost \hat{c}_ℓ/\hat{c}_h . We denote this collection by $\hat{\mathcal{B}} = \{\hat{B}_\beta\}_{\beta \in [0,1]}$. Suppose that both groups have the same total mass and that the budgets are distributed according to the strictly increasing cumulative distribution functions G and \hat{G} , respectively.

We are interested in how researchers from the first collection will match with researchers from the second to form teams. For simplicity, we assume that each team will work on a separate project, so the problem reduces to a classic matching problem in the spirit of Becker (1973). In particular, we assume what Becker calls a "rigid" split of the team surplus, where the agent from collection \mathcal{B} always receives a fraction t of the team's surplus and the agent from collection $\hat{\mathcal{B}}$ receives fraction 1-t instead. Again, the production function of the team is the indirect utility with respect to the team's joint budget. Since the cost ratios are fixed in this section, we simply denote the production function by $v(B_{\alpha}, \hat{B}_{\beta})$. We show in the appendix that this indirect utility is submodular and that its marginals are positive (Lemma 5).

A matching is a bijection $\tau \colon \mathcal{B} \to \hat{\mathcal{B}}$, equivalently represented by the corresponding bijection on the indices of the collections, denoted by $\tau_a \colon [0,1] \to [0,1]$. We refer to the matching τ such that $\tau_a(\alpha) = \hat{G}^{-1}(G(\alpha))$ as the positive assortative matching, and the matching τ such that $\tau_a(\alpha) = \hat{G}^{-1}(1 - G(\alpha))$ as the negative assortative matching. A matching is welfare-optimal if it is the solution

$$\tau^* \in \arg \sup_{\tau} \int_0^1 v(B_{\alpha}, \tau(B_{\alpha})) dG(\alpha).$$

Instead, a matching τ is individually optimal, or *stable*, if, for all $\alpha \in [0,1]$,

$$v(B_{\alpha}, \hat{B}_{\beta}) > v(B_{\alpha}, \tau(B_{\alpha}))$$
 for some $\hat{B}_{\beta} \in \hat{\mathcal{B}}$

implies that

$$v(\tau^{-1}(\hat{B}_{\beta}), \hat{B}_{\beta}) > v(B_{\alpha}, \hat{B}_{\beta}).$$

That is, a stable sorting respects individual incentives by avoiding blocking pairs.

The following is an immediate corollary of the submodularity of the indirect utility that we establish in Lemma 5 in Appendix A.

Proposition 9. The negative assortative matching is welfare-optimal. The positive assortative matching is individually-optimal.

The results identifies a market failure in the formation of research teams. If left up to individual incentives, stable teams that form will generally not maximize total welfare. The researcher's production function in our setting exhibits decreasing returns to scale, which severely limits the social benefit of "superstar" teams, documented in the literature (see, for example, Ahmadpoor and Jones, 2019, who identify positive assortative matching across scientific fields as well as in patenting) in settings with decentralized matching procedures. A benevolent designer could reshuffle teams to increase total surplus, by matching the lowest ability researcher in one group with the highest ability researcher in the other. Anecdotally, this corresponds to the formal matching processes found in private sector research units, where the newest engineers are paired with the most knowledgeable mentors.

5.3 Hiring

We leverage our simple model of how teams may form in equilibrium as well as under the influence of a benevolent planner to speak to the question of whether teams should form. So far we have treated the team's joint budget set as the Minkowski sum of the budget sets of the individual researchers, but running a team entails administrative and managerial costs. We model these costs as a reduction in the feasible choices of the manager, who now faces a tradeoff - whether to hire a researcher and reduce their research efforts in order to manage them (become a "manager"), or whether to remain an "independent contributor".

¹²Of course, this abstracts from any considerations of the type of research produced (e.g., its novelty) and focuses entirely on the total amount of research output.

Formally, consider a researcher, which we call the PI (Principal Investigator), with budget B and cost vector c. He can choose to hire another researcher from a pool $B = \{(c_1, B_1), \dots, (c_n, B_n)\}$ by giving up B_0 units of budget. If they hire agent B_i , they receive utility $v(B - B_0, B_i)$. If they instead decide not to hire, they receive utility v(B).

The first result is immediate.

Proposition 10. The researcher will never hire if all agents in \mathcal{B} have absolute disadvantage in both dimensions with respect to an agent with cost c and budget B_0 . The researcher will always hire if there exists an agent in \mathcal{B} with absolute advantage in both dimensions with respect to an agent with cost c and budget B_0 .

Aside from these extreme cases, the decision of when to hire will be based on comparative advantages of the researchers in \mathcal{B} with respect to the original researcher's budget set with income B_0 . To reduce the dimensionality of the problem, we make the simplifying assumption that the PI hires from a pool of candidates with comparative disadvantage on idea generation. Parametrically, we assume in the following that all researchers have the same budget B_0 and that the absolute disadvantage on ideas of any researcher k is equal to her absolute advantage on methods, that is, we assume for all $k \geq 2$ that

$$\underbrace{c_{\ell}^{k} - c_{\ell}}_{\ell} \qquad = \qquad \underbrace{c_{h} - c_{h}^{k}}_{\ell} \qquad \geq 0 \tag{11}$$

comp. disadvantage on ideas comp. advantage on methods

The next proposition answers the question of whether the PI should hire, and if so who.

Proposition 11. If the PI hires, she hires the researcher with the greatest absolute/comparative advantage on methods. The PI always hires if her optimal solution in isolation (ℓ^*, h^*) is such that

$$h^* \ge \frac{B_0}{c_h + c_\ell}$$

This proposition helpfully characterizes many features of optimal hiring policies. For example, more efficient researchers (ones such that B_0 is smaller) will, all else equal, hire more often. This reflects a sorting effect of comparatively better managers into managerial positions.

Constrained researchers don't hire, that is, when $B < \overline{B}$ from Proposition 3, because such researchers always choose $h^* = 0$, the researcher will not form a team with any of the available agents in \mathcal{B} .

5.4 Human-AI Collaboration

In this part, we consider a different type of collaboration: the interaction of human researchers with technology. Our main focus will be on the impact of AI tools on the human researcher's optimal research mix between ideas and methods. To this end, we model the human-AI interaction as an expansion of the feasible set of (h,ℓ) -combinations. We assume, in line with the findings in Vaccaro et al. (2024), that AI is particularly valuable for creative tasks, implying that access to AI reduces the cost of generating ideas. Specifically, we model the researcher's cost of idea generation to be

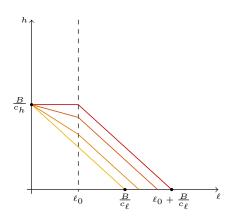
$$c_{\ell}^{H-AI}\ell = \begin{cases} c_{\ell}^{AI}\ell, & \text{if } \ell \leq \ell_0 \\ c_{\ell}^{AI}\ell_0 + c_{\ell}(\ell - \ell_0), & \text{if } \ell > \ell_0, \end{cases}$$
 (12)

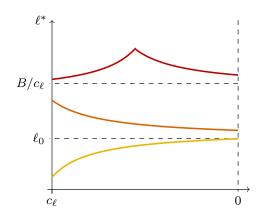
where $c_\ell^{AI} \le c_\ell$ measures the AI performance in idea generation, while ℓ_0 provides an upper bound on the amount of ideas that AI can support generating. Thus, we obtain the budget line of the human-AI team as

$$h^{H-AI} = \frac{B}{c_h} - \begin{cases} \frac{c_\ell^{AI}}{c_h} \ell, & \text{if } \ell \le \ell_0 \\ c_\ell^{AI} \ell_0 - \frac{c_\ell}{c_h} (\ell - \ell_0), & \text{if } \ell > \ell_0. \end{cases}$$
(13)

Figure 4a illustrates the impact of adopting AI as idea-generation tool on the researcher's set of feasible research mixes. As the AI becomes more powerful, the set expands, and eventually generates ℓ_0 ideas for free. Intuitively, the AI tool acts like a team member with a comparative advantage on idea generation relative to the researcher.¹³ The next proposition shows how the adoption of AI technologies as idea-generation tools affects different researchers or fields heterogeneously.

¹³However, it is a team member that cannot contribute any methods to the team, implying that the upper bound on methods remains B/c_h . Moreover, the AI-collaborator's budget is such that it cannot produce more than ℓ_0 ideas.





- (a) Human-AI collaboration budget sets
- (b) Idea generation with AI tool

Figure 4: Human-AI collaboration.

Proposition 12. Adopting increasingly powerful AI tools for idea-generation leads to the following adjustments of the optimal research mix.

- (i) An increase in ideas generated and a decrease potentially followed by an increase in methods mastered for researchers who generate few ideas ($\ell^* < \ell_0$) without the AI tool.
- (ii) A decrease in ideas generated and an increase in methods mastered for researchers who generate many ideas $(\ell^* > \ell_0)$ and master some methods $(h^* > 0)$ without the AI tool.
- (iii) An initial increase and potentially a decrease in ideas generated and an increase in methods mastered for researchers who generate many ideas $(\ell^* > \ell_0)$ but do not master any methods $(h^* = 0)$ without the AI tool.

Proposition 12 highlights the heterogeneous impact of adopting AI tools on different researchers. Figure 4b illustrates this heterogeneity. Researchers whose work relies relatively little on ideas without the AI tool ($\ell^* < \ell_0$) will shift their optimal research mix in favor of ideas over methods. Intuitively, for these researchers, the AI collaboration reduces the marginal cost of generating ideas, corresponding to a rotation of their budget line. Proposition 6 therefore implies that the researcher will generate more ideas and master fewer methods. As the AI becomes more powerful, the researcher eventually hits the AI-boundary ℓ_0 , and

will invest any further technology improvements in mastering additional methods. In contrast, researchers who relied relatively more on ideas while still mastering some methods ($\ell > \ell_0$, $h^* > 0$) will generate fewer ideas while mastering more methods. For these researchers, the AI adoption can be interpreted as a budget increase, which, by Proposition 4 implies that they will generate fewer ideas while expanding their methods. The remaining set of researchers ($\ell^* > \ell_0$, $h^* = 0$) is constrained before the technology adoption and will therefore only generate more ideas initially. Once the researchers become unconstrained, they will act as the ones with ($\ell^* > \ell_0$, $h^* > 0$).

Thus, our results can have fundamentally different implications for the work of different types of researchers. Some may perceive their work as less inspiring and creative, as they rely less on ideas while focusing more on methods, while others may enjoy the additional focus on a broader set of ideas. Moreover, and perhaps counterintuitively, the adoption of idea-generating tools can make the final output *less* idea intensive.

6 Conclusion

In this paper, we provide a novel conceptualization of discovery as a two-dimensional matching problem between ideas and methods embedded in an epistemic land-scape. Modeling the mapping from ideas to required methods as a Brownian motion, we obtain that the researchers' preferences over idea-method combinations are strictly convex. Thus, we can analyze the problem of a resource-constrained researcher using the consumer-theory toolkit and characterize the researcher's Marshallian demands for ideas and methods. With this rich yet tractable and portable model of the knowledge-production process, we can shed light on many understudied aspects of the corporate and academic research. We obtain sharp comparative statics that speak to important policy questions and that can unify a variety of phenomena in science and innovation.

While we touch on several aspects of knowledge production, our framework lends itself to more intensive study of various questions that are beyond the scope of the present paper. For example, consider the case of grant or resource allocation. We take on the perspective of an individual researcher maximizing the probabil-

ity of discovery. However, a funding institution might consider how to allocate a limited budget to a set of heterogeneous researchers. Our team formation results suggest that the weakest researchers should receive additional resources.¹⁴ However, beyond the probability of success the preferences of the funding institution may incorporate also the *type* of successful research, for example, its novelty in the idea or in the method space.

A similarly promising direction is to think about how the head of an R&D division should allocate his heterogeneous researchers (that is, researchers with different relative skills c_{ℓ}/c_h) to problems with different degrees of complexity σ .

Finally, a central opportunity lies in the close empirical analogues of our model ingredients. The *idea breadth* ℓ can be proxied by topical dispersion and the *method breadth* h by methodological dispersion (for a recent approach of measuring a paper's question- and method-novelty, see, for example, Luo et al., 2022). The field-level epistemic volatility σ can be estimated from the observed research mix (ℓ , h) within a field, as our model has immediate and monotone comparative statics with respect to the volatility parameter σ . Hence, our model's predictions are straightforwardly testable and policy recommendations can be tailored to fields and individual researchers' characteristics. The tractability and flexibility of our model allows to easily develop new testable hypotheses in the context of research and innovation.

¹⁴To see this, note that the grants can be thought of as the second group in our team formation problem with identical relative costs. The limited budget can be incorporated by some members of the grant collection having a budget of zero.

A Proofs

A.1 Proof of Lemma 1

Item 1. Consider any rectangle with some $L = [\ell_0, \ell_0 + \ell]$ and $H = [h_1, h_1 + h]$. We need to show that the probability that there exists a pair (x, y(x)) that belongs to $L \times H$ is maximized when $h_1 = -\frac{h}{2}$. We leverage the same characterization of the probability of hitting a given rectangle we develop for Theorem 1. In particular, we know that the above probability for our rectangle will be

$$\begin{split} \int_{h_1}^{h_1+h} d\Phi \left(\frac{z}{\sigma \sqrt{\ell_0}} \right) + 2 \int_{h_1+h}^{+\infty} \left(1 - \Phi \left(\frac{y - (h_1 + h)}{\sigma \sqrt{\ell}} \right) \right) d\Phi \left(\frac{z}{\sigma \sqrt{\ell_0}} \right) \\ + 2 \int_{-h_1}^{+\infty} \left(1 - \Phi \left(\frac{y - (-h_1)}{\sigma \sqrt{\ell}} \right) \right) d\Phi \left(\frac{z}{\sigma \sqrt{\ell_0}} \right) \end{split}$$

Since we want to maximize this, let us set its derivative with respect to h_1 to zero. We get

$$\frac{2}{\sigma^2 \sqrt{\ell \ell_0}} \left(\int_{h_1 + h}^{+\infty} \phi \left(\frac{y - h_1 + h}{\sigma \sqrt{\ell}} \right) \phi \left(\frac{y}{\sigma \sqrt{\ell_0}} \right) dy - \int_{-h_1}^{+\infty} \phi \left(\frac{y + h_1}{\sigma \sqrt{\ell}} \right) \phi \left(\frac{y}{\sigma \sqrt{\ell_0}} \right) dy \right) = 0$$

which is satisfied exactly when $h_1 = -\frac{h}{2}$, and hence $h_1 + h = \frac{h}{2}$.

Item 2. Let $L = [\ell_0, \ell_0 + \ell]$. Given the first item, the sum above can be written as

$$2\int_0^\infty K(y)d\Phi\left(\frac{y}{\sigma\sqrt{\ell_0}}\right)$$

where

$$K(y) = \mathbb{1}[y \le h/2] + \mathbb{1}[y > h/2] 2\left(1 - \Phi\left(\frac{y - (h/2)}{\sigma\sqrt{\ell}}\right)\right).$$

Note that K(y) is decreasing over y and vanishes at infinity. We are interested in

$$\frac{\partial}{\partial \ell_0} 2 \int_0^\infty K(y) d\Phi \left(\frac{y}{\sigma \sqrt{\ell_0}} \right) = 2 \int_0^\infty K(y) \frac{\partial}{\partial \ell_0} \frac{1}{\sigma \sqrt{\ell_0}} \phi \left(\frac{y}{\sigma \sqrt{\ell_0}} \right) dy$$

Since the standard normal distribution satisfies the Heat equation $\frac{\partial f}{\partial \sigma} = \sigma \frac{\partial^2 f}{\partial y^2}$, this

is equivalent to

$$2\int_{0}^{\infty}K(y)\frac{1}{2}\frac{\partial^{2}}{\partial y^{2}}\frac{1}{\sigma\sqrt{\ell_{0}}}\phi\left(\frac{y}{\sigma\sqrt{\ell_{0}}}\right)dy = -\int_{0}^{\infty}\frac{\partial}{\partial y}K(y)\frac{1}{\sigma\sqrt{\ell_{0}}}\frac{\partial}{\partial y}\phi\left(\frac{y}{\sigma\sqrt{\ell_{0}}}\right)dy,$$

where the equality comes from an integration by parts. The latter integral is the integral of the first derivatives of two decreasing functions: K(y) and $\phi(y)$, respectively. Hence, it is positive, and the initial negative sign proves that the derivative is negative. This implies that the constrained optimal value for ℓ_0 is the minimum feasible value, Δ .

A.2 Remaining Steps in the Proof of Theorem 1

Following the steps in the main text, we can rewrite (6) using the cdf $G_{\ell}(\cdot)$ of the supremum of the Brownian motion initialized at zero as

$$2\int_{-\infty}^{0} (1 - G_{\ell}(-h/2 - z)) d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right),$$

as for any $z \le 0$, we have $G_{\ell}(z) = 0$. By symmetry and simplifying the integral, we obtain that this is equivalent to

$$1 - 2 \int_0^\infty G_{\ell}(z - h/2) d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right).$$

Integration by parts allows us to rewrite the last line as

$$2\int_0^\infty \Phi\left(\frac{z}{\sigma\sqrt{\Lambda}}\right) dG_\ell(z-h/2) - 1,$$

as the boundary terms at ∞ , that is, $G(\infty)\Phi(\infty)$, evaluate to one and at zero, that is, $G(-h/2)\Phi(0)$, to zero, because the supremum of the Brownian initialized at zero must be non-negative. A change of variables w = z - h/2 and observing that G has no mass on [-h/2, 0] implies that this expression is equal to

$$2\int_0^\infty \Phi\left(\frac{h/2+w}{\sigma\sqrt{\Delta}}\right)dG_\ell(w)-1,$$

which is expression (7) from the main text.

Finally, we apply a change of variables $p = G_{\ell}(w)$, which requires the substitution $w = G^{-1}(p) =: Q_{\ell}(p)$, where $Q_{\ell}(p)$ is the quantile function of the supremum M_{ℓ} . Hence, we obtain the final expression (8) from the main text

$$2\int_0^1 \Phi\left(\frac{h/2 + Q_{\ell}(p)}{\sigma\sqrt{\Delta}}\right) dp - 1.$$

A.3 Technical Lemma on Standard Normal Hazard Rate

In several proofs we make use of the following lemma that derives properties of the hazard of a standard normal distribution.

Lemma 3. The standard normal hazard rate function H(x) satisfies the following properties for positive arguments x > 0:

- 1. H(x) > x,
- 2. H'(x) = H(x)(H(x) x),
- 3. $0 \le H'(x) < 1$.
- 4. X(H(X) X) is increasing.

A.3.1 Proof of Lemma 3

Item 1. Instead of proving H(x) > x directly, we prove that $\frac{1}{H(x)} < \frac{1}{x}$. This is equivalent to proving that $f(x) = 1 - \Phi(x) - \frac{\phi(x)}{x} < 0$. Note that $\lim_{x \to \infty} f(x) = 0$, and

$$f'(x) = -\phi(x) - \left(-\frac{\phi(x)}{x^2} + \frac{\phi'(x)}{x}\right)$$

$$= -\phi(x) - \left(-\frac{\phi(x)}{x^2} + \frac{-x\phi(x)}{x}\right)$$

$$= -\phi(x) + \frac{\phi(x)}{x^2} + \phi(x) = \frac{\phi(x)}{x^2} > 0.$$

Since the derivative is everywhere increasing and the limit for $x \to \infty$ is 0, it must be that f(x) < 0, which concludes the proof of the first item.

Item 2. This is a simple derivation:

$$H'(x) = \frac{\phi'(x)(1 - \Phi(x) - \phi(x)(-\phi(x))}{(1 - \Phi(x))^2}$$

$$= \frac{\phi(x)^2 - x\phi(x)(1 - \Phi(x))}{(1 - \Phi(x))^2}$$

$$= H(x)\frac{\phi(x) - x(1 - \Phi(x))}{1 - \Phi(x)} = H(x)(H(x) - x).$$

Item 3. Combining Items 1 and 2 from this lemma it is easy to see that H'(x) > 0. To show H'(x) < 1, consider the Mills' Ratio $R(x) = \frac{1}{H(x)}$. The inequality can be reformulated as

$$H(x)(H(x)-x) = \frac{1-xR(x)}{R^2(x)} < 1,$$

or

$$R^2(x) + xR(x) - 1 > 0.$$

We then need to prove that $R(x) > \frac{\sqrt{x^2+4}-x}{2}$. Let the quantity on the right hand side be L(x). Then the inequality can be rewritten once again as

$$G(x) = 1 - \Phi(x) - L(x)\phi(x) > 0.$$

Note first that $\lim_{x\to\infty} G(x) = 0$, and consider the derivative G'(x):

$$G'(x) = -\phi(x) - L'(x)\phi(x) - L(x)\phi'(x)$$

= $\phi(x)(xL(x) - L'(x) - 1)$.

Since $L'(x) = -L(x)\frac{1}{\sqrt{x^2+4}}$, the derivative becomes

$$\phi(x)L(x)\left(\frac{x\sqrt{x^2+4}+1}{\sqrt{x^2+4}}-1\right)$$

We claim that G'(x) < 0. To show this, suppose by contradiction that $G'(x) \ge 0$. Because $\phi(x)$ is always positive, this is equivalent to

$$L(x) \ge \frac{\sqrt{x^2 + 4}}{x\sqrt{x^2 + 4} + 1}.$$

Plugging in L(x) we get

$$\frac{\sqrt{x^2+4}-x}{2} \ge \frac{\sqrt{x^2+4}}{x\sqrt{x^2+4}+1} \iff$$

$$\iff \frac{2}{\sqrt{x^2+4}+x} \ge \frac{\sqrt{x^2+4}}{x\sqrt{x^2+4}+1} \iff$$

$$\iff 2 \ge \frac{x\sqrt{x^2+4}+x^2+4}{x\sqrt{x^2+4}+1} \iff$$

$$\iff 0 \ge \frac{x^2+2-x\sqrt{x^2+4}}{x\sqrt{x^2+4}+1} \iff$$

$$\iff 0 \ge x^2+2-x\sqrt{x^2+4} \iff$$

$$\iff x\sqrt{x^2+4} \ge x^2+2 \iff$$

$$\iff x^4+4x^2 \ge x^4+4x^2+4 \iff$$

$$0 \ge 4$$

which is a contradiction. Thus G'(x) < 0, which implies that G(x) > 0.

Item 4. Item 2 of this lemma shows that X(H(X) - X) is increasing if and only if $-\frac{XH'(X)}{H(X)}$ is decreasing. Recall the Mills' ratio expression as $R(X) = \frac{1}{H(X)}$, and note that

$$-\frac{XH'(X)}{H(X)} = XH(X)\left(-\frac{H'(X)}{H^2(X)}\right)$$
$$= XH(X)\frac{d}{dX}\left(\frac{1}{H(X)}\right)$$
$$= \frac{XR'(X)}{R(X)}.$$

Theorem 2.5, part (b) of Baricz (2008) shows that the last term is decreasing over the positive domain, completing our proof.

A.4 Proof of Proposition 8

Items (i) and (iii) follow directly from observing that any alternative task allocation shrinks the set of feasible (ℓ, h) -combinations. By monotonicity of the success probability, these task allocations are dominated by the one based on compara-

tive advantages. To see this, note that for any feasible ℓ^t , the task allocation that maximizes the feasible h^t follows from

$$\max_{h^1, h^2, \ell^1, \ell^2} h^t$$
s.t. $\ell^1 + \ell^2 = \ell^t$

$$h^1 + h^2 = h^t$$

$$c_\ell^i \ell^i + c_h^i h^i \le B^i.$$

Due to monotonicity, the budget constraints must be binding. Thus, we can (i) substitute h^t by the individual choices, and (ii) substitute individual choices using the budget constraints to obtain

$$\max_{\ell^1,\ell^2} \sum_{i=1}^2 \left(B^i / c_h^i - c_\ell^i / c_h^i \ell^i \right)$$
s.t. $\ell^1 + \ell^2 = \ell^t$.

Finally, we can replace ℓ^1 by the aggregate constraint on ℓ^t , yielding

$$\max_{\ell^2} \ B^1/c_h^1 + B^2/c_h^2 - c_\ell^1/c_h^1\ell^t + \underbrace{(c_\ell^1/c_h^1 - c_\ell^2/c_h^2)}_{<0} \ell^2.$$

Thus, the team chooses the smallest feasible ideas for researcher 2, who has a comparative disadvantage on ideas and the largest feasible amount of ideas for researcher 1.

Item (ii) is a direct consequence of Propositions 3 and 4. If researcher 1's autarky solution ($\ell^{1,*}$, $h^{1,*}$) is interior, then $\ell^{1,*} < B^1/c_\ell^1$ and, holding $\ell^{1,*}$ fixed, the team's budget line at $\ell^{1,*}$ has the same slope c_ℓ^1/c_h^1 but is shifted up from the researcher 1's autarky budget line by B^2/c_h^2 by item (iii). Hence, Proposition 4 implies that the team's optimal choice satisfies $\ell^t < \ell^{1,*}$. Hence, $\ell^2 = 0$, $h^2 = B^2/c_h^2$.

A.5 Lemma 5 and Proof of Proposition 9

First we prove a technical lemma.

Lemma 4. Let $\mathcal{X}(B, B')$ be the joint budget set of the team composed of two researchers with cost vectors c, c' and budgets B and B'. Fix the cost vectors c, c' and consider two budget levels for each researcher, B, B' and \hat{B}, \hat{B}' . For any $x \in \mathcal{X}(B, B')$ and any $y \in \mathcal{X}(\hat{B}, \hat{B}')$ we have

$$\alpha x + (1 - \alpha)y \in \mathcal{X}(\alpha B + (1 - \alpha)\hat{B}, \alpha B' + (1 - \alpha)\hat{B}').$$

Proof. By definition, $x \in \mathcal{X}(B, B')$ if there exist x^1, x^2 such that $x^1 + x^2 = x$ and $c \cdot x^1 \le B$, $c' \cdot x^2 \le B'$, and similarly for $y = y^1 + y^2$. Then clearly $c \cdot \left(\alpha x^1 + (1 - \alpha)y^1\right) \le B$, and $c' \cdot \left(\alpha x^2 + (1 - \alpha)y^2\right) \le B'$. The bundle $\alpha x + (1 - \alpha)y$ is equal to the bundle $\alpha(x^1 + x^2) + (1 - \alpha)(y^1 + y^2)$, which is the sum of the two bundles above, both feasible under their respective budget sets, and hence feasible under the joint budget set.

Lemma 5. The team's indirect utility function $v(B_{\alpha}, \hat{B}_{\beta})$ is non-decreasing in α , β , and submodular.

Proof. For the first part, note that if $\alpha' > \alpha$ then $B_{\alpha'} \supset B_{\alpha}$, and hence the joint budget set $\mathcal{X}(B_{\alpha'}, \hat{B}_{\beta})$ contains $\mathcal{X}(B_{\alpha}, \hat{B}_{\beta})$. The indirect utility function optimizes over a larger set and hence is monotone non-decreasing. The same argument holds for the second collection.

To show submodularity, it is sufficient to prove that the indirect utility function is jointly concave. Denote by **B** the team with budgets $(B_{\alpha}, \hat{B}_{\beta})$ for some α and β in [0,1]. The indirect utility of team **B** is

$$v(\mathbf{B}) = \max_{x \in \mathcal{X}(\mathbf{B})} F(x)$$

where F(x) is the utility function characterized in Sections 3 and 4. Take another team, denoted by \mathbf{B}' , with budgets $(B'_{\alpha}, \hat{B}'_{\beta'})$. Let x^* be the solution to $\max_{x \in \mathbf{B}} F(x)$ and x'^* the solution to $\max_{x \in \mathbf{B}'} F(x)$. Then,

$$v(a\mathbf{B} + (1-a)\mathbf{B}') \ge F(ax^* + (1-a)x'^*)$$

 $\ge aF(x^*) + (1-a)F(x'^*)$
 $= aV(\mathbf{B}) + (1-a)V(\mathbf{B}')$

where the first inequality is the definition of the indirect utility function together

with Lemma 4, the second inequality is the concavity of the utility function shown in Theorem 1 and the last equality is the definition of indirect utility function. The inequality chain shows that the indirect utility function is concave. Since the indirect utility is twice continuously differentiable in the two income levels B_{α} and \hat{B}_{β} , concavity is equivalent to submodularity on the lattice \mathbb{R}^2 .

Now we are ready to conclude the proof of Proposition 9. That the negative assortative matching is welfare-optimal follows the arguments in Becker (1973), and standard optimal transport literature with submodular objectives. Instead it is immediate to see that the stable match must be comonotone from the stability condition itself. Consider B_{α} matched with \hat{B}_{β} (that is, $\tau_{a}(\alpha) = \beta$. It is always true that $v(B_{\alpha}, \hat{B}_{\beta'}) > v(B_{\alpha}, \tau(B_{\alpha}))$ for a $\beta' > \beta$, by virtue of the first claim of Lemma 5. It must be then that every $\beta' > \beta$ prefers their match to B_{α} , hence $\tau_{a}^{-1}(\beta') > \alpha$ for all $\beta' > \beta$. This is exactly the condition for positive assortment.

A.6 Proof of Proposition 12

Item (i) follows from Proposition 6 and observing that the AI-adoption (and further technological improvements) locally corresponds to a reduction in the costs of idea-generation. If the researcher's choice becomes $\ell^* = \ell_0$ eventually, the optimal solution will remain at ℓ^* by Propositions 4 and 6, as the further reduction in idea-generation costs c_ℓ^{AI} pushes the researcher to increase ℓ . However, at ℓ_0 , the slope discontinuously increases from c_ℓ^{AI} to c_ℓ . Without AI, the researcher's optimal choice was below ℓ_0 , and thus, any optimal choice on a higher budget line but with slope c_ℓ must feature $\ell < \ell_0$. Hence, on both segments of the human-AI budget line, we have a corner solution with $\ell^* = \ell_0$. Hence, h^* increase with further improvements in AI capabilities whenever $\ell^* = \ell_0$.

Items (ii) and (iii) are an immediate consequence of Propositions 3 and 4, as the optimal choice features $\ell^* > \ell_0$, and therefore, the relevant segment of the budget line does not change slope with AI adoption and improvements in AI capabilities.

References

- Mohammad Ahmadpoor and Benjamin F Jones. 2019. Decoding team and individual impact in science and invention. *Proceedings of the National Academy of Sciences* 116, 28 (2019), 13885–13890.
- Pierre Azoulay. 2019. Small research teams 'disrupt'science more radically than large ones., 330-332 pages. Issue 566.
- Pierre Azoulay and Danielle Li. 2020. *Scientific grant funding*. Vol. 26889. National Bureau of Economic Research Cambridge, MA.
- Martino Banchio and Suraj Malladi. 2025. Rediscovery. arXiv preprint arXiv:2504.19761 (2025).
- Arjada Bardhi. 2024. Attributes: Selective learning and influence. *Econometrica* 92, 2 (2024), 311–353.
- Arjada Bardhi and Nina Bobkova. 2023. Local evidence and diversity in minipublics. *Journal of Political Economy* 131, 9 (2023), 2451–2508.
- Arjada Bardhi and Steven Callander. 2026. Learning in a Correlated World. *Annual Review of Economics* (2026).
- Árpád Baricz. 2008. Mills' ratio: Monotonicity patterns and functional inequalities. *J. Math. Anal. Appl.* 340, 2 (2008), 1362–1370.
- Gary S Becker. 1973. A theory of marriage: Part I. *Journal of Political economy* 81, 4 (1973), 813–846.
- Catherine Bobtcheff, Jérôme Bolte, and Thomas Mariotti. 2017. Researcher's dilemma. *The Review of Economic Studies* 84, 3 (2017), 969–1014.
- Kevin A Bryan and Jorge Lemus. 2017. The direction of innovation. *Journal of Economic Theory* 172 (2017), 247–272.
- Steven Callander. 2011. Searching and learning by trial and error. *American Economic Review* 101, 6 (2011), 2277–2308.

- Steven Callander and Tom S Clark. 2017. Precedent and doctrine in a complicated world. *American Political Science Review* 111, 1 (2017), 184–203.
- Steven Callander, Nicolas Lambert, and Niko Matouschek. 2025. Innovation and competition on a rugged technological landscape. *working paper* (2025).
- Steven Callander and Niko Matouschek. 2019. The risk of failure: Trial and error learning and long-run performance. *American Economic Journal: Microeconomics* 11, 1 (2019), 44–78.
- Christoph Carnehl, Marco Ottaviani, and Justus Preusser. 2025. Designing Scientific Grants. *Entrepreneurship and Innovation Policy and the Economy* 4, 1 (2025), 139–178.
- Christoph Carnehl and Johannes Schneider. 2025. A quest for knowledge. *Econometrica* 93, 2 (2025), 623–659.
- Richard Michael Cyert and James G March. 1963. A behavioral theory of the firm. Prentice-Hall.
- Giada Di Stefano, Alfonso Gambardella, and Gianmario Verona. 2012. Technology push and demand pull perspectives in innovation studies: Current findings and future research directions. *Research policy* 41, 8 (2012), 1283–1295.
- Umberto Garfagnini and Bruno Strulovici. 2016. Social experimentation with interdependent and expanding technologies. *Review of Economic Studies* 83, 4 (2016), 1579–1613.
- Marc Gruber, Ian C MacMillan, and James D Thompson. 2013. Escaping the prior knowledge corridor: What shapes the number and variety of market opportunities identified before market entry of technology start-ups? *Organization science* 24, 1 (2013), 280–300.
- Ryan Hill and Carolyn Stein. 2025. Race to the bottom: Competition and quality in science. *The Quarterly Journal of Economics* 140, 2 (2025), 1111–1185.
- Ryan Hill, Yian Yin, Carolyn Stein, Xizhao Wang, Dashun Wang, and Benjamin F Jones. 2025. The pivot penalty in research. *Nature* (2025), 1–8.

- Hugo Hopenhayn and Francesco Squintani. 2021. On the direction of innovation. *Journal of Political Economy* 129, 7 (2021), 1991–2022.
- Benjamin F Jones. 2021. The rise of research teams: Benefits and costs in economics. *Journal of Economic Perspectives* 35, 2 (2021), 191–216.
- Boyan Jovanovic and Rafael Rob. 1990. Long waves and short waves: Growth through intensive and extensive search. *Econometrica: Journal of the Econometric Society* (1990), 1391–1409.
- Corinne Kruger and Nigel Cross. 2006. Solution driven versus problem driven design: strategies and outcomes. *Design Studies* 27, 5 (2006), 527–548.
- Zhuoran Luo, Wei Lu, Jiangen He, and Yuqi Wang. 2022. Combination of research questions and methods: A new measurement of scientific novelty. *Journal of Informetrics* 16, 2 (2022), 101282.
- Mary Maher and Hsien-Hui Tang. 2003. Co-evolution as a computational and cognitive model of design. *Research in Engineering design* 14, 1 (2003), 47–64.
- Mary Lou Maher and Josiah Poon. 1996. Modeling design exploration as coevolution. *Computer-Aided Civil and Infrastructure Engineering* 11, 3 (1996), 195–209.
- Suraj Malladi. 2025. Searching in the dark and learning where to look. *working paper* (2025).
- Kyle Myers and Wei Yang Tham. 2023. Money, time, and grant design. *arXiv* preprint arXiv:2312.06479 (2023).
- David Ricardo. 1821. On the principles of political economy. J. Murray London.
- Can Urgun and Leeat Yariv. 2025. Contiguous Search: Exploration and Ambition on Uncharted Terrain. *Journal of Political Economy* 133, 2 (2025), 522–567.
- Michelle Vaccaro, Abdullah Almaatouq, and Thomas Malone. 2024. When combinations of humans and AI are useful: A systematic review and meta-analysis. *Nature Human Behaviour* 8, 12 (2024), 2293–2303.

Eric Von Hippel and Georg Von Krogh. 2016. Crossroads—Identifying viable "need–solution pairs": Problem solving without problem formulation. *Organization science* 27, 1 (2016), 207–221.

Martin L. Weitzman. 1979. Optimal Search for the Best Alternative. *Econometrica* 47, 3 (1979), 641–654.