Pivot or Persist

Martino Banchio* Christoph Carnehl†

Abstract

A resource-constrained researcher seeks to advance knowledge by choosing the ideas to explore and the methods to apply to them. Similar ideas require similar methods, introducing correlation in the researcher's two-dimensional search problem. Our main result proves that the researcher's problem is convex, allowing us to draw parallels to the canonical consumer choice problem. We obtain a rich and tractable framework to study the properties of the optimal research mix that we leverage to generate novel insights in the context of scientific discovery and innovation. We find that ideas are an inferior good, implying that more resources can lead to fewer ideas generated. Applying our results to research teams, we show that task allocation is based on comparative advantages, leading to specialization. Team formation relies on absolute advantages as well, generating maximally or minimally diverse teams. When studying the implications of technology adoption, we find that tools supporting idea generation, such as LLMs, have ambiguous effects on the optimal amount of ideas produced.

1 Introduction

Innovation is often modeled as an iterative process of search, with progress coming from exploring novel ideas or methods. Yet, not all search is equal. 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

*Università Bocconi, IGIER, and Google Research. Email: martino.banchio@unibocconi.it
†Università Bocconi and IGIER. Email: christoph.carnehl@unibocconi.it

source was obvious, but Edison's lightbulb took 6000 trials with different materials and structures.[1] Instead, other innovations begin with a novel method, and the question is how to use it. Theodore Maiman described his invention of the LASER technology as "*a solution looking for a problem.*" Only later did researchers identify applications of the technique to the most disparate fields, such as medicine, telecommunications, and consumer electronics.

These examples illustrate an intrinsic duality of discovery. Progress requires ideas or products, as well as the methods or technologies capable of bringing those ideas and products to life. An idea without its solution method remains unrealized, while a method without an idea remains unused. Discovery is thus not simply about choosing the "best" among a set of options, as in traditional models of search. Instead, discovery relies on a fundamental *matching* across two dimensions, ideas and their solution methods. Progress happens when an idea and a method line up. Researchers thus face a trade-off in deciding how to allocate their research activity. They can devote more effort to explore ideas, or they can master increasingly many methods, knowing that success requires a fit between the two.

In this paper, we develop a model that formalizes this trade-off. We represent the fit between ideas and the methods they require, known as the *epistemic landscape* in the epistemology literature, as a continuous function. We impose a Brownian prior over the epistemic landscape to capture correlation between similar ideas and their solution methods. A researcher, constrained by the resources available, chooses how broadly to explore ideas and how many methods to master, both costly activities. Discovery occurs if at least one explored idea can be matched with one of the methods the researcher mastered. This formulation recasts discovery as a two-dimensional search problem with a matching structure.

Our main results characterize how researchers select their optimal research mix, the balance between ideas and methods, when looking for a single discovery. We first show that preferences over ideas and methods are convex, which allows us to represent the problem in the canonical consumer theory framework. Building on this representation, we derive several results about how the optimal allocation responds to changes in the environment. First, ideas and methods are gross substitutes, thus additional ideas reduce the marginal value of methods and vice versa.

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

Methods are normal while ideas are often inferior: expanding the feasible set of a researcher may lead that researcher to shift their focus to mastery of methods.

By leveraging applications of standard consumer theory, our model can provide sharp answers to several important domains surrounding innovation and research. We show how grants, teaching discounts, and other increases in a researcher's budget may lead to fewer ideas but mastery of more methods.[2] If a researcher can flexibly invest resources in improving her methods or ideas skills (or a graduate school decides how to allocate classes between methods and ideas classes), we find that it is always optimal to invest exclusively in improving the relatively more employed skill.

Further, we study how research collaborations shape discovery, and how such collaborations form.[3] As in trade theory, we find that researchers divide their tasks according to comparative advantages, with the most technically inclined researchers mastering methods alone, and a single researcher exploring ideas. Team formation in contrast is based on absolute advantages, leading to either maximally heterogeneous or maximally homogeneous teams. Finally, we study what technology advancements imply for the researcher. A technology that lowers the cost to master methods (i.e. Mathematica, better testing technology for rapid prototyping) unambiguously leads to the mastery of more methods. Instead, the effect of a technology that lowers the cost of exploring ideas (possibly Large Language Models, as argued by Vaccaro et al. (2024)) is ambiguous. Methods-focused researchers explore more ideas, while more idea-focused researchers may rely increasingly on methods.

Our novel search paradigm delivers tractability and simple testable economic implications to the discovery process. We think such a framework could be applied elsewhere, where the multidimensional-fit aspect of the search process is similarly relevant. Some examples include technology adoption, labor markets, and product-market fit.

[2] Our results speak to the diverse empirical findings on the impact of grants, as in, for example, Myers and Tham (2023).

[3] The increasing role of teams in research and innovation, as well as their diversity is discussed, for example, in Jones (2009) and Jones (2021).

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.[4] 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. 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.[5] Our departure changes the economic problem the researcher faces. 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.

[4]The earlier paper Jovanovic and Rob (1990) axiomatizes a search setting that delivers the driftless Brownian motion as underlying process.

[5]Other papers drawing on related modeling tools, consider different search targets and different processes. 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 assumes an Ornstein-Uhlenbeck process. Callander and Clark (2017) studies search for roots of a Brownian motion. However, none of these papers captures the nption of a two-dimensional search for a match.

Our main applications relate to science and innovation. In the philosophy of science, the seminal paper Kitcher (1990) introduced the concept of the *cognitive division of labor*, which supposes that for science to advance efficiently, scientists should diversify across different approaches rather than focusing only on the most promising approach. Our model studies an individual researchers cognitive division of labor across methods *and* ideas.[6]

In the economics of innovation, several papers address the direction of innovation (see, for example, Hopenhayn and Squintani, 2021; Bryan and Lemus, 2017). Our paper is, to the best of our knowledge, the first to distinguish between two alternative but related directions of innovation.

Papers in economics that focus on academic and corporate researcher's incentives have mainly focused on competition or specific design issues.[7] We take a more fundamental approach to the researcher's problem and provide a framework that can shed light on many related aspects.[8]

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 techniques 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 $y(0) = 0$.

[6] Another strand of literature in philosophy studying the cognitive division of labor is an agent-based analysis of search on epistemic landscapes similar to the one underlying our model (see, for example, Weisberg and Muldoon, 2009; Alexander et al., 2015).

[7] Bobtcheff et al. (2017) and Hill and Stein (2025) study how the race for priority lets researchers compromise on quality of their work. Adda and Ottaviani (2024) analyzes how different aspects of grant design induce self-selection of researchers into grant applications.

[8] Stephan (1996) provides an excellent introduction into the economics of science.

The link between ideas and solution methods is inherently uncertain. We model their relationship as a driftless Brownian motion with volatility parameter $\sigma^2 > 0$, $(y(x))_{x\geq 0}$, 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.

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:

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

Second, the researcher faces a novelty constraint, as the market for discoveries typically does not reward incremental discoveries.[9] 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.

[9]For examples, new patents have to be sufficiently different from existing patents. Academic papers have to be sufficiently distinct from the prior literature, etc.

The problem that the researcher solves is

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

$$\text{s.t.} \quad c_\ell \ell + c_h h \leq B. \tag{3}$$

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

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

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

$$\text{s.t.} \quad c_\ell \ell + c_h h \leq B. \tag{5}$$

3 Convex Tastes for Discovery

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

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

$$(\ell, h) \succ (\ell', h') \iff F(\ell, h) > F(\ell', h'),$$

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, $A(\ell, h + \delta) \supset A(\ell, h)$. 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 $c_\ell, c_h > 0$, 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 (the Marshallian demand) 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 delivers the proof of Theorem 1 and the corresponding intuition.

3.1 Proof of Theorem 1

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 $\succ$ is strictly convex.*

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

Second, we break down the utility function in three components that additively contribute to its value. We can write the probability of the event $A(\ell, h)$ as the integral of the conditional probability of the event happening given the value of the epistemic landscape at Δ, $y(\Delta)$:

$$F(\ell, h) = \int_{\mathbb{R}} \mathbb{P}\Big(A(\ell, h) \mid y(\Delta) = z\Big)\, d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right),$$

where the Brownian assumption yields the distribution of the value of the landscape at Δ, $\Phi(\cdot)$. Recall that the event $A(\ell, h)$ is the event that the epistemic landscape "passes through" the rectangle $L \times H$. We can thus decompose the probability as

$$\begin{aligned} F(\ell, h) = &\int_{-\infty}^{-h/2} \mathbb{P}\Big(A(\ell, h) \mid y(\Delta) = z\Big)\, d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) + \int_{-h/2}^{h/2} 1\; d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) \\ &+ \int_{h/2}^{\infty} \mathbb{P}\Big(A(\ell, h) \mid y(\Delta) = z\Big)\, d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) \\ = &2\int_{-h/2}^{0} d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) + 2\int_{-\infty}^{-h/2} \mathbb{P}\Big(A(\ell, h) \mid y(\Delta) = z\Big)\, d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right), \end{aligned}$$

where the first equality uses the fact that the hitting probability of the rectangle is 1 when the epistemic landscape falls within the rectangle at location Δ, and the second uses the symmetry of the landscape around the x-axis.

Consider now 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 initialized

at 0 over the range $[\Delta, \Delta+\ell]$ by M_ℓ and by $G_\ell(x)$ its distribution function, we have

$$\begin{aligned}
F(\ell,h) =& 2\int_{-h/2}^{0} d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) + 2\int_{-\infty}^{-h/2} (1-G_\ell(-h/2-z))\, d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) \\
=& 2\int_{-\infty}^{0} d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) - 2\int_{-\infty}^{-h/2} G_\ell(-h/2-z)\, d\Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) \\
=& 1 - 2\int_{-\infty}^{-h/2} \Phi\left(\frac{z}{\sigma\sqrt{\Delta}}\right) dG_\ell(-h/2-z) \\
=& 2\int_{0}^{\infty} \Phi\left(\frac{h/2+z}{\sigma\sqrt{\Delta}}\right) dG_\ell(z) - 1 = 2\mathbb{E}_{M_\ell}\left[\Phi\left(\frac{h/2+M_\ell}{\sigma\sqrt{\Delta}}\right)\right] - 1.
\end{aligned}$$

One way to see this integral is that we are averaging over future headroom (the value of M_ℓ) the probability that the value $y(\Delta)$ is within headroom of the barrier $-h/2$. Concavity remains difficult to show directly, since the argument of the integral depends on h and ℓ through a product. However, with canonical arguments, the last integral can be written as

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

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 $p\in[0,1]$.

Since the argument is separable and linear in h, all we are left to check 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 concludes the proof.

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, the Marshallian demand is unique.

We first obtain a convenient representation of the marginal rate of substitution between methods and ideas, which we rely on in the characterization of the researcher's optimal choice and the analysis of its properties.

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

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

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

Equivalently, the MRS can be written as

$$\frac{K}{\ell+\Delta}\mathbb{E}\Big[W - X | W > X\Big]$$

where W is a standard normal distribution. Intuitively, the MRS compares the density of $y(\Delta)$ at the threshold $h/2$ (i.e., the marginal effect of increasing h) to the expected slack *beyond* the threshold (i.e., the marginal effect of increasing ℓ, which affects all values $y(\Delta) > h/2$). That is, increasing the length of the rectangle works through the entire overshoot $y(\Delta) - h/2$ through its expectation.

This characterization leads us to the following observation about the Marshallian demand.

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 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 $e(c_l, c_h, p)$ is increasing in p, thus for all $B \leq e(c_l, c_h, \overline{p})$ the Marshallian demand is $(\ell^*, h^*) = (\ell(p), 0)$. Conversely, for all $B > e(c_h, c_l, \overline{p})$ the Marshallian demand will have $h^* > 0$.

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

By 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 $B > \overline{B}$:

$$\begin{cases} M(\ell^*, h^*) = \frac{c_\ell}{c_h} \\ c_\ell \cdot \ell^* + c_h \cdot h^* = B. \end{cases} \tag{6}$$

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.

4.1 Income and Price Effects

The first part toward understanding the optimal research mix is to understand how the researcher's choice of methods and ideas 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 1. Instead, when* $B > \overline{B}$, *ideas are an* inferior *good: 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 our utility function is twice differentiable, all we need is 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 (6), we get the system of linear equations

$$\underbrace{\begin{pmatrix} M_h & M_l \\ c_h & c_l \end{pmatrix}}_{A} \begin{pmatrix} \partial_B h^* \\ \partial_B l^* \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

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

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

Let us find 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.

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

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

We will make repeated use of the following technical Lemma, which we prove in the Appendix.

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. $0 \leq H'(x) < 1$.
4. $X(H(X) - X)$ *is increasing.*

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

The inferiority of ideas will drive many interesting results in the paper. While we will interpret it further in the next section, note here that this inferiority is related to the dispersion of methods. As 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]

A second natural comparative static to perform is one with respect to costs. Some are immediate from the so-called "law of demand":

Proposition 5. *An increase in the cost c_h of methods 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 ℓ are inferior. □

We can, in fact, prove that ideas and methods are *gross substitutes* by showing a similar result for increases in the cost of ideas c_ℓ.

Proposition 6. *An increase in the cost c_ℓ of ideas leads to*

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

[10] In 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 it is beyond the scope of this paper.

Proof. The proof is again trivial if the Marshallian is in a corner solution. Instead, if it is interior, we can again take the total derivative of the system (6) with respect to c_ℓ and use Cramer's rule to get that $\text{sign}(\partial_{c_\ell}\ell^*) = \text{sign}(-\ell M_h - 1)$. 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 results in this section provide the fundamental properties of the researcher's optimization problem, which we use in the following to generate novel insights in various topics related to innovation and science. To do so, we leverage the tractable reformulation of the researcher's problem as a consumer choice problem together with its comparative statics. Doing so allows us to reinterpret these non-trivial economic questions as relatively straightforward applications of our theoretical results.

5 Applications

In this section, we use our model and main characterization of the optimal research mix to glean some insights into scientific and corporate researchers' incentives. First, as a direct application of the comparative statics derived in Section 4.1, we consider the consequences of investments in either relaxing the time constraint or enhancing the researchers' method-learning or idea-generation skills. Second, we further leverage the simple structure of the optimal research mix to investigate how researchers allocate their limited resources in collaborative settings, where the collaborators may be other human researchers or technological tools, such as AI assistants.

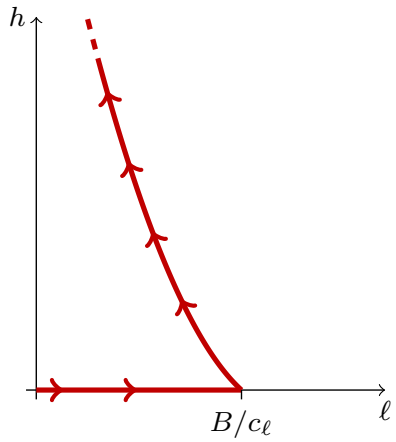


Figure 1: Income-allocation curve

5.1 Investments in Budget and Skills

In various settings, researchers have the opportunity to invest additional resources in spending more on existing projects or in enhancing their own abilities. For decades, engineers at 3M, for example, have been able to spend up to 15% of their time on personal projects.[11]. Of course, some may use this time to advance their projects directly, while others may decide to spend it to learn a new skill or improve existing ones. In academic research, grants allow researchers to buy out teaching hours that they can instead spend on ongoing projects. Alternatively, grant money can be invested in novel technology, making it easier to try out more ideas or methods.

Within our framework, we can interpret these flexible investments as changes in the researcher's budget B or reductions in the skill-dependent costs c_h and c_ℓ, as we characterized in Section 4.1.

Budget changes First, we consider how the researcher's optimal allocation between ideas and methods changes with a marginal increase in her budget.

Proposition 4 shows that methods are a normal good. Any increase in the researcher's budget will unambiguously lead to the researcher attempting a larger set of solutions methods. However, the comparative statics of the ideas inves-

[11] https://www.wired.com/1998/01/the-15-percent-solution/

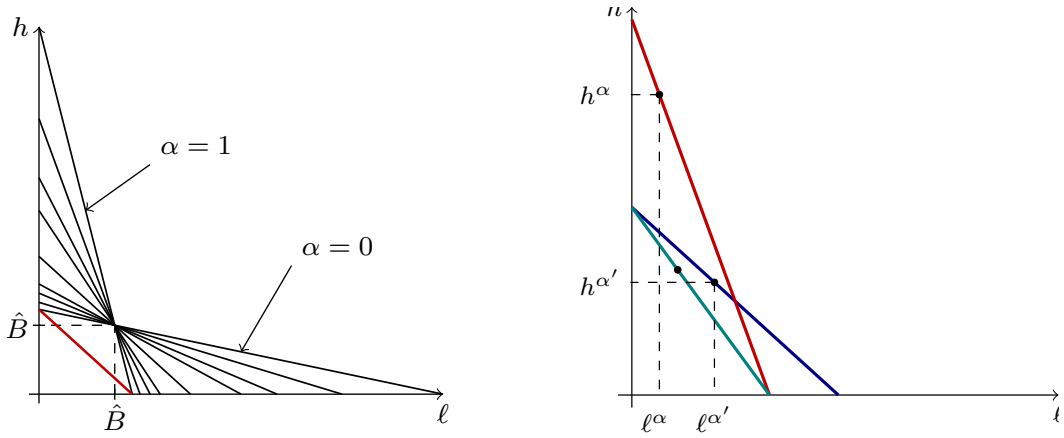
tigated are more subtle. Unless the researcher was depleting all her budget on ideas, a budget increase triggers a reduction in the ideas she generates, in favor of more and more methods. Ideas will only increase if the researcher is choosing the minimal set of methods, $\{h = 0\}$.

Figure 1 depicts the researcher's "Engel curve": the optimal research mix follows the curve as the budget increases. When an agent's budget is small, additional budget will lead to more ideas. The indifference curves are very steep at this budget level, so much so that at current prices adding an idea is cheaper than adding the amount of methods necessary to reach the same probability of success. Instead, a researcher with a large budget will try less ideas if given additional budget. The marginal rate of substitution is decreasing in methods: adding more methods *reduces* the relative value of additional ideas.

This reduction in the marginal rate of substitution is driven by a substitutability of methods and ideas in the researcher's utility function, $F_{hl} < 0$. As the researcher attempts more methods, the value of trying more ideas decreases: the likelihood that the set of methods attempted matches already one of the existing ideas increases, reducing the conditional likelihood that an additional idea leads to discovery.

One interpretation of Proposition 4 is that receiving a grant has ambiguous effects on the "novelty" of a researcher's work. Researchers with relatively constrained budgets, who focus on idea generation, will generate more ideas in response to a budget increase. Researchers with relatively unconstrained budgets, who have a more balanced research mix between ideas and methods, will generate fewer ideas instead. The ambiguous response to a more relaxed budget constraint is in line with the mixed evidence of how grants affect a researchers' question selection, see, for example, Myers and Tham (2023).

Skill changes Consider a researcher characterized by some initial cost/skill vector (c_h, c_ℓ). To understand the consequences of an investment into skill improvements, i.e., reductions in c_h and/or c_ℓ, recall first the comparative statics of the researcher's optimal choice with respect to changes in c_h and c_ℓ of Propositions 5 and 6. Ideas and methods are both net and gross substitutes. Moreover, despite ideas being inferior good for higher budgets, they are not a Giffen good. Hence, the researcher reduces her use of ideas if generating them becomes more costly.



(a) Available budget lines (b) Choice among investments

Figure 2: *Skill investments.*

Suppose now that the researcher can invest an amount $I \in (0, \bar{I})$ in her skills.[12] In particular, she can freely allocate the investment 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, $(h^I = \hat{B}, \ell^I = \hat{B})$, where $\hat{B} := \frac{B}{c_h + c_\ell - I}$.

The next proposition summarizes the optimal investment decision of the researcher.

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 $\alpha = 0$.*

The researcher's choice of α amounts to choosing between budget lines with different slopes, each crossing in the same point $\hat{B}$. If the researcher were to invest equally in both skills, $\alpha = 1/2$, her new budget line would have the identi-

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

cal slope $-c_\ell/c_h$ as her original budget line. Proposition 7 shows that the solution is always bang-bang. If her initial relative skills were sufficiently in favor of methods, she prefers investing in methods. Otherwise, she will specialize further in ideas-generation skills. To prove the first item, notice that the indirect utility function, $v(c_\ell, c_h, B) = F(h^*(c_\ell, c_h, B), \ell^*(c_\ell, c_h, B)$ is quasi-convex in its arguments, a well-known result from consumer theory. For the second result, simply note that Roy's identity relates the indirect utility gain of reducing a cost to the Marshallian demand:

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

Thus, the researcher will invest entirely in the good it currently consumes the most.

To gain some intuition, note that we can derive the researcher's preference relation between any two budget lines identified by investment allocations $\alpha > \alpha'$ following a simple two-step method. Consider first the steeper budget line corresponding to α which has slope $-c_\ell^\alpha/c_h^\alpha$, and identify its optimal allocation (h^α, ℓ^α). Then, consider an increase in the cost of methods to $c_h^{\alpha'} > c_h^\alpha$ but *keeping fixed* the cost of ideas to c_ℓ^α. The new budget line has slope $-c_\ell^\alpha/c_h^{\alpha'}$ it's flatter than the previous one, as one can see in Panel 2 of Figure 2. By Proposition 5, the Marshallian demand with this new budget line will be $h^{\alpha\alpha'} < h^\alpha$ and $\ell^{\alpha\alpha'} > \ell^\alpha$. Since the new budget line lies inside the original feasible allocation set, $v(c_\ell^\alpha, c_h^\alpha, B) \geq v(c_\ell^\alpha, c_h^{\alpha'}, B)$.

Now, consider the budget line with slope $-\frac{c_\ell^{\alpha'}}{c_h^\alpha}$, which is even flatter than what we have considered so far. By Proposition 6 the Marshallian under this budget line will be $h^\alpha < h^{\alpha\alpha'}$ and $\ell^\alpha > \ell^{\alpha\alpha'}$, hence $h^\alpha > h^{\alpha'}$ and $\ell^\alpha > \ell^{\alpha'}$. Again, the feasible set under this budget line contains the previous feasible set, thus $v(c_\ell^{\alpha'}, c_h^{\alpha'}, B) \geq v(c_\ell^\alpha, c_h^{\alpha'}, B)$.

To compare whether the researcher prefers α over α', suppose that the Marshallian $(h^{\alpha'}, \ell^{\alpha'})$ lies to the north-west of $(\hat{B}, \hat{B})$. Then, so does (h^α, ℓ^α), and it follows that the researcher prefers α over α', because the Marshallian under α' is feasible under α. Hence, among all investment levels α that have optimal allocations with $h^\alpha > \hat{B}$ and $\ell^\alpha < \hat{B}$, the researcher prefers the highest investment in methods-learning skills, i.e., $\alpha = 1$. Similarly, for all investment levels that feature $h^\alpha < \hat{B}$ and $\ell^\alpha > \hat{B}$, the researcher prefers the highest investment in idea-generation skills, i.e., $\alpha = 0$. Hence, we can conclude that the solution is bang-bang and all that

remains to compare is whether the researcher prefers $\alpha = 0$ or $\alpha = 1$. 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 higher the initial relative cost, the more valuable are further 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.

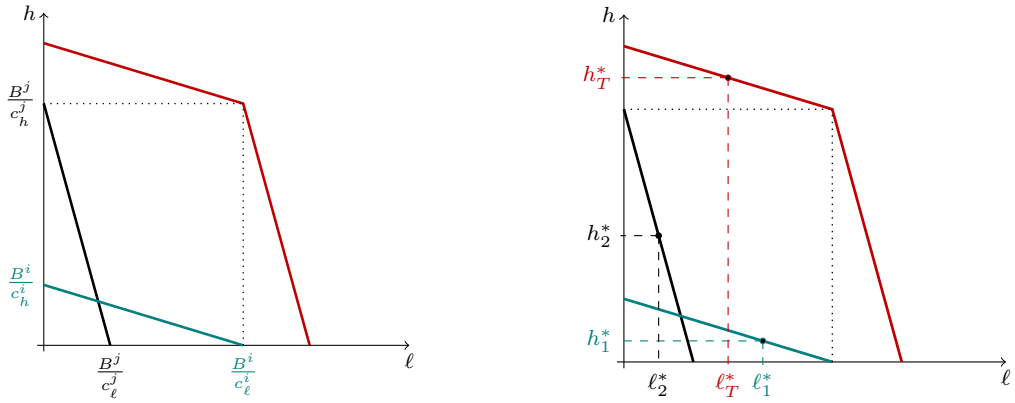
Thus, we may expect that researchers in fields with more complex epistemic landscapes have greater incentives to invest in idea-generation skills than researchers in less complex, more predictable fields.

An alternative interpretation of these results is the optimal design of curricula. For example, 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 complex underlying epistemic landscapes, the optimal curriculum should focus relatively more on asking questions and generating ideas rather than on teaching more methods.

5.2 Task Allocation, Team Formation, and Human-AI Collaboration

In this part, we study how collaboration affects a researchers mix between idea-generation and methods-learning. First, we investigate how tasks are allocated within a team of researchers and how such teams may form. Second, we analyze the consequences of human-AI collaboration on the creativity of a researcher's work.

Task Allocation in Teams Consider a set of n researchers, characterized by their relative skill levels, $\frac{c_\ell^1}{c_h^1} < \frac{c_\ell^2}{c_h^2} < \cdots < \frac{c_\ell^n}{c_h^n}$ and their resources B^i, $i = 1, \ldots, n$. This notation ensures that researchers are sorted by their comparative advantage in idea-generation, ℓ, i.e., a researcher i has a comparative advantage in ℓ with respect to researcher $j > i$. Whenever a two-researcher team forms, we assume that the team jointly optimizes their individual allocations on ideas and methods to maximize their probability of discovery.



(a) Team budget

(b) Optimal allocation

Figure 3: *Team allocation and specialization.*

Proposition 8. *Teams allocate tasks according to comparative advantages and task specialization occurs. Suppose, without loss of generality, that $i < j$. Then, the following hold.*

(i) The team's budget line is

$$h^T = \begin{cases} \frac{B^j}{c_h^j} + \frac{B^i}{c_h^i} - \frac{c_\ell^i}{c_h^i}\ell^T, & \text{if } \ell \leq \frac{B^i}{c_\ell^i} \\ \frac{B^j}{c_h^j} - \frac{c_\ell^j}{c_h^j}\left(\ell^T - \frac{B^i}{c_\ell^i}\right), & \text{if } \ell > \frac{B^i}{c_\ell^i}. \end{cases} \tag{7}$$

(ii) Researcher j does not generate any ideas ($\ell^j = 0$) if researcher i's autarchy solution is interior, that is, if $B^i \geq \overline{B}^i$.

(iii) Researcher i does not learn any methods ($h^i = 0$) if researcher j generates some ideas ($\ell^j > 0$), which can only happen if researcher i's autarchy solution is a corner solution, that is if $B^i < \overline{B}^i$.

Proposition 8 highlights that in our model of research and discovery, central insights from international trade carry over into the setting of research teams. Researcher specialize in tasks and allocate their efforts according to comparative rather than absolute advantages and individual budgets. We use this observation in the following paragraph to study the formation of research teams.

Team formation To understand which research teams form, we consider the following scenario. There is lab manager or a principal investigator, who considers hiring a new collaborator. We assume that this researcher only considers hiring researcher with a comparative advantage on methods, that is, we index the hiring researcher as researcher 1 and denote him as PI in the following. To simplify the analysis, we assume in the following that all researchers have the same budget 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

$$\underbrace{c_\ell^k - c_\ell^1}_{\text{comp. disadvantage on ideas}} = \underbrace{c_h^1 - c_h^k}_{\text{comp. advantage on methods}} \geq 0 \quad \text{and} \quad B^k = B^1 =: B. \tag{8}$$

The next proposition answers the question who the PI should hire to maximize the probability of success, and illustrates the role of absolute advantages in determining team formation.

Proposition 9. *The PI will hire either the researcher with the greatest* absolute *advantage on ideas or the researcher with the greatest* absolute *advantage on methods. Moreover, the following hold.*

(i) If the PI is unconstrained, that is, if $B > \overline{B}^1$, the PI hires the researcher with the greatest absolute *advantage on methods, i.e.,* $\arg\min_{i\in\{2,\dots,n\}} c_h^i$.

(ii) If the PI is constrained and if she could hire a duplicate of herself, and that duplicate would choose $h_{dup} \geq \ell_{dup}$, then the PI would hire the researcher with the greatest absolute *advantage on methods, i.e.,* $\arg\min_{i\in\{2,\dots,n\}} c_h^i$.

(iii) If the PI hires the researcher with the greatest absolute *advantage on ideas, that is,* $\arg\min_{i\in\{2,\dots,n\}} c_\ell^i$, *then it must be the case that a duplicate of PI would choose $\ell_{dup} > h_{dup}$.*

Proposition 9 highlights how absolute advantages are the basis for team formation. Intuitively, absolute advantages play the main role in team formation follows from the specialization result in Proposition 8 conditional on a given team. As specialization will obtain in the new team formed, the PI will choose a collaborator that is maximally skilled in the task that the collaborator will specialize in.

Our results highlight that very often the most successful teams are highly diverse in their background.[13] While it is frequently argued that diverse backgrounds contribute to the success of diverse teams, our analysis adds an additional, simple, perspective: diverse teams often consist of members with different trainings and experiences and thus different skills. As individual team optimally tend to specialize, the best performance in making discoveries is often observed in teams with strong backgrounds in different domains.

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 that, 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

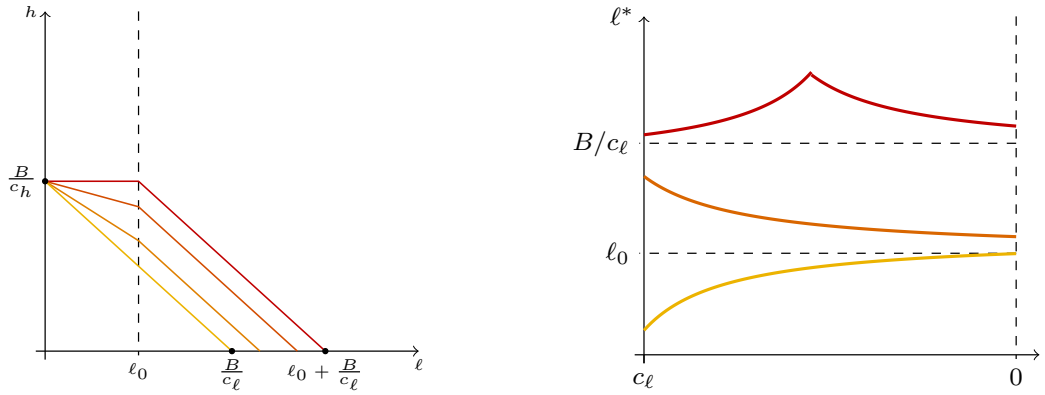
$$c_\ell^{H-AI}\ell = \begin{cases} c_\ell^{AI}\ell, & \text{if } \ell \le \ell_0 \\ c_\ell^{AI}\ell_0 + c_\ell(\ell-\ell_0), & \text{if } \ell > \ell_0, \end{cases} \tag{9}$$

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 set of the human-AI team as

$$h^{H-AI} = \frac{B}{c_h} - \begin{cases} \frac{c_\ell^{AI}}{c_h}\ell, & \text{if } \ell \le \ell_0 \\ c_\ell^{AI}\ell_0 - \frac{c_\ell}{c_h}(\ell-\ell_0), & \text{if } \ell > \ell_0. \end{cases} \tag{10}$$

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

[13] For a discussion of the relevance of teams and their diversity, see, for example, Jones (2021).



(a) Human-AI collaboration budget sets

(b) Human idea generation

Figure 4: *Human-AI collaboration.*

the researcher.[14] The next proposition shows how the adoption of AI technologies as idea-generation tools affects different researchers or fields heterogeneously.

Proposition 10. *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 10 highlights the heterogeneous impact of adopting AI tools on different researchers. Figure 4b illustrates this heterogeneity. Researchers whose work without AI tools relies relatively little on ideas without the AI tool ($\ell^* < \ell_0$),

[14] However, 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.

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 she will generate fewer methods while expanding her methods. The remaining set of researchers ($\ell^* > \ell_0,\ h^* = 0$) is initially constrained and will therefore only generate more ideas initially. Once the researcher becomes unconstrained, she will act as one 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.

A Proofs

A.1 Proof of Lemma 1

tba

A.2 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

$$\begin{aligned} f'(x) &= -\phi(x) - \left(-\frac{\phi(x)}{x^2} + \frac{\phi'(x)}{x}\right) \\ &= -\phi(x) - \left(-\frac{\phi(x)}{x^2} + \frac{-x\phi(x)}{x}\right) \\ &= -\phi(x) + \frac{\phi(x)}{x^2} + \phi(x) = \frac{\phi(x)}{x^2} > 0. \end{aligned}$$

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:

$$\begin{aligned} H'(x) &= \frac{\phi'(x)(1-\Phi(x)-\phi(x)(-\phi(x))}{(1-\Phi(x))^2} \\ &= \frac{\phi(x)^2 - x\phi(x)(1-\Phi(x))}{(1-\Phi(x))^2} \\ &= H(x)\frac{\phi(x)-x(1-\Phi(x))}{1-\Phi(x)} = H(x)(H(x)-x). \end{aligned}$$

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 $G'(x)$:

$$\begin{aligned} G'(x) &= -\phi(x) - L'(x)\phi(x) - L(x)\phi'(x) \\ &= \phi(x)\,(xL(x) - L'(x) - 1)\,. \end{aligned}$$

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 $G'(x) \geq 0$.

Because $\phi(x)$ is always positive, this is equivalent to

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

Plugging in $L(x)$ we get

$$\begin{aligned}
\frac{\sqrt{x^2+4}-x}{2} &\geq \frac{\sqrt{x^2+4}}{x\sqrt{x^2+4}+1} \iff \\
\iff \frac{2}{\sqrt{x^2+4}+x} &\geq \frac{\sqrt{x^2+4}}{x\sqrt{x^2+4}+1} \iff \\
\iff 2 &\geq \frac{x\sqrt{x^2+4}+x^2+4}{x\sqrt{x^2+4}+1} \iff \\
\iff 0 &\geq \frac{x^2+2-x\sqrt{x^2+4}}{x\sqrt{x^2+4}+1} \iff \\
\iff 0 &\geq x^2+2-x\sqrt{x^2+4} \iff \\
\iff x\sqrt{x^2+4} &\geq x^2+2 \iff \\
\iff x^4+4x^2 &\geq x^4+4x^2+4 \iff \\
0 &\geq 4
\end{aligned}$$

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

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

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

A.3 Proof of Proposition 8

Item (i) follows directly from observing that any alternative task allocation shrinks the set of feasible (ℓ, h)-combinations. Items (ii) and (iii) are a direct consequence of Propositions 3 and 4.

A.4 Proof of Proposition 9

Item (i) follows from Propositions 3 and 4. Items (ii) and (iii) follow from Proposition 7.

A.5 Proof of Proposition 10

Item (i) follows from Proposition 6 and observing that the AI-adoption locally corresponds to a reduction in c_ℓ. If the researcher's choice $\ell^* = \ell_0$, the optimal solution will remain at ℓ^* by Proposition 4, and thus h^* will increase.

Items (ii) and (iii) are an immediate consequence of **??** and proposition 4.

References

Jérôme Adda and Marco Ottaviani. 2024. Grantmaking, grading on a curve, and the paradox of relative evaluation in nonmarkets. *The Quarterly Journal of Economics* 139, 2 (2024), 1255–1319.

Jason McKenzie Alexander, Johannes Himmelreich, and Christopher Thompson. 2015. Epistemic landscapes, optimal search, and the division of cognitive labor. *Philosophy of Science* 82, 3 (2015), 424–453.

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.

Árpád Baricz. 2008. Mills' ratio: Monotonicity patterns and functional inequalities. *J. Math. Anal. Appl.* 340, 2 (2008), 1362–1370.

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 and Johannes Schneider. 2025. A quest for knowledge. *Econometrica* 93, 2 (2025), 623–659.

Umberto Garfagnini and Bruno Strulovici. 2016. Social experimentation with interdependent and expanding technologies. *Review of Economic Studies* 83, 4 (2016), 1579–1613.

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.

Hugo Hopenhayn and Francesco Squintani. 2021. On the direction of innovation. *Journal of Political Economy* 129, 7 (2021), 1991–2022.

Benjamin F Jones. 2009. The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies* 76, 1 (2009), 283–317.

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.

Philip Kitcher. 1990. The division of cognitive labor. *The journal of philosophy* 87, 1 (1990), 5–22.

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).

Paula E Stephan. 1996. The economics of science. *Journal of Economic literature* 34, 3 (1996), 1199–1235.

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.

Michael Weisberg and Ryan Muldoon. 2009. Epistemic landscapes and the division of cognitive labor. *Philosophy of science* 76, 2 (2009), 225–252.

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