

A QUEST FOR KNOWLEDGE

CHRISTOPH CARNEHL

Dept. of Economics & IGIER, Bocconi University

JOHANNES SCHNEIDER

Department of Economics, Universidad Carlos III de Madrid

Is more novel research always desirable? We develop a model in which knowledge shapes society's policies and guides the search for discoveries. Researchers select a question and how intensely to study it. The novelty of a question determines both the value and difficulty of discovering its answer. We show that the benefits of discoveries are nonmonotone in novelty. Knowledge expands endogenously step-by-step over time. Through a dynamic externality, moonshots—research on questions more novel than what is myopically optimal—can improve the evolution of knowledge. Moonshots induce research cycles in which subsequent researchers connect the moonshot to previous knowledge.

KEYWORDS: Science of Science, Search on Brownian Path, Knowledge Creation.

1. INTRODUCTION

IN A LETTER TO FRANKLIN D. ROOSEVELT, Vannevar [Bush \(1945\)](#) pleads with the president to preserve freedom of inquiry by federally funding basic research—the “pacemaker of technological progress.” That letter paved the way for the creation of the National Science Foundation (NSF) in 1950. The NSF today, like the vast majority of governments and scientific institutions, cherishes scientific freedom and allows academic researchers to select research projects independently.

With scientific freedom comes the responsibility to select the right research questions. However, what are the right questions? Biologist and Nobel laureate Peter [Medawar \(1967\)](#) famously notes that “research is surely the art of the soluble. ... Good scientists study the most important problems they think they can solve.” Finding the most important yet soluble question is nontrivial. One reason is that both importance and solubility depend on the current state of knowledge.

In this paper, we develop a microfounded model of knowledge creation through research. We conceptualize research as a costly search process that may fail. The researcher picks a question and an intensity with which to search for its answer. The cost of search depends on how well existing knowledge guides the researcher's efforts. If the researcher

Christoph Carnehl: christoph.carnehl@unibocconi.it

Johannes Schneider: jschneid@eco.uc3m.es

We thank Arjada Bardhi, Antonio Cabrales, Steven Callander, Marco Celentani, Rahul Deb, Philipp Denter, Florian Ederer, Chiara Franzoni, Alex Frug, William Fuchs, Ben Golub, Ryan Hill, Toomas Hinosaar, Nenad Kos, Jorge Lemus, Igor Letina, David Lindequist, Gerard Llobet, Ignacio Ortuño, Marco Ottaviani, Nicola Pavoni, Harry Pei, Konrad Stahl, Carolyn Stein, Armin Schmutzler, Carlo Schwarz, Adrien Vigier, Ludo Visschers, Ralph Winkler, and various audiences for comments. Johannes Schneider gratefully acknowledges financial support from Horizon Europe MSCA Project 101061192; the Agencia Estatal de Investigación (MICIU/AEI/10.13039/501100011033) through grants PID2019-111095RB-I00, PID2020-118022GB-I00, and CEX2021-001181-M; the German Research Foundation (DFG) through CRC TR 224 (Project B03); and Comunidad de Madrid, grant MAD-ECON-POL-CM H2019/HUM-5891. Funding for APC: Universidad Carlos III de Madrid (Agreement CRUE-Madroño 2024). Part of the research was carried out while Johannes Schneider was visiting the University of Bern. He thanks them for their hospitality and inspiration.

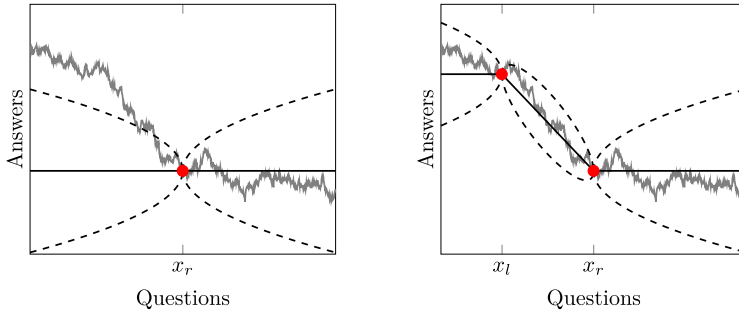


FIGURE 1.—Existing knowledge and conjectures.

discovers the answer, her gross benefits depend on how much society’s decision-making improves due to the discovery—both through answering a particular question and the spillovers on related questions. The discovery is then added to the stock of knowledge and future researchers can build on the additional knowledge. We characterize the researcher’s optimal choice for arbitrary existing knowledge, and find that researchers work on questions that are too narrow and fail too often. The questions chosen are neither novel enough to maximize the instantaneous value of knowledge for decision-making nor novel enough to inspire future generations. By incentivizing distant discoveries, a designer can reduce the researchers’ failure rate and improve the evolution of knowledge.

We model the value of knowledge as the extent to which knowledge improves decision-making. We represent society by a single decision maker who can use existing knowledge to address a variety of problems. Existing knowledge is the set of questions to which the answer has already been discovered. Because answers to similar questions are correlated, knowledge also provides the decision maker with conjectures regarding questions to which the answer is yet undiscovered. The precision of a conjecture depends on the question’s location relative to existing knowledge.¹

We conceptualize the correlation by assuming that answers to questions follow the realization of a Brownian path. Figure 1 depicts that idea. Questions are on the horizontal axis, and the gray line represents the answers to all questions. Red dots (•) represent existing knowledge. Because of the assumption of a Brownian path, all conjectures follow a normal distribution. The mean and the variance depend on existing knowledge. The solid black lines in Figure 1 represent the mean; the dashed lines provide the band of the 95% prediction interval.²

Our first contribution is to characterize how adding a discovery improves the value of knowledge. To see what we have in mind, consider the left panel of Figure 1. Only the answer to question x_r is known. Assume that a researcher discovers the answer to question x_l . Knowledge is now depicted in the right panel. Decision-making improves in three ways. First, the decision maker has precise knowledge about the answer to x_l . Second, her conjectures about all questions to the left of x_l improve. Third, her conjectures improve

¹The protein folding problem in structural biology provides a case in point. Spillovers from other proteins led to Moderna’s development of the COVID-19 vaccine, which “took all of one weekend” (only). For more on the protein folding problem, see Hill and Stein (2024, 2025).

²The 95% prediction intervals describe the following relation: for each question, with a probability of 95%, the answer lies between the respective dashed lines given existing knowledge. For visual clarity, we selected a Brownian path that leaves the 95% prediction interval in the negative branch when we condition on it passing through the origin only.

in the area $[x_l, x_r]$. For all questions in that area, the decision maker now has two pieces of knowledge that help her form the conjecture.

How much decision-making benefits from discovering the answer to x_l depends on x_l 's distance from x_r . The decision maker cares about the range of questions she has a good conjecture about and the actual precision of each conjecture. Increasing the distance between x_l and x_r is similar to the effect of output expansion by a monopolist. Consider two alternatives for a discovery $x'_l < x_l$. Knowing x'_l instead of x_l implies that more questions lie inside the area $[x'_l, x_r]$ than in $[x_l, x_r]$ —a marginal gain in the range of questions with precise conjectures. At the same time, the conjectures about each question in $[x_l, x_r]$ become less precise with x'_l —an inframarginal loss. The benefits of a discovery are maximized at an intermediate distance.

If both x_l and x_r are known initially, discoveries *expand knowledge* beyond the frontier if the discoveries concern questions $x \notin [x_l, x_r]$. Expanding beyond the frontier works in the manner described in the previous paragraph. Alternatively, discoveries *deepen knowledge* if they concern questions $x \in [x_l, x_r]$. Depending on the distance between x_l and x_r , expanding knowledge or deepening knowledge may be optimal. If x_l and x_r are close, the conjecture about any question in $[x_l, x_r]$ is already precise. In this case, expanding knowledge provides larger benefits than deepening knowledge does. If x_l and x_r are far apart, conjectures about questions in $[x_l, x_r]$ are imprecise. Obtaining an answer to a question $x \in [x_l, x_r]$ divides this single area of imprecise conjectures into two areas with precise conjectures. In this case, deepening knowledge provides larger benefits than expanding knowledge.

Overall, the decision maker benefits most from deepening knowledge between distant, yet not too distant, pieces of knowledge. Expanding knowledge beyond the frontier beats deepening knowledge only if all available areas are short.

Our second contribution is to characterize a researcher's optimal choice of which question to tackle and how much effort to invest in studying that question for arbitrary existing knowledge. We assume that the researcher's benefits of a finding are proportional to the benefits for the decision maker we have discussed above. In addition, we conceptualize the research process as the search for an answer. We assume that it requires effort to search for an answer and that the cost of effort is increasing and convex. We propose a microfounded cost function that captures the following idea: For a given amount of effort, search is more likely to result in a discovery if the conjecture about the answer is more precise. Because conjectures depend on the distance to existing knowledge, our cost function naturally links *novelty* (the distance to existing knowledge) and research *output* (the probability that search results in a discovery).

Regarding the researcher's choices, we show that novelty and output are nontrivially related. Depending on the structure of existing knowledge, the two can substitute or complement each other. Thus, in some cases selecting a more novel question actually decreases the risk of failure. Overall, the researcher creates more output when she deepens knowledge than when she expands beyond the frontier. Output peaks for deepening knowledge in areas of intermediate length. This suggests that researchers prefer areas that contain a limited number of questions, even though they benefit from improving decision-making across a wide range of problems. They mainly do this for two reasons: First, discoveries in smaller areas lead to more significant improvements on each question in that area. Second, it is less costly to obtain a discovery.

Our third contribution is to apply our previous insights to study the evolution of knowledge. We turn to a setting in which research is done sequentially and show that if a researcher expands knowledge, no future researcher will deepen knowledge. Therefore, the evolution of knowledge takes on a ladder structure.

Starting from this observation, we study a simple policy intervention. Suppose a long-lived designer can direct the choices of one researcher. The designer aims to improve decision-making by enhancing the evolution of knowledge. Under natural conditions, the designer induces a research cycle by encouraging a moonshot discovery—a discovery far from existing knowledge. Moonshot discoveries are suboptimal in the short run. They create knowledge that is too disconnected from existing knowledge and, therefore, provides little immediate benefit. However, moonshots guide future researchers aiming at questions between the moonshot and previously existing knowledge. As a result of the moonshot, future researchers increase their output and pursue different questions. The knowledge they create becomes more valuable than otherwise. If the designer is patient and the cost of research is intermediate, the positive dynamic externality of moonshots dominates the implied myopic loss. Thus, moonshots have the flavor of an infrastructure investment. By themselves, they provide little value, yet they enable future researchers to thrive by setting the stage.

To summarize, we make three contributions. First, we offer a microfounded framework to study knowledge and research in a complex world. We quantify the value of a discovery when society extrapolates from knowledge to address a wide range of questions. Second, we shed light on the nontrivial relation between the novelty of a question and the probability that a researcher discovers its answer. Novelty and the probability of discovery are endogenously linked through a microfounded cost function. Third, we provide novel insights into a classical question in the economics of science funding: should a funder incentivize research far beyond the frontier even if the immediate benefits of such a discovery are low? Yes, if the cost of research is intermediate and society is patient. The research cycle that such a moonshot induces leads to researchers addressing more novel questions and producing more output.

Related Literature. The premise of our model is rooted in a large literature originating in the philosophy of science that studies how scientific research evolves; [Kuhn \(1962\)](#) is the most prominent contribution therein. [Kuhn \(1962\)](#) distinguishes between two phases in the evolution of knowledge: normal science and scientific revolutions. While his focus and, therefore, his hypotheses differ from ours, our findings share some similarities with his description of the dynamics of science. Our main distinction to the literature building on [Kuhn \(1962\)](#) is that we study how past discoveries shape today's choice of research questions and the search for answers rather than how an overreaching theory is used to tackle a given question. Building on [Kuhn \(1962\)](#), [Kitcher \(1990\)](#) introduces a model of researchers choosing between different theories. Further developing that idea, [Brock and Durlauf \(1999\)](#) show that social factors, such as the desire for conformity, affect whether a superior theory will outcompete an inferior theory over time. [Strevens \(2003\)](#) argues that the priority reward system, as described in [Merton \(1957\)](#), can achieve an efficient allocation of resources across approaches. [Akerlof and Michailat \(2018\)](#) provide a dynamic economic model of competing paradigms and show that a false paradigm may survive due to homophily across generations of researchers. Relatedly, philosophers of science have studied how researchers uncover the most significant approach to a research topic ([Alexander, Himmelreich, and Thompson \(2015\)](#), [Weisberg and Muldoon \(2009\)](#)). They formalize approaches as locations in an epistemic landscape and their significance as the unknown height at that location. None of these ideas is our focus, yet our model connects naturally and provides a complementary perspective on cumulative research.

Closer to an economic approach is a growing literature that leverages large-scale data sets to shed light on researchers' decision-making (see [Fortunato et al. \(2018\)](#), for an

overview). Foster, Rzhetsky, and Evans (2015) empirically analyze researchers' strategies in biomedical chemistry and find that conservative research strategies, akin to little novelty and deepening knowledge in our setting, are more widespread than risky strategies, akin to high novelty and expanding knowledge in our setting. They emphasize a tension between being productive and making novel contributions. Building on these ideas, Rzhetsky, Foster, Foster, and Evans (2015) simulate how risk-taking speeds up the evolution of knowledge. Myers (2020) discusses the effectiveness of policy instruments to encourage more novelty in research. We view our work as complementary to this literature. We provide a theoretical model that, in its baseline static version, generates outcomes in line with their empirical findings. The tractability of our model allows us to generate further insights from making it dynamic and by adding policy tools to the picture.

An adjacent literature on innovation in corporate R&D studies the allocation of resources to different innovative strategies. Our key departure from this literature lies in the value of the findings. In corporate R&D, discovery is rewarded by a product's market value. In science, the value of a finding is the knowledge it generates, for example, to guide decision-making. In the innovation literature, the direct spillovers of findings on the value of related questions are rarely the main driver of incentives. However, in our model the benefits of a finding specifically derive from improved knowledge about related questions. Hopenhayn and Squintani (2021), for example, study the competitive dynamic allocation of researchers to different questions, but assume the value of projects is independent of the stock of completed projects. Bryan and Lemus (2017) analyze a similar problem. In their setting, innovations may relate, but, in contrast to our setting, the sequence in which innovations can be achieved is exogenously fixed.

A key aspect of our model is that the benefits and costs of addressing a research question depend on existing knowledge. The theoretical literature on scientific discoveries does not explicitly model this aspect, yet it incidentally captures parts of the scientific process we have in mind. Aghion, Dewatripont, and Stein (2008) consider a setting in which they assume that knowledge evolves in an exogenous step-by-step structure, whereas Bramoullé and Saint-Paul (2010) provide a model of research cycles albeit without considering an intensive margin. In our framework, intensive margins are relevant.

Conceptually, we contribute to the literature modeling search as discovery on a Brownian path that builds on Callander (2011). Our focus on modeling scientific research leads us to depart from the canonical ideas in the existing literature—most notably in two dimensions.

First, our decision maker aims to understand the entire Brownian path, as any question can become a potential problem to act on. In contrast, in Callander (2011), Garfagnini and Strulovici (2016) knowing the location of the optimal realization along the path suffices for their decision maker. Closer to us are Bardhi (2024) and Callander and Clark (2017). Still, in Bardhi (2024), being informed about a summary statistic of the Brownian path suffices to make an optimal decision. In Callander and Clark (2017), being informed about the roots of the Brownian path suffices to make an optimal decision. In all of their models, the discovery of realizations beyond the frontier declines in value over time. Callander, Lambert, and Matouschek (2022) is an exception where market competition slows this decline. Nevertheless, knowledge expansion eventually halts. In our model, there is a constant and endogenous desire to expand knowledge.³

³Jovanovic and Rob (1990) study a related problem. There, expanding knowledge implies an i.i.d. draw at a fixed cost, while deepening knowledge is costless. Here, all questions are connected. See Callander and Hummel (2014), Callander, Lambert, and Matouschek (2021), Bardhi and Bobkova (2023), Callander and Matouschek (2019), Urgan and Yariv (2025) for unrelated applications in a related framework.

Second, we conceptualize discovering the realization of the Brownian motion at a particular point as a costly search process that may fail. This generates an endogenous link between novelty and output, leading to a trade-off: should the researcher choose more novel questions or higher research output? The existing literature ignores this link between novelty and output.

Combining both differences to the literature, a new reason why learning stops emerges: the stochastic process taking an unexpected turn. Such turns prevent researchers from discovering the resulting unexpected answers, as their search optimally focuses on the most likely answers. This reason differs from those present in the literature where search tends to stop because further search becomes less valuable; for example, due to very informative or valuable discoveries.

Finally, the endogenous growth literature is related in that research and innovation generate value for society. Typically, the value of successful research derives from improvements in the product market, which is usually modeled in one of two ways: by expanding the set of available varieties (e.g., [Romer \(1990\)](#)) or by climbing a quality ladder, that is, by replacing old products or processes with improved ones (e.g., [Grossman and Helpman \(1991\)](#), [Aghion and Howitt \(1992\)](#)). In our model, a ladder structure of knowledge expansion somewhat reminiscent of the quality ladder model as described in [Klette and Kortum \(2004\)](#) arises as well. However, there is a crucial difference. Knowledge advances horizontally in step sizes, but new steps leave the value of old ones unaffected. In that sense, the ladder structure arising in our model is closer to the expansion of product varieties as in [Romer \(1990\)](#). Yet, we show that the ladder structure may be suboptimal.⁴

2. A MODEL OF KNOWLEDGE AND RESEARCH

There are two players: society—represented by a single decision maker—and a researcher. The researcher observes initial knowledge \mathcal{F}_k and chooses two actions: a question, x , and a probability, ρ , with which she discovers the answer, $y(x)$, to question x . With probability ρ , knowledge is augmented by the question–answer pair and becomes $\mathcal{F}_k \cup \{(x, y(x))\}$. With complementary probability, knowledge remains \mathcal{F}_k . Finally, the decision maker observes current knowledge and decides for every question whether to apply knowledge or to select an outside option.

2.1. Knowledge and Conjectures

Questions and Answers. We represent the universe of questions by the real line. A question is an element $x \in \mathbb{R}$. Each question x has a unique answer, $y(x) \in \mathbb{R}$.

Truth and Knowledge. The answer $y(x)$ to question x is determined by the truth. The truth is the graph of the realization of a random variable $Y(x)$ that follows a standard Brownian motion defined over the entire real line.⁵ This assumption captures the following notion: the answer to question x is likely to be similar to the answer to a close-by question x' . As the distance between x and x' increases, the uncertainty increases. However, a correlation remains.

⁴It is important to keep in mind that questions in our model are related and have informational spillovers. This relation is absent in models with product varieties.

⁵As in [Callander \(2011\)](#), the realized truth Y is a random draw from the space of all possible paths \mathcal{Y} generated by a standard Brownian motion going through an initial knowledge point $(x_0, y(x_0))$.

Knowledge is the finite collection of known question–answer pairs. We denote it by $\mathcal{F}_k = \{(x_i, y(x_i))\}_{i=1}^k$. For notational convenience, we assume that \mathcal{F}_k is ordered such that $x_i < x_{i+1}$. We refer to x_1 and x_k as the *frontiers* of knowledge. Knowledge \mathcal{F}_k determines a partition of the real line consisting of $k + 1$ elements

$$\mathcal{X}_k := \{(-\infty, x_1), [x_1, x_2), \dots, [x_{k-1}, x_k), [x_k, \infty)\}.$$

Each element of the partition \mathcal{X}_k is an *area*. We call $(-\infty, x_1)$ area 0, $[x_1, x_2)$ area 1, and so on until area k , which is $[x_k, \infty)$. The length of area $i \in \{1, \dots, k-1\}$ is $X_i := x_{i+1} - x_i$, and $X_0 = X_k = \infty$.

Conjectures. A conjecture is the cumulative distribution function $G_x(y|\mathcal{F}_k)$ of the answer $y(x)$ to question x given knowledge \mathcal{F}_k . Because $Y(x)$ is determined by a Brownian motion, the conjecture about $y(x)$ is a cumulative distribution function of a normal distribution with mean $\mu_x(Y|\mathcal{F}_k)$ and variance $\sigma_x^2(Y|\mathcal{F}_k)$. Both μ_x and σ_x^2 follow from the properties of the Brownian motion.

PROPERTY 1—Expected Value: Given \mathcal{F}_k , the conjecture $G_x(y|\mathcal{F}_k)$ has mean:

$$\mu_x(Y|\mathcal{F}_k) = \begin{cases} y(x_1) & \text{if } x < x_1, \\ y(x_i) + \frac{x - x_i}{X_i}(y(x_{i+1}) - y(x_i)) & \text{if } x \in [x_i, x_{i+1}), i \in \{1, \dots, k-1\}, \\ y(x_k) & \text{if } x \geq x_k. \end{cases}$$

PROPERTY 2—Variance: Given \mathcal{F}_k , the conjecture $G_x(y|\mathcal{F}_k)$ has variance:

$$\sigma_x^2(Y|\mathcal{F}_k) = \begin{cases} x_1 - x & \text{if } x < x_1, \\ \frac{(x_{i+1} - x)(x - x_i)}{X_i} & \text{if } x \in [x_i, x_{i+1}), i \in \{1, \dots, k-1\}, \\ x - x_k & \text{if } x \geq x_k. \end{cases}$$

2.2. Actions and Payoffs

Decision Maker. For each question x , the decision maker either applies knowledge, or takes an outside option. The decision-maker's payoff from taking the outside option is normalized to zero. If she applies knowledge, her expected payoff is

$$1 - \frac{\sigma_x^2(Y|\mathcal{F}_k)}{q},$$

with $q > 0$ exogenously given.⁶ Abstracting from any prioritization of questions, the decision maker values all questions equally. Her value of knowing \mathcal{F}_k is

$$v(\mathcal{F}_k) := \int_{-\infty}^{\infty} \max\left\{1 - \frac{\sigma_x^2(Y|\mathcal{F}_k)}{q}, 0\right\} dx.$$

⁶One microfoundation to attain these payoffs is to assume that the decision-maker's preferences are represented by an affine transformation of a quadratic loss function and she aims to match the truth. She would then choose $\mu(x)$ as her best reply. Normalizing the outside option to zero, choosing an intercept of one and a slope of $1/q$ provides our payoff formulation.

The decision-maker's outside option. A key element of our model is that the precision of knowledge determines whether the decision maker applies it to address problems or prefers the outside option, \emptyset . This feature captures what is referred to as the “precautionary principle”: if uncertainty is large, prudence trumps risking poor application of knowledge. How often the decision maker reverts to the outside option is governed by the parameter q , the size of which is not essential to our results as long as $q \in (0, \infty)$. Yet, q gives knowledge creation a meaning. To see this, consider the two limiting cases. As $q \rightarrow 0$, the decision maker does not apply knowledge unless she is certain about the answer. As knowledge is finite, but problems are uncountably infinite, the decision maker chooses the outside option almost everywhere for any knowledge: knowledge becomes irrelevant to decision-making. At the other extreme, $q \rightarrow \infty$, the decision maker prefers to apply knowledge no matter how vague the conjecture. In that case, the emphasis the decision maker puts on the precision of her answers must be low: knowledge becomes irrelevant to decision-making.

Researcher. The researcher builds on initial knowledge \mathcal{F}_k and decides to search for an answer to a question x . Another key element of our model is that the researcher decides how much effort to exert to find the answer. Given a choice of question x , we posit a one-to-one relationship between the level of effort exerted and the resulting probability ρ of discovering the answer $y(x)$. Thus, (x, ρ) is a sufficient statistic to summarize the researcher's choice.

To save on notation, we allow the researcher to choose ρ directly. We aim to capture that increasing the probability of discovery requires costly effort. In Section 4, we provide details and a microfoundation. For now, we assume the researcher's cost to be

$$\hat{c}(\rho; x) := \eta(\text{erf}^{-1}(\rho))^2 \sigma_x^2(Y|\mathcal{F}_k),$$

where $\eta \geq 0$ is an exogenous cost parameter and $\text{erf}^{-1}(\cdot)$ is the inverse error function of the normal distribution. The cost scales in uncertainty: discovering $y(x)$ with a given probability ρ is more costly if the conjecture about $y(x)$ is less precise.

On the benefits side, we assume that the researcher is aligned with the decision maker. The researcher's total payoff is

$$\rho(v(\mathcal{F}_k \cup \{(x, y(x))\}) - v(\mathcal{F}_k)) - \eta(\text{erf}^{-1}(\rho))^2 \sigma_x^2(Y|\mathcal{F}_k).$$

Economically, our assumption entails that the researcher benefits from her discovery through its impact on the decision-maker's payoff. We revisit this assumption and its implications when we get to the researcher's problem (Section 4) and again in a dynamic version of the model in Section 5.

3. THE BENEFITS OF DISCOVERY

In this section, we describe the benefits of a particular discovery. That is, we ignore how the discovery came about and focus on the marginal value it creates.

Discovery occurs whenever an answer is found and the new question–answer pair is added to existing knowledge, \mathcal{F}_k . The (gross) benefits of discovering $y(x)$ are the marginal increase in the value of knowledge from adding $(x, y(x))$ defined as

$$V(x; \mathcal{F}_k) := v(\mathcal{F}_k \cup \{(x, y(x))\}) - v(\mathcal{F}_k).$$

We distinguish two scenarios: expanding knowledge and deepening knowledge. A discovery $y(x)$ *expands* knowledge if $x \notin [x_1, x_k]$. A discovery $y(x)$ *deepens* knowledge in area i if $x \in [x_i, x_{i+1}]$.

The benefits of a discovery are determined by the length of the research area, X , the discovery occurs in, and how distant the question, x , is from existing knowledge.

DEFINITION 1—Distance: The distance of question x from knowledge \mathcal{F}_k is the minimal Euclidean distance to a question to which the answer is known:

$$d(x) := \min_{\xi \in \{x_1, x_2, \dots, x_k\}} |x - \xi|.$$

DEFINITION 2—Variance: The variance of a question with distance d in an area of length X is $\sigma^2(d; X) := d(X - d)/X$.

Note that $\sigma^2(d; X) = \sigma_x^2(Y|\mathcal{F}_k)$ for $d(x) = d$ and x in an area of length X . Abusing notation, we define for any $X \in \mathbb{R} \cup \{\infty\}$ and $d \in [0, X/2]$,

$$\begin{aligned} V(d; X) := & \frac{1}{6q} (2X\sigma^2(d; X) + \mathbf{1}_{d > 4q} \sqrt{d}(d - 4q)^{3/2} \\ & + \mathbf{1}_{X - d > 4q} \sqrt{X - d}(X - d - 4q)^{3/2} \\ & - \mathbf{1}_{X > 4q} \sqrt{X}(X - 4q)^{3/2}), \end{aligned}$$

where $V(d; \infty) := \lim_{X \rightarrow \infty} V(d; X)$.

PROPOSITION 1: $V(d; X)$ describes the benefits of a discovery $(x, y(x))$ with distance $d(x) = d$ to existing knowledge when the question x lies in an area of length X .

Proposition 1 shows that the benefits of a discovery depend only on the parameters d and X of a question, not its precise location. To gain some intuition on $V(d; X)$, note that the terms without an indicator function measure the direct reduction in uncertainty about answers due to the discovery. The indicator that enters negatively becomes active if the decision maker chose the outside option for some question in the area *before* discovery. The indicators that enter positively become active only if the decision maker chooses the outside option for some question *after* discovery.

Expanding Knowledge. Suppose we expand knowledge by discovering a question-answer pair $(x, y(x))$ with $x < x_1$. Our discovery pushes the knowledge frontier to the left generating the area $[x, x_1]$. The benefits of this knowledge expansion come from the value of the new area $[x, x_1]$ —the dark-shaded area in Figure 2's left panel.⁷

The value of the new area depends on the amount of questions with conjectures based on two discoveries and the degree of improvement in those conjectures relative to the outside option. The benefits-maximizing question resolves a marginal-inframarginal trade-off. Increasing the length of the newly created area has two opposing effects on the value: (i) a positive marginal effect because the amount of questions in the new area increases, and (ii) a negative inframarginal because the degree of improvement decreases.

⁷More precisely, the conjectures about questions to the left of the old frontier are replaced by conjectures inside the new research area, and conjectures to the left of the new frontier also become more precise. As can be seen in the left panel of Figure 3, the variance reduction to the left of the frontier is always the same. Hence, the benefits are the same as if only the new area was added.

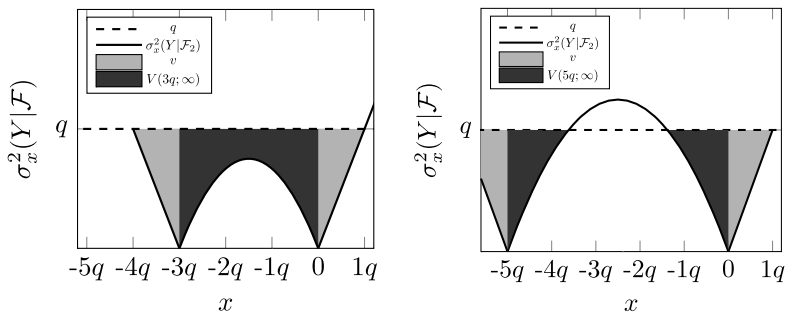


FIGURE 2.—Benefits-maximizing (left) and too large (right) distance of x given \mathcal{F}_1 .

Figure 2 and Figure 3 illustrate the benefits of discovery from creating ideal (left panel of Figure 2), too large (right panel of Figure 2), and too short (left panel of Figure 3) areas. The largest benefits come at an intermediate level at which all conjectures have a variance strictly smaller than q .

Deepening Knowledge. Deepening knowledge differs conceptually. Instead of creating a new area, discoveries replace an existing area with two shorter ones. Discovering the answer to $y(x)$ replaces an area, $[x_i, x_{i+1})$ with the new areas $[x_i, x)$ and $[x, x_{i+1})$.

Areas of length $3q$ provide the largest benefits, the one depicted in the left panel of Figure 2. Thus, the benefits-maximizing discovery inside an area of length $X_i = 6q$ is naturally at the midpoint, replacing the area by two areas of length $3q$.

Finding the benefits-maximizing discovery for areas of length $X_i \neq 6q$ is less straightforward. There are two forces at play. First, there is a benefit to replacing the old area with two symmetric new areas. The intuition echoes that of expanding knowledge: the inframarginal loss increases when an area becomes too large. Choosing symmetric area lengths reduces the inframarginal losses compared with asymmetric area lengths. Second, benefits decline if the area length is greater than $3q$ because conjectures inside the area become imprecise.

If the initial area length X_i was small, the first force would dominate. Selecting the midpoint at $d = X_i/2$ is optimal. However, if X_i was large, the trade-off would be resolved in favor of creating one high-value area at the cost of having imprecise conjectures in the other. A cutoff, \tilde{X}^0 , determines which force dominates.

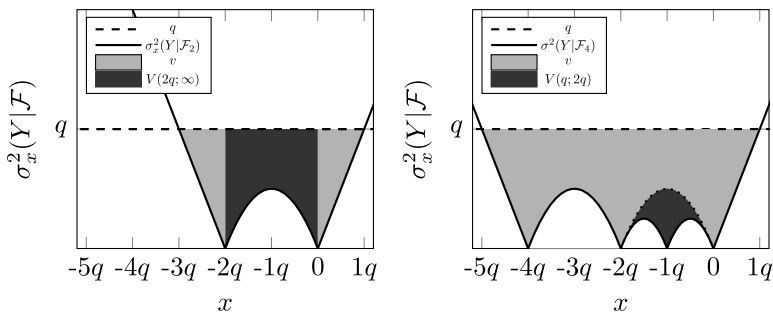


FIGURE 3.—The benefits of discovery. The dark-shaded area illustrates the benefits of discovery; the light-shaded areas illustrate the value of initial knowledge. In the *left panel*, knowledge is expanded. In the *right panel*, knowledge is deepened.

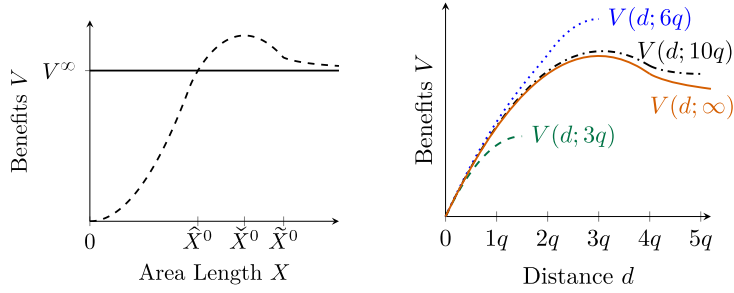


FIGURE 4.—The benefits of discovery. The dashed line in the *left panel* plots the benefits of discovery $V(d^0(X); X)$ against $X < \infty$. The solid line shows the maximum benefits from expanding knowledge. The *right panel* plots $V(d; X)$ for different X .

Expanding vs. Deepening Knowledge. Finally, there is a trade-off between expanding knowledge and deepening knowledge. On the one hand, expanding knowledge means that no area is replaced and a new area is created. On the other hand, deepening knowledge means creating two areas with precise conjectures. Expanding knowledge provides higher benefits than deepening knowledge in some existing area only if all existing areas are shorter than a cutoff \hat{X}^0 . Figure 4 illustrates this observation.

We now summarize our findings. We define the optimal distance as

$$d^0(X) := \arg \max_d V(d; X).$$

PROPOSITION 2: *A distance of $d^0(\infty) = 3q$ maximizes the benefits of expanding knowledge. When deepening knowledge in an area with length below a cutoff \hat{X}^0 , the midpoint of the area maximizes the benefits, $d^0(X) = X/2$. Otherwise, a distance between $3q$ and the midpoint maximizes the benefits of deepening knowledge, $d^0(X) \in (3q, X/2)$.*

Expanding knowledge is only benefits-maximizing if all available bounded areas are shorter than a cutoff, \hat{X}^0 , with $\hat{X}^0 < \tilde{X}^0$. The benefits from optimally deepening knowledge are single peaked in the area length with the peak at $\tilde{X}^0 \in (\hat{X}^0, \tilde{X}^0)$.

4. THE RESEARCHER

4.1. The Researcher's Objective

In this section, we analyze the researcher's optimal choice. Recall that the researcher only benefits from research if it culminates in a discovery.⁸ The researcher's expected payoff given a choice of question x and success probability ρ can be written as

$$u_R(d, \rho; X) := \rho V(d; X) - \underbrace{\eta \tilde{c}(\rho) \sigma^2(d; X)}_{=c(\rho, d; X)},$$

where $\tilde{c}(\rho) := (\text{erf}^{-1}(\rho))^2$. To obtain this expression, we replaced the question x by its sufficient statistics (d, X) .

⁸This is a direct consequence of the assumption that the decision maker only has access to \mathcal{F}_k when addressing problems. One rationale is a moral-hazard concern: science is complex, and it is impossible to distinguish the absence of a finding from the absence of a (proper) search.

The cost of research, $\eta c(\rho, d; X)$, derives from conceptualizing research as the search for an answer. We assume that, given a question x , the researcher chooses a sampling interval $[a, b] \in \mathbb{R}$ in the y -dimension. She discovers the answer if and only if $y(x) \in [a, b]$. We interpret the interval length as the amount of effort the researcher invests in finding an answer. For simplicity, we assume a quadratic cost $\eta(a - b)^2$.

LEMMA 1: *For knowledge \mathcal{F}_k , probability ρ , and question x , the minimal cost of obtaining an answer to question x with probability ρ is*

$$\eta c(\rho, d; X) = \eta \tilde{c}(\rho) \sigma^2(d; X).$$

The intuition behind Lemma 1 is straightforward. Conjectures are normally distributed and it is optimal to center the search effort around the mean. Given a success probability ρ , the more precise the conjecture, the shorter the expected search around the mean, and thus, the lower the cost. For a given variance, the higher the desired success probability, the larger the respective sampling interval.

Our cost function exhibits the following properties: It is (i) multiplicatively separable in ρ and $(d; X)$, (ii) increasing in d and X , and (iii) concave in d ; the concavity decreases in X and becomes linear in d in the limit as $X \rightarrow \infty$.⁹

The cost of research links output and novelty. To see this, consider a researcher who chooses a question x and aims to discover its answer with probability ρ . If that researcher increases ρ by a given amount, ε , she needs to increase her search effort, that is, expand her sampling interval, $[a, b]$, which is optimally centered around the mean. The additional effort required to increase ρ depends on the variance of the conjecture about $y(x)$, and hence, on the novelty of the question. If the variance is low, the density at the boundaries of the sampling interval is high—a moderate increase in effort gets the success probability to $\rho + \varepsilon$. If, instead, the variance is high, the density is low and the increase in effort must be larger.

4.2. The Researcher's Choice

We now characterize the researcher's optimal choice and elaborate on the resolution of the novelty-output trade-off. The researcher solves

$$\max_{X \in \{X_0, \dots, X_k\}} \underbrace{\max_{\substack{d \in [0, X/2], \\ \rho \in [0, 1]}} \rho V(d; X) - \eta c(\rho, d; X)}_{=: U_R(X)}.$$

Without cost ($\eta = 0$), we can apply Proposition 1 to derive the researcher's optimal choice. For any research area of length X , the researcher selects distance $d^0(X)$ and discovers an answer with certainty.

However, for positive cost, $\eta > 0$, the researcher's optimal decision on output, $\rho^\eta(X)$, is nontrivial and linked with her decision on novelty, $d^\eta(X)$. Choosing a question close to existing knowledge allows for a high probability of discovery at a low cost. The researcher's initial conjecture about the answer is already precise. Nevertheless, her payoff is low, as such a discovery provides little benefits. By increasing the distance, the

⁹More complex cost functions are possible, and we discuss two alternatives in the Supplemental Appendix (Carnehl and Schneider (2025)).

researcher increases both benefits and cost, *ceteris paribus*. The effect on the optimal probability of discovery is ambiguous. Depending on which effect dominates, the distance and the probability of discovery are substitutes (the researcher optimally reduces the success probability when answering a more novel question) or complements (she increases the success probability when answering a more novel question).

Optimal Choice Within a Research Area. The following proposition captures the key aspects of the researcher's choice within a research area.

PROPOSITION 3: *Suppose $\eta > 0$. Researchers fail with positive probability, $\rho^\eta(X) \in (0, 1)$. There is a cutoff area length $\tilde{X}^\eta < \tilde{X}^0$ such that researchers choose the benefits-maximizing distance, $d^\eta(X) = d^0(X)$, in area X if and only if $X \leq \tilde{X}^\eta$. Otherwise, researchers choose a question strictly less novel than that, $d^\eta(X) < d^0(X)$.*

Proposition 3 shows that when expanding knowledge the researcher chooses a question closer to existing knowledge than the benefits-maximizing distance $3q$. This is because novelty and output are substitutes. The marginal cost of increasing ρ rises with d , while the marginal benefits of increasing d approach zero as $d \rightarrow 3q$. The researcher balances novelty and output and selects a question less novel than $3q$.

That trade-off is less pronounced when the researcher deepens knowledge. The reason is that inside an area moving away from one boundary implies moving closer to the other boundary. Thus, the marginal cost of the success probability flattens in distance and becomes zero at $d = X/2$. Whether this effect is strong enough to make novelty and output complements depends on area length.

Optimal Choice Among Areas. The following proposition characterizes the researcher's optimal choice among intervals.

PROPOSITION 4: *Suppose $\eta > 0$. There are cutoffs $2q < \hat{X}^\eta < \dot{X} < \check{X}^\eta < \tilde{X}^\eta < 8q$ such that the following claims hold:*

1. *The researcher expands knowledge if and only if knowledge is dense, that is, if and only if all bounded areas are shorter than $\hat{X}^\eta < \hat{X}^0$.*
2. *Both novelty and output are nonmonotone in area length. Novelty attains a maximum at \tilde{X}^η . Output attains a maximum at \dot{X} .*
3. *The researcher's expected payoff from conducting research in an area X , $U_R(X)$ is single peaked and attains a maximum at \tilde{X}^η .*

Proposition 4 shows that the pattern in distance is qualitatively the same as in Proposition 2. However, the cost adds another dimension: the researcher's choice of success probability interacts with both the choice of distance and research area.

Consider a short area. The scope of improving the decision-maker's policies is small because conjectures are already precise and investing in discovery thus provides a limited payoff. Consequently, the researcher does not invest much in the search for an answer despite the low cost. She opts for a low success probability. A marginal increase in the area length provides larger increase in the benefits than in the cost. In response, the researcher increases both distance and success probability.

By contrast, consider a large area. The benefits of discovery trump those of the small area, yet the cost is higher. The researcher does not invest much in discovery due to the high marginal cost of increasing the success probability. As a result, the probability

of discovery is low. If, in this case, the area length increases marginally, the researcher responds by decreasing both the distance and the success probability.

Finally, consider an area of intermediate length. The benefits of discovery are high, yet the associated cost is limited. The return on investment is large, and the probability of discovery is high. As the area length further increases, the marginal return of increasing the distance declines, while the marginal cost rises. Eventually, the researcher faces a trade-off: should she reduce the success probability to maintain maximal distance? It turns out that she should. While the researcher wants to remain at a boundary in her choice of distance, she mitigates the increased cost by lowering the success probability.

The researcher's preferred area length, \tilde{X}^η , is in a region in which a trade-off between output and novelty exists. While the researcher would prefer a larger research area to increase the benefits of research, she would prefer a smaller research area to reduce her cost. Thus, distance is increasing and the success probability is decreasing at the point at which the researcher's payoff is maximal.

Thus far, we have not taken into account which research areas are available, which is determined by existing knowledge \mathcal{F}_k . Computing the optimal among the available areas is straightforward.

5. THE EVOLUTION OF KNOWLEDGE, MOONSHOTS, AND RESEARCH CYCLES

In this section, we consider a dynamic extension of our baseline model to study the endogenous evolution of knowledge. Through the lens of our dynamic model, we find that forward-looking interventions, which induce research cycles through moonshot discoveries, can improve the evolution of knowledge.

5.1. Sequential Research

Our starting point is a setting in which knowledge is $\mathcal{F}_0 = \{(x_0, y(x_0))\}$. At any time $t = 1, 2, \dots$, a short-lived researcher R_t arrives, observes current knowledge, \mathcal{F}_{t-1} , and selects (x, ρ) . If a discovery occurs, knowledge updates and a decision maker updates the application of knowledge accordingly.

To retain focus, we make three assumptions: All researchers have the same cost parameter $\eta > 0$, break ties identically, and condition their decision (x, ρ) *only* on current knowledge. While we impose the first two assumptions for simplicity only, the last assumption is more meaningful.¹⁰ We discuss our modeling choice in Section 6.

Moving forward, keep in mind that, given our symmetry assumptions, we can invoke Proposition 4 to describe knowledge as *dense* whenever \mathcal{F}_k is such that all bounded areas are shorter than \hat{X}^η . For any dense knowledge \mathcal{F}_k , each researcher's incentives are identical and in particular identical to the incentives of a researcher facing the initial knowledge \mathcal{F}_0 . As a benchmark, we first show that knowledge evolves in knowledge-expanding steps.

PROPOSITION 5: *For any $\eta \geq 0$, knowledge is dense in any period. Specifically, every researcher aims to expand knowledge by the same distance $d^\eta(\infty) \in (2q, 3q]$ with the same probability of discovery $\rho^\eta(\infty) \leq 1$. Moreover, both the distance $d^\eta(\infty)$ and the probability of discovery $\rho^\eta(\infty)$ strictly decrease in the cost parameter η .*

¹⁰In particular, once a researcher fails, all subsequent ones ignore those failures because they condition their decision only on current knowledge, and—by symmetry—fail again.

Proposition 5 shows that no short-lived researcher endogenously inspires future researchers to deepen knowledge. The intuition follows from Propositions 1 and 4: Researchers are rewarded for their immediate contribution to the value of knowledge. Therefore, no researcher expands knowledge beyond the benefits-maximizing $d^0(\infty) = 3q$. Deepening knowledge within areas of $X \leq 3q$, however, is unattractive as it creates too little value. Thus, all researchers aim to expand knowledge, and the stepsize depends on η . The larger η , the smaller the stepsize $d^\eta(\infty)$. Moreover, if $\eta > 0$, researchers fail at each step with probability $1 - \rho^\eta(\infty) > 0$ if the Brownian takes an unlikely turn. If one researcher fails, all later researchers fail, too. By symmetry, they make the exact same choice as the initially failing researcher and thus fail to discover the realization.

5.2. Interventions

In reality, we see that societies invest in affecting scientists' choices. The most prominent incentive schemes are large ex ante cost reductions through grants from funding institutions, such as the NSF or the ERC, and high-prestige ex post rewards for successful research, such as the Nobel Prize or the Fields Medal. These incentives, however, are often awarded to scientists pushing the frontier considerably.

A natural question is whether a designer (a funder, a government, ...) has an incentive to interfere with the knowledge production process in our model. Here, we consider a designer whose per-period payoff is a weighted average of the decision makers' and the researchers' payoffs,

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} ((1-\alpha)v(\mathcal{F}_{t+1}) + \alpha(v(\mathcal{F}_{t+1}) - v(\mathcal{F}_t) - \eta\hat{c}(d, \rho))) \right], \quad (1)$$

where $\alpha \in [0, 1]$ measures the designer's weight on the researchers' payoffs. This formulation introduces two potential frictions between the designer and the researcher.

First, their per-period payoffs differ. The only meaningful difference between the researcher's and the decision-maker's payoff is that the researcher bears the cost of research.¹¹ Therefore, we can think of the designer as having a modified cost parameter $\hat{\eta} \in [0, \eta]$, where $\hat{\eta} = 0$ corresponds to full alignment with the decision maker, and $\hat{\eta} = \eta$ to full alignment with the researcher. We can interpret $\hat{\eta}$ either as a degree of research cost internalization or, alternatively, as a measure of the appropriability of the benefits of research. Second, the time horizons of the researcher and the designer differ. Researchers are rewarded for their discoveries only through the immediate benefits they generate in the value of knowledge. However, a researcher's discovery alters the landscape of knowledge indefinitely and thereby the future value of knowledge as well. We assume that the designer discounts the future value of knowledge by $\delta \in [0, 1)$.

For now, we assume that the designer can intervene at most once and proceed by studying a set of natural benchmarks to understand the forces at play. First, we isolate the time-horizon friction by assuming that research is costless ($\eta = 0$) and that the designer is forward-looking ($\delta > 0$). Second, we isolate the cost friction by considering positive research cost ($\eta > 0$) but a myopic designer ($\delta = 0$).

¹¹Note that the term $\alpha v(\mathcal{F}_t)$ is, within a period, a constant as the knowledge at the beginning of period t , \mathcal{F}_t is exogenously given.

Forward-Looking Designer and Costless Research. First, consider a forward-looking designer, $\delta \in (0, 1)$, in an environment with costless research, $\eta = 0$. Note that this designer's payoff is

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \delta^{t-1} (v(\mathcal{F}_{t+1}) - \alpha \cdot v(\mathcal{F}_t)) \right]. \quad (2)$$

A researcher without cost chooses $d^0(\infty) = 3q$ and $\rho^0(\infty) = 1$ (Proposition 2). Moreover, by Proposition 5, knowledge remains dense under these choices. From Proposition 2, we know that adding an area of length $3q$ to any given set of areas \mathcal{X}_k dominates adding a single bounded area of any other length. Because every researcher adds such an area, it follows that knowledge expands in $3q$ -steps and without failure. Irrespective of (α, δ) , the designer has no reason to intervene, and a sequence of short-lived researchers without cost implements the forward-looking designer's optimum. Thus, it is not the researcher's short-livedness alone that causes inefficiency in the evolution of knowledge.

Myopic Designer and Costly Research. Second, consider a myopic designer, $\delta = 0$, in an environment with costly research, $\eta > 0$, that is not fully aligned with the researcher, $\alpha < 1$. Such a designer has, in any period t , payoffs similar to this period's researcher, albeit with a smaller cost parameter, $\hat{\eta}$. It follows that, in each period, the researcher's choices are suboptimal from the designer's perspective. In particular, the designer prefers a larger distance and a higher success probability.

Straightforwardly, such a myopic designer wants to align the researcher's cost parameter, η , with his own, $\hat{\eta}$, whenever given the chance. If the designer aligns the cost, the researcher implements the designer's preferred research. Whether the designer aligns η and $\hat{\eta}$ depends on whether she needs to pay an additional cost to do so. If not, the designer intervenes by reducing the researcher's cost parameter directly, yet the evolution of knowledge remains qualitatively unchanged.¹² Knowledge evolves through step-by-step expansions albeit with larger stepsizes and fewer failures. This observation also provides an answer to the question of how heterogeneous researchers would change the picture. In fact, heterogeneous researchers are researchers that incorporate the cost at different weights, thus observationally indistinguishable from cost-internalizing designers.

To achieve a better outcome, the designer could, for example, award a grant to the researcher, which effectively lowers that researcher's cost by allowing her to hire personnel or reduce her teaching load. However, such subsidies may come at a cost to the designer not reflected in (1). If such a cost was present, the designer would subsidize the researcher only until the marginal cost of doing so exceeds the marginal benefits.¹³

Moonshots. The preceding benchmarks provide a rationale for moderate interventions at most that encourage researchers to invest more into the search for answers to more distant questions. These observations highlight that it is neither the researcher's short-livedness nor her cost of research alone that warrants more substantial interference with the scientific production process. To rationalize more drastic interventions that qualitatively change the evolution of knowledge, we now turn to a forward-looking designer in an environment with research cost.

¹²Because η governs the relative weight of benefits and cost, increasing the researcher's benefits or decreasing her cost is the same.

¹³A realistic model would need to develop a theory of supply of grants, awards, or teaching reductions, which we expect to result in a non-trivial cost function to the designer. We thus leave it for future research.

We focus on a particular simple model in which the designer does not internalize the researcher's cost ($\alpha = 0$). We allow the designer to fully control the first researcher's actions at no cost. We are interested in whether such a designer has an incentive to induce a *moonshot*, a discovery more distant than the myopically optimal distance, $3q$, to the knowledge frontier. Because a distance of exactly $3q$ maximizes the value of the area generated (by Proposition 2), a moonshot can only be beneficial if it “inspires” future researchers to deepen knowledge in the newly created area; that is, if it is larger than \hat{X}^η . The next proposition states that a designer finds it optimal to launch a moonshot if and only if he is sufficiently patient and the researcher's cost friction is intermediate.

PROPOSITION 6: *Suppose $\alpha = 0$. There are cost parameters $0 < \underline{\eta} < \bar{\eta} < \infty$ and a critical discount factor $\underline{\delta}(\eta) < 1$ such that for $\eta \in (\underline{\eta}, \bar{\eta})$ the designer optimally launches a moonshot if and only if $\delta > \underline{\delta}(\eta)$. If $\eta = 0$ or $\eta \rightarrow \infty$, a moonshot is suboptimal for any δ .*

Proposition 6 incorporates the no-cost case. As we have argued above, moonshots are of no use in that case. The same is true if the researchers' costs are high. Then ρ^η is low, and already the second researcher most likely fails to discover an answer. Thus, the designer expects little progress from $t = 2$ onward and focuses on maximizing the value of knowledge at the end of $t = 1$ inducing the myopic optimum $d = 3q$ which, unless $\eta = 0$, still implies more novelty than what the researcher chooses herself.

For intermediate cost parameters, however, moonshots generate valuable spillovers. Without intervention, intermediate cost parameters lead to successful but too narrow research that fails too often. Researchers expand knowledge in a ladder-type structure but select questions too close to existing knowledge and too low success probabilities. In anticipation, the designer initiates a *research cycle* through a moonshot. The next researcher builds on both the initial knowledge \mathcal{F}_0 and the moonshot discovery and deepens knowledge in the newly generated area in an attempt to close the gap between moonshot and initial knowledge.

A completed research cycle may lead to a more valuable landscape as Figure 5 illustrates. Once the cycle is complete and knowledge is dense again, knowledge is in a better state than with one optimal area of length $3q$ and subsequent knowledge-expanding steps following Proposition 5. As time moves on, the better landscape remains and generates further rents.

Figure 6 illustrates a second positive effect of a moonshot. The probability that the evolution of knowledge ends because the next researcher fails is smaller after the moonshot than after a myopic disclosure. The moonshot opens a research area of considerable size. In $t = 2$, the second researcher aims to fill it, but can use both the initial frontier and the moonshot discovery to fine-tune her research. This reduces her cost—in particular

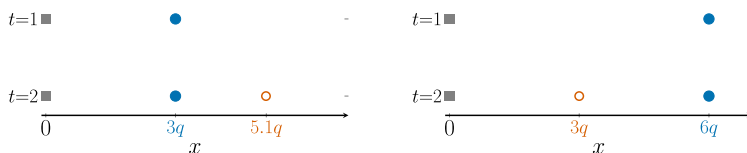


FIGURE 5.—Evolution of knowledge for different choices in $t = 1$. The dots show which questions have a known answer at each time t , assuming that discovery has been successful. The designer's choice for R_1 (●) is given, R_2 's choice (○) is a best response. $F_0 = 0$, $y(0)$, $\eta = 1$. The left panel assumes discovery of question $x = 3q$ in period $t = 1$, the right panel $x = 6q$.

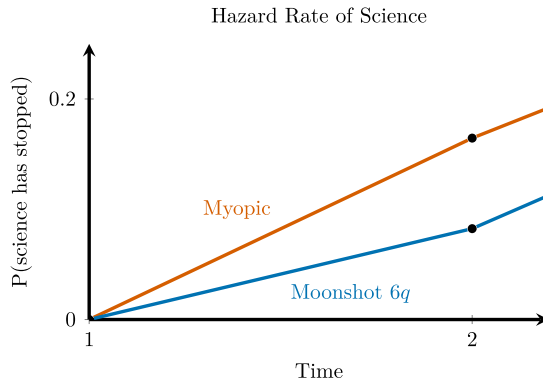


FIGURE 6.—Hazard rate of science. Cumulative probability that science stops by time t for different initial disclosures given $\eta = 1$.

when the moonshot is not too far from initial knowledge. As a consequence of this logic, “marsshots” very far from existing knowledge are not optimal. They are too disconnected to inspire the following researchers to be productive.

While moonshots always reduce the hazard rate of science, sometimes they do not lead to a better landscape conditional on success. Suppose, for example, that η is close to but below $\bar{\eta}$ from Proposition 6 (and δ not too small). Then the risk of not producing anything beyond the first discovery is large. However, a moonshot significantly reduces the risk of failure in $t = 2$, in particular, if it is only marginally above the cutoff for deepening research, \hat{X}^η . Such a moonshot does not improve the landscape, but focuses on making R_2 succeed. If R_2 completes the cycle, the knowledge landscape is worse than completing the counterfactual two steps $3q$ (through the designer) and $d^\eta(\infty)$ (through R_2). Yet, because R_2 becomes more likely to succeed after a moonshot, the moonshot remains profitable. If, instead, η is small but above $\underline{\eta}$, a moonshot will provide benefits both in output and in the value of knowledge.

Taking Costs Into Account. One, seemingly crucial, assumption of our moonshot analysis is that the designer does not bear the costs of the initial moonshot or does not directly care about the costs exerted by future researchers. Taking these costs into account complicates the model significantly, as the designer needs to trade off several paths with potentially different costs. Analytically, we consider it beyond scope for this paper. However, numerically, it turns out that incorporating these costs and/or the costs of future researcher generations does not affect Proposition 6.¹⁴

Research Cycles. As we have seen above, with infrequent interventions, a designer may improve the evolution of knowledge by inducing a research cycle through an initial moonshot, which provides a higher payoff than the same number of steps when expanding knowledge stepwise. A natural question to ask is what happens if the designer has multiple but infrequent opportunities to intervene. There are several ways to model infrequent interventions. We opt for the simplest, which is that, in each period, the designer gets to

¹⁴This and other numerical results we allude to here can be obtained using the code we provide as Supplementary Material to this paper. Since our static model has effectively only one (relevant) parameter, η , we can provide comprehensive numerical results.

control the researcher with probability $\lambda > 0$. For the sake of clarity, we are particularly interested in the case in which such opportunities are rare. This assumption allows us to circumvent that future opportunities play a large role in the designer's decision today, which would obfuscate our trade-off of interest.

For our analytical result, we consider environments that (i) are *promising* in the sense that the discount factor δ and the cost parameter η are such that $\delta\rho^\eta(\infty) > 1/2$, and (ii) feature optimal one-time moonshots of intermediate length, that is, the optimal moonshot has length $X \in [4q, \min\{2\hat{X}^\eta, \tilde{X}^\eta\}]$.¹⁵ Then, after a moonshot in t , the next researcher, R_{t+1} , aims to complete the cycle and chooses the midpoint of the area created by that moonshot.

Having fixed the environment for the single moonshot, continuity implies that, for rare moonshot opportunities, the designer starts a moonshot whenever given the opportunity as long as science did not get stuck before due to failed research. However, the question is less clear when science is stuck. Specifically, we may consider the following scenario. The designer had created a moonshot in the past, but due to an unexpected turn of the truth, the researchers have failed to complete the research cycle. In that case, should the designer abandon the incomplete cycle and start a new one, or, should he use his opportunity to complete the incomplete cycle?

PROPOSITION 7: *Assume $\alpha = 0$ and a promising environment with moonshots of intermediate length. The designer strictly prefers completing an open cycle over initiating a new research cycle when interventions are sufficiently rare.*

The intuition behind Proposition 7 is straightforward. The environment is promising and researchers are unlikely to fail while expanding research. If, in addition, the designer is patient, his effective discount factor is high. A high effective discount factor implies that leaving a cycle incomplete leads to a severe reduction in long-run research output. Moreover, the immediate payoff from completing a cycle is large.

Providing analytical solutions for the general case proves difficult, because the researcher's problem is two-dimensional and the disclosure probabilities $\rho^\eta(\infty)$ can only be implicitly defined. Numerically, however, we can show that the logic of Proposition 7 holds if the initial optimal moonshot is sufficiently large. The reason is intuitive. If the moonshot is large, there are two benefits. First, once filled up, the knowledge landscape is better than if it evolved through knowledge-expanding steps. Second, a moonshot lowers the probability that researchers get stuck by increasing the success probability of the next researchers. If, instead, the moonshot is short, the sole reason behind the moonshot is to lower the risk of researchers getting stuck. But then, if the optimal moonshot is short, researchers are likely to fail outside of moonshots. Thus, completing the cycle provides little benefits and likely leads to immediate failure thereafter. A new moonshot, on the other hand, improves the chances of a discovery in the next period. Thus, the designer opts for the latter option.

6. FINAL REMARKS

We propose a tractable and flexible model based on three simple premises: (i) the pool of available research questions is large, (ii) questions close to existing knowledge are eas-

¹⁵Numerically, there is a wide range of parameters (δ, η) satisfying both conditions. For example, when $\delta = 0.95$, $\eta \in [0.01, 0.5]$ suffices. The lower bound, $4q$, is directly implied by the environment being promising.

ier to answer than questions far from existing knowledge, and (iii) society applies knowledge when selecting policies. Our model endogenously links novelty and research output and highlights the importance of existing knowledge for research and knowledge accumulation. A dynamic extension delivers rich insights into how interventions can improve the accumulation of knowledge over time.

We close with a discussion along two dimensions: First, we revisit our modeling choices, and discuss extensions and alternative assumptions. Then we discuss implications of our findings for designing institutions and research environments.

6.1. *Modeling Choices*

To model the evolution of knowledge as close as possible to the static setting, we made a set of modeling assumptions that may appear strong. We now revisit these assumptions and briefly discuss alternatives.

Beyond the Real Line. For illustrative purposes, we assume that the set of questions is the real line, which implies that there are exactly two directions in which the research frontier may be expanded. In reality, we would expect that, at least in some fields, there is a plethora of directions in which a given researcher could expand the frontier. It turns out, however, that as long as the partitioning of the question space into areas is attainable, our analysis and findings remain unchanged. We provide a formal extension in the Supplemental Appendix (Carnehl and Schneider (2025)). There, we also discuss the case of “seminal discoveries” that open up new fields of questions.

Observing Failure. In the limit, when researchers can select from a continuum of directions, it becomes hard for others to find previous failed attempts. For example, suppose that there is some cost of finding out whether a past researcher has done expanding research in the direction the current researcher contemplates. Because there is an entire ocean of potential directions, the probability that someone has worked in the contemplated direction is zero. Hence, it is not worth paying the cost and past failures remain unobserved.

While our baseline model is perhaps too restrictive on the set of directions, an ocean may be too open. Thus, indeed, there are settings in which observing past failure becomes relevant. Generally, observing a failed attempt provides information about an unexpected turn of the Brownian motion and, therefore, increases the variance at that point. Assuming some coherence in the search for an answer (e.g., by restricting search to a connected interval of answers), digging deeper on previous failed attempts proves unattractive: The failed attempt has revealed that the answer is complicated. In response, researchers would aim to answer a different question rather than resolving that failed attempt.

Getting Knocked off the Ladder. Proposition 5 shows that knowledge expands in steps without interventions. Because $\hat{X}^\eta > 3q$ for all η , researchers, even if heterogeneous, will not deviate from the ladder structure. Instead, with heterogeneous researchers, we would only get ladders with unequal step sizes.

Serendipity, on the other hand, may lead to researchers leaving the ladder for some time. Imagine some random discovery far from existing knowledge. Then the next generations will work to connect it to the existing bulk of knowledge by Proposition 4 until knowledge is dense again. Only then do researchers return to the knowledge-expanding steps of Proposition 5.

Other forces with the potential to break the ladder structure are long-lived researchers who would act similar to a cost-internalizing designer and have an incentive to set moonshots, or exogenous shocks to the importance of questions. In our baseline, all questions have the same relevance. However, real-world shocks may cause an elevated interest in particular questions. Our model predicts that such temporary interest can have long-lasting effects. Once the interest fades, researchers use the improved conjectures to bridge old knowledge and new findings.

Other Frictions. We assume that, besides the cost friction, the market for ideas works well and researchers are paid their marginal contribution to the value of knowledge. In reality, several known frictions hinder that process. Well-known examples include publication bias (Andrews and Kasy (2019)), the emphasis on priority (Hill and Stein (2025), Bobtcheff, Bolte, and Mariotti (2017), Hill and Stein (2024)), or researchers' career concerns (Akerlof and Michailat (2018), Heckman and Moktan (2020)). While the question of optimal market design is beyond our scope, our framework is flexible enough to incorporate these frictions. It may thus be a stepping stone toward structural models of science funding that include such frictions.

6.2. Implications

Although our dynamic model is stylized, we reconcile empirical findings in the economics of science (see, e.g., Rzhetsky et al. (2015), Fortunato et al. (2018), documenting a lack of novelty in research) and the economics of innovation (e.g., Jaffe, Fogarty, and Banks (2003), documenting the inspiring nature of the original moonshot). In reality, a researcher's value of any given discovery, especially in the basic sciences, depends on the institutional framework a researcher operates in. Our findings in Section 5 suggest that focusing only on immediate policy relevance when designing researchers' incentives is suboptimal for patient societies. In the following, we discuss some alternative incentive structures.

Future-Oriented Rewards. No short-lived researcher selects a question novel enough to initiate a research cycle, even when research is costless to her. The reason is that researchers exclusively care about the value of their research today but not about its indirect value in guiding future researchers.

To incentivize moonshot discoveries, researchers need future-oriented rewards that go beyond the instantaneous value of their findings. Apart from prizes for novel findings, the value of citations for promotion decisions or scientific reputation may serve such purpose. That insight is reminiscent of recent empirical work on firm-level R&D. Frankel, Krieger, Li, and Papanikolaou (2024) estimate the value of dynamic spillovers from discoveries of drugs. In line with our model, they provide suggestive evidence that the lack of appropriability of these spillovers harms novelty in pharmaceutical innovation.

Research Consortia. The idea of research consortia has been put forward in some fields of basic science to improve the evolution of knowledge. Research consortia formed by scientists of different backgrounds operate on missions different from the "publish or perish" or "marketability" paradigms. Hill and Stein (2024) document that the incentives and choices of consortia in structural biology differ from those of university researchers. In line with our model, they find that consortia provide more novel discoveries but also discoveries of lower immediate value.

Our findings suggest that the underlying reason is not particular to structural biology. Establishing institutions that alleviate some researchers from the need to provide immediate benefits can guide those in traditional incentive schemes. In fact, our findings in Section 5 suggest that a mix in incentives may be key to improving the evolution of knowledge.

Direction of Science. As we noted above, funding measures that only target the researcher's cost do not alter the way knowledge progresses qualitatively. However, it may alter the direction of science. Recall that a researcher's question choice in our model can be interpreted as choosing from a set of directions and then picking novelty along that direction. Reducing the cost of research in one of the available directions distorts the marginal researcher in favor of the subsidized direction.

Recent empirical work is consistent with this observation. Nagaraj and Tranchero (2024) suggest that data availability affects the direction of sciences. Kim (2023) finds that novel technologies can incentivize researchers to focus on more explored rather than unexplored areas. Myers (2020) shows that topic-specific grants can lead to a change in the direction of science. These findings can be interpreted in our model as a change in the cost of research: some directions become less cost intensive than others.

When a researcher chooses a moonshot in some direction, she affects future generations through the implied research cycle, suggesting that the effect is persistent. Moonshots determine the direction science takes in the medium run. That coordination may provide additional network benefits, which are currently outside our model. A potential counterforce has been identified in the literature in corporate R&D, which emphasizes the role of competition.¹⁶ While beyond our scope, combining competition and dynamic spillovers within our framework is an exciting avenue for future research.

APPENDIX A: PROOFS

Pure algebraic reformulations are relegated to the Supplemental Appendix (Carnehl and Schneider (2025)). There, we also review the properties of $\tilde{c}(\rho) = (\text{erf}^{-1}(\rho))^2$. We use subscripts to denote partial derivatives; $\frac{df(x,y)}{dx}$ for the total derivative; and omit function arguments when clarity is preserved.

A.1. Proof of Proposition 1

The value of knowing \mathcal{F}_k is $\int_{-\infty}^{\infty} \max\{1 - \sigma_x^2(y|\mathcal{F}_k)/q, 0\} dx$. No matter which point of knowledge $(x, y(x))$ is added to \mathcal{F}_k , the value of knowledge outside the frontiers is identical for both \mathcal{F}_k and $\mathcal{F}_k \cup \{(x, y(x))\}$. Area lengths $X_1 = X_k = \infty$ do not depend on \mathcal{F}_k and neither does the variance for a question $x < x_1$ or $x > x_k$ with a given distance d to \mathcal{F}_k . The conjectures about all questions outside $[x_1, x_k]$ deliver a value of $2 \int_0^q (q-x)/q dx = q$, which is independent of \mathcal{F}_k .

Moreover, if a question $\hat{x} \in [x_i, x_{i+1}]$ is answered, it deepens knowledge and only affects questions in that area, that is, $G(x|\mathcal{F}_k) = G(x|\mathcal{F}_k \cup \{(\hat{x}, y(\hat{x}))\}) \forall x \notin (x_i, x_{i+1})$.

The value of an area $[x_i, x_{i+1}]$ is (with abuse of notation)

$$v(X) = \int_0^X \max\left\{\frac{q - \frac{d(X-d)}{X}}{q}, 0\right\} dd.$$

¹⁶See, for example, Bryan and Lemus (2017), Hill and Stein (2024, 2025), Hopenhayn and Squintani (2021).

Note that whenever $X \leq 4q$, $d(X - d)/X \leq q$. Hence, we can directly compute the value of any area with length $X \leq 4q$ as $v(X) = X - X^2/(6q)$. Whenever $X > 4q$, a positive value is generated only on a subset of points in the area. As the variance is a symmetric quadratic function with midpoint $X/2$, there is a symmetric area around $X/2$, which has a variance exceeding q . On those points, the decision-maker's losses are limited to zero. The points with variance equal to q are $\bar{d}_{1,2} = X/2 \pm 1/2\sqrt{X}\sqrt{X - 4q}$. Hence, the value of an area with $X > 4q$ is (due to symmetry)

$$v(X) = 2 \int_0^{\bar{d}_1} \frac{q - \frac{d(X-d)}{X}}{q} dd = X - \frac{X^2}{6q} + \frac{X-4q}{6q} \sqrt{X}\sqrt{X-4q}.$$

If knowledge expands, a new area is created and no area is replaced. The value created is

$$V(d; \infty) = v(d) = d - \frac{d^2}{6q} + \begin{cases} 0 & \text{if } d \leq 4q, \\ \frac{d-4q}{6q} \sqrt{d}\sqrt{d-4q} & \text{if } d > 4q. \end{cases}$$

If a knowledge point is added inside an area with length X with distance d to the closest existing knowledge, it generates two new areas with length d and $X - d$ that replace the old area with length X . The total value of the two new intervals is

$$\begin{aligned} v(d) + v(X-d) &= d - \frac{d^2}{6q} + \begin{cases} 0 & \text{if } d \leq 4q, \\ \frac{d-4q}{6q} \sqrt{d}\sqrt{d-4q} & \text{if } d > 4q \end{cases} \\ &\quad + X - d - \frac{(X-d)^2}{6q} \\ &\quad + \begin{cases} 0 & \text{if } X-d \leq 4q, \\ \frac{X-d-4q}{6q} \sqrt{X-d}\sqrt{X-d-4q} & \text{if } X-d > 4q. \end{cases} \end{aligned}$$

The benefits of discovery are then $V(d; X) = v(d) + v(X-d) - v(X)$. Noticing that $\sigma^2(d; X) = d(X-d)/X$ and replacing results in the expression in the proposition. Taking the limit $X \rightarrow \infty$ implies the expanding value.

A.2. Proof of Proposition 2

Expanding Knowledge. The first-order condition for $d \leq 4q$ immediately delivers $d = 3q$. Moreover, the benefits decrease in d for $d > 4q$, which can be seen from the negative d -derivative (invoking Lemma 27 in the Supplemental Appendix (Carnehl and Schneider (2025)))

$$V_d(d; \infty | d > 4q) = -\frac{d}{3q} + 1 + \sqrt{\frac{d-4q}{d}} \frac{d-q}{3q} < 0.$$

Deepening Knowledge.

LEMMA 2: $d^0(X) = X/2$ if $X \leq 6q$.

PROOF: 1. Assume $X \leq 4q$. The benefits of discovery are $V(d; X|X \leq 4q) = (Xd - d^2)/(3q)$, which increase in d for $d \in [0, X/2]$ and are maximized at $d = X/2$. Moreover, $V(X/2; X|X \leq 4q) = X^2/(12q)$, which increases in X .

2. Assume $X \in (4q, 6q]$. Case (i). $d \geq X - 4q$ implies (since $d \leq 3q$)

$$V(d; X) = \frac{1}{6q}(2dX - 2d^2 - \sqrt{X}(X - 4q)^{3/2}),$$

which are the same as in the first case up to the constant $-\sqrt{X}(X - 4q)^{3/2}$. Thus, the optimal d conditional on $d \geq X - 4q$ is $d = X/2$.

Case (ii). For $d \leq X - 4q$, the benefits and their derivative are

$$V(d; X) = \frac{2dX - 2d^2 + \sqrt{X-d}(X-d-4q)^{3/2} - \sqrt{X}(X-4q)^{3/2}}{6q},$$

$$V_d(d; X) = \frac{1}{3q} \left(X - 2d - (X-d-q) \sqrt{\frac{X-d-4q}{X-d}} \right) > 0.$$

Lemma 28 in the Supplemental Appendix (Carnehl and Schneider (2025)) implies the last claim. Then $V_d(d; X|d \leq X - 4q, X \in [4q, 6q]) > 0 \forall d, X$ in the considered domain; $d = X - 4q$ maximizes $V(d; X|d \leq X - 4q, X \in (4q, 6q])$ and by (i) $d = X/2$ maximizes $V(d; X|X \in (4q, 6q])$. *Q.E.D.*

LEMMA 3: For any $X < \infty$, $V_d(X/2; X) = 0$.

PROOF:

$$V(d; X) = \frac{1}{6q} \underbrace{(2X\sigma^2(d; X))}_{(I)} + \underbrace{\mathbf{1}_{d>4q} \sqrt{d}(d-4q)^{3/2}}_{(II)} + \underbrace{\mathbf{1}_{X-d>4q} \sqrt{X-d}(X-d-4q)^{3/2}}_{(III)} - \underbrace{\mathbf{1}_{X>4q} \sqrt{X}(X-4q)^{3/2}}_{(IV)}.$$

At $d = X/2$, either both (II) and (III) are active or neither is. Moreover, (IV) is independent of d , and we have $\partial(\text{II})/\partial d = -\partial(\text{III})/\partial d$, and $\partial(\text{I})/\partial d = 0$. *Q.E.D.*

LEMMA 4: If $X > 8q$, then $d^0(X) \neq X/2$. If $d^0(X) \neq X/2$, then $d^0(X) \leq 4q$.

PROOF: Consider $d = \bar{d} = 4q < X/2$. We obtain

$$V(\bar{d}; X|\cdot) - V(X/2; X|\cdot) = \frac{1}{6q} \frac{(X-8q)^{3/2}}{2} (2\sqrt{X-4q} - \sqrt{X} - \sqrt{(X-8q)}),$$

which is positive if $4(X-4q) > 2X-8q \Leftrightarrow X > 4q$, which holds by assumption.

To establish the second part of the lemma, note that $d > 4q$ only occurs for $X > 8q$. We will show that $V_d(d; X) < 0$ for all $d > 4q$ when $X > 8q$. Toward this, observe

$$V_{ddX}(d; X > 8q) = -\frac{4q^2}{(X-d)^{\frac{5}{2}}(X-d-4q)^{\frac{3}{2}}} < 0,$$

because $X > 8q$ and $X-d \geq X/2 > 4q$. Thus, a lower bound for $V_{dd}(d; X > 8q)$ is

$$V_{dd}(d; X > 8q)|_{\lim_{X \rightarrow \infty}} = \frac{1}{3q} \frac{d^2 - d^{\frac{3}{2}}\sqrt{d-4q} - 2q(d+q)}{d^{\frac{3}{2}}\sqrt{d-4q}} > 0.$$

Thus, $V_d(d; X)$ is highest for $d = X/2$ which, by Lemma 3, is 0. Hence, $V(d; X)$ decreases in d for $X > 8q$ and $d > 4q$. Q.E.D.

LEMMA 5: $d^0(X) < X/2 \Rightarrow \frac{dV(d^0(X); X)}{dX} < 0$.

PROOF: By the envelope theorem, $\frac{dV(d^0(X); X)}{dX} = V_X(d^0(X); X)$, which is negative for $X \geq 4q$ and $d \in [0, X-4q]$ by Lemma 29 in Supplemental Appendix (Carnehl and Schneider (2025)). If $X \geq 8q$, then $d \leq X-4q$ by definition and Lemma 29 applies. If $X < 8q$, by Lemma 2, we know that $d^0(X) \neq X/2$ only if $X \geq 6q$. Moreover,

$$V_d(d; X|X/2 > d > X-4q, X < 8q) = \frac{X-2d}{3q} > 0$$

Hence, $d^0(X) \neq X/2 \Rightarrow d^0(X) \leq X-4q$. Lemma 29 applies proving Lemma 5. Q.E.D.

LEMMA 6: $d^0(X) < X/2$ for some $X \in [6q, 8q) \Rightarrow d^0(X) < X/2$ for all $X' > X$.

PROOF: It suffices to consider $X' < 8q$ by Lemma 4. We prove the claim by showing that $V(d_c^0(X); X)$ for any interior critical point $d_c^0(X) < X/2$ cuts $V(X/2; X)$ from below at any potential intersection. Thus, there is at most one switch from $d^0(X) = X/2$ to $d^0(X) < X/2$ and no switch back. Continuity then implies the statement.

$V(d; X)$ is a continuously differentiable function in X and d . Thus, any interior (local) optimum $d_c^0(X)$ is continuous and so are $V(d_c^0(X); X)$ and $V(X/2; X)$. We now show that if $V(d_c^0(X); X) = V(X/2; X)$ for some local optimum $d_c^0(X) < X/2$ and $X \in [6q, 8q]$, then $dV(d_c^0(X); X)/dX > dV(X/2; X)/dX$. Because the first-order condition holds for an interior critical point, the envelope theorem applies and $dV(d_c^0(X), X)/dX < 0$. By Lemma 2, we know that for $X \leq 6q$ $V(X/2, X) \geq V(d_c^0(X), X)$ if $d_c^0(X)$ exists. Hence, at the first intersection $V(d_c^0(X), X)$ must cross from above. Now, because $V(d_c^0(X), X)$ decreases, the first intersection can only occur in a region where $V(X/2, X)$ decreases and must be such that $dV(X/2, X)/dX < dV(d_c^0(X), X)/dX$. We prove that this is the only potential intersection in Lemma 30 in the Supplemental Appendix (Carnehl and Schneider (2025)). There we show that $d^2V(X/2, X)/(dX)^2 < 0$ and $d^2V(d_c^0(X), X)/(dX)^2 > 0$. Q.E.D.

LEMMA 7: $V(d^0(X); X)$ is continuous in X . As $X \rightarrow \infty$, it converges uniformly to $V(d; X)$ and $d^0(X) \rightarrow d^0(\infty)$. For any $X > 6q$, we have $d^0(X) > 3q$ and $V(d^0(X), X) > V(3q, \infty)$.

PROOF: Continuity follows because $V(d^0(X); X) = \max_d V(d; X)$ with $V(d; X)$ continuous in both $d \in [0, X/2]$ and X . Now take any sequence of increasing X_n with $\lim_{n \rightarrow \infty} X_n = \infty$. For any $\delta(d)$, $\exists n$ such that $V_n(d; X_n) - V(d; \infty) < \delta(d)$ as can be seen from the formulation in the proof of Proposition 1. Hence, $V(d; X_n)$ converges uniformly to $V(d; \infty)$. By uniform convergence, the maximizer $d^0(X_n)$ of $V(d; X_n)$ converges, too. Convergence from above follows from $V(3q; X) > V(3q; \infty)$ for $X \in (6q, \infty)$.

Finally, recall that $V(d; \infty)$ describes the value of an area of length d . That value increases for $d < 3q$ and decreases for $d > 3q$. Now suppose $X > 6q$ and $d^0(X) < 3q$. Then by increasing d both areas created get closer to $3q$, and thus, increase in value. A contradiction to $d^0(X)$ being the maximizer. Q.E.D.

LEMMA 8: $V(d^0(X); X)$ is single peaked with peak at $\check{X}^0 \approx 6.2q$ and $d^0(\check{X}^0) \approx 3.1q$.

PROOF: By Lemma 7, $V(d^0(X); X)$ is continuous. By Lemma 5, it is decreasing if $d^0(X) < X/2$. By Lemmata 2, 4, and 6, and especially the arguments used in the proof of Lemma 6, the (single) switch from $d^0(X) = X/2$ to $d^0(X) < X/2$ happens for some $X \in (6q; 8q]$ and at a point where $V(X/2; X)$ is already decreasing. Thus, we can compute the peak by considering the first-order conditions of $V(X/2; X)$ with respect to X . The peak is the (real) solution to

$$\frac{X}{X - q} = 2 \frac{\sqrt{X - 4q}}{\sqrt{X}}. \quad (3)$$

Defining $m := X/q$, the above reduces to $m/(m - 1) = 2\sqrt{(m - 4)/m}$. For $m > 4$, the LHS decreases and the RHS increases in m . The solution is $m \approx 6.204$. Q.E.D.

LEMMA 9: *Expanding trumps deepening knowledge if and only if $X < \hat{X}^0 \approx 4.338q$.*

PROOF: $V(3q; X) > V(3q; \infty)$ for $X \geq 6q$ by direct comparison. For $X \in [0, 6q]$, we need to consider only $d^0(X) = X/2$ by Lemma 2. We compare $V(X/2; X)$ with $V(3q; \infty)$. Using $m = X/q$ from the previous proof, the two functions intersect if $m^2/12 - \sqrt{m}/6(m - 4)^{(3/2)} - 3/2 = 0$, which has a unique solution such that $m \leq 6$ at $m \approx 4.338$. Q.E.D.

LEMMA 10: $4q < \hat{X}^0 < 6q < \check{X}^0 < \tilde{X}^0 < 8q$.

PROOF: The first two inequalities follow from Lemma 9, the third from Lemma 8. The fourth inequality follows from Lemma 6 and Lemma 4 implies the last. Q.E.D.

A.3. Proof of Proposition 3

Throughout, we make use of the first-order necessary conditions for interior solutions.

$$\eta \tilde{c}_\rho(\rho) \sigma^2(d; X) = V(d; X), \quad (\text{FOC}^\rho)$$

$$\rho V_d(d; X) = \eta \tilde{c}(\rho) \sigma_d^2(d; X). \quad (\text{FOC}^d)$$

Part 1: Expanding Knowledge.

LEMMA 11: *There is a nontrivial optimal choice with $\infty > d > 0$, $1 > \rho > 0$ on any interval with positive length, $X \in \mathbb{R}^+ \cup \{\infty\}$. The first-order condition, (FOC^ρ) , is necessary for optimality of $\rho^\eta(X)$.*

PROOF: The researcher can guarantee a nonnegative payoff by choosing $d = 0$ or $\rho = 0$. Hence, her value is bounded from below by zero. Next, note that $u_R(\rho = 0, d > \varepsilon; X) = 0$ for any $\varepsilon > 0$ and that $(\partial u_R(\rho, d; X))/(\partial \rho)|_{\rho=0, d=\varepsilon} = V(\varepsilon, X) > 0$ by Proposition 1. Therefore, on any interval X a maximum with $d > 0$, $\rho > 0$ exists.

By Lemma 8, the benefits of discovery are bounded $V(d, X) \leq M < \infty$ and $\lim_{\rho \rightarrow 1} \tilde{c}(\rho) = \infty$. Therefore, the optimal $\rho < 1$. Finally, $V(d, \infty)$ decreases in d for d large enough while the cost $\eta \tilde{c}(\rho) \sigma^2(d, \infty)$ increases in d . Hence, the optimal distance is bounded $d^\eta(\cdot) \leq D < \infty$.

Because the optimal choice is interior and the objective is continuously differentiable, a necessary condition for $\rho^\eta(X)$ is that it solves (FOC^ρ) . Q.E.D.

LEMMA 12: *When expanding knowledge, the optimal choice is characterized by the first-order conditions (FOC^d) and (FOC^ρ) . These FOCs suffice and $d^\eta(\infty) \in (2q, 3q)$. The researcher's value is strictly positive $U_R(\infty) > 0$.*

PROOF: We proceed in three steps. First, we show that the distance is at most $3q$. Second, we show that the first-order conditions are sufficient when expanding knowledge. Third, we characterize the optimal choice of the researcher.

Step 1. $d \leq 3q$. Fix any $\rho \geq 0$. Since $\sigma^2(d; \infty)$ increases in d , it is immediate that the researcher's utility is nonincreasing in d if $V(d; \infty)$ decreases in d . Combining this observation with Proposition 2, it is sufficient to restrict attention to $d \leq 3q$.

Step 2. FOCs sufficient. By Lemma 11, the researcher's optimal choice is interior, and hence, characterized by the first-order conditions. To see the sufficiency of the first-order conditions, note that the first principal minor of Hessian is $\rho V_{dd} - \eta c \sigma_{dd}^2 = -\rho \frac{1}{3q} < 0$ as $\sigma_{dd}^2 = 0$ and that the second principal minor is given by the determinant of the Hessian at the critical point:

$$\begin{aligned} & -\rho V_{dd}(d; \infty) \eta \tilde{c}_{\rho\rho}(\rho) \sigma^2(d; \infty) - (V_d - \eta \tilde{c}_\rho(\rho) \sigma_d^2(d; \infty))^2 \\ &= \rho \frac{\tilde{c}_{\rho\rho}(\rho)}{\tilde{c}_\rho(\rho)} \frac{V(d; \infty)}{3q} - \left(-\frac{d}{3q} + 1 - \frac{V(d; \infty)}{\sigma^2(d; \infty)} \right)^2. \end{aligned} \quad (4)$$

The equality follows from $V_{dd} = -\frac{1}{3q}$, $\sigma^2(d; \infty) = d$, and using the necessary condition (FOC^ρ) via $\eta \sigma^2(d; \infty) = V(d; \infty)/\tilde{c}_\rho(\rho)$. Substituting for $V(d; \infty) = d - d^2/(6q)$ (as $d \leq 3q$ by Step 1) yields as condition for a positive second principal minor:

$$\rho \frac{\tilde{c}_{\rho\rho}(\rho)}{\tilde{c}_\rho(\rho)} > \frac{d}{2(6q - d)}.$$

The inequality holds as the properties of \tilde{c} imply LHS ≥ 1 while RHS $\leq \frac{1}{2}$ for $d \leq 3q$.

Step 3. Characterization. Substituting the expressions for $V(d; \infty)$ and $\sigma^2(d; \infty)$ for expanding knowledge into the first-order condition (FOC^d) yields $\rho(1 - d/(3q)) = \eta \tilde{c}(\rho)$.

Replacing η via equation (FOC ^{ρ}) and solving for d , we obtain

$$d^\eta(\infty) = 3q \left(1 - \frac{\tilde{c}(\rho)}{2\tilde{c}_\rho(\rho)\rho - \tilde{c}(\rho)} \right) \in (2q, 3q),$$

where the bounds follow from the properties of \tilde{c} .

Q.E.D.

Part 2. Deepening Knowledge.

LEMMA 13: *The researcher's optimal choice of distance is $d^\eta(X) = \frac{X}{2}$ for $X \leq \tilde{X}^\eta$ and $d^\eta(X) < X/2$ otherwise. At \tilde{X}^η , the payoff $U_R(X)$ decreases. Further, $\lim_{X \rightarrow \infty} d^\eta(X) = d^\eta(\infty)$ and the convergence is from above. Any optimal distance satisfies $d^\eta(X) \leq 4q$.*

PROOF: Define $d^b := X/2$ —the boundary solution—and d^i as the solution d to (FOC ^{d}) assuming $d < X/2$ (whenever it exists)—the interior solution.

Step 1. d^b always a candidate solution. Note first that the choice d^b always constitutes a local maximum as the marginal cost of distance is zero at this point, $\partial \sigma^2(d, X)/\partial d = 1 - 2d/X$. Moreover, by Lemma 3 also the marginal benefit is zero at $d = X/2$. Finally, for any choice of d , there is a unique ρ that solves (FOC ^{ρ}), because given d , (FOC ^{ρ}) has a continuous, strictly increasing, left-hand side with $\tilde{c}_\rho(0) = 0$ and $\lim_{\rho \rightarrow 1} \tilde{c}_\rho(\rho) = \infty$, and has a constant right-hand side. Hence, the boundary solution with d^b is always a candidate solution.

Step 2. $d^\eta(X) = X/2$ if $X \leq 4q$. Recall (FOC ^{ρ}) and (FOC ^{d}). Assuming an interior solution d^i , replacing η via (FOC ^{ρ}) in (FOC ^{d}) we obtain

$$\frac{V_d(d, X)}{\sigma_d^2(d, X)} \bigg/ \frac{V(d, X)}{\sigma^2(d, X)} = \frac{\tilde{c}(\rho)}{\rho} \bigg/ \tilde{c}_\rho(\rho).$$

It follows from the properties of $\tilde{c}(\rho)$ that the RHS $\in [0, 1/2]$ and decreasing. Thus, if the LHS $> 1/2$, it is beneficial to increase d if possible and the boundary choice d^b is optimal. For short areas, $X \leq 4q$, the boundary choice is optimal as

$$\frac{V_d}{\sigma_d^2} \bigg/ \frac{V}{\sigma^2} = \frac{2(X - 2d)}{X - 2d} \bigg/ \frac{2(dX - d^2)}{d(X - d)} = 1.$$

Step 3. $d^\eta(X) < X/2$ if $X > 8q$. Note first that the variance of the boundary question is always greater than for any interior question. Hence, if the benefits of a discovery, V , are larger for an interior question than for the boundary question, the researcher can obtain a higher payoff by choosing an interior question with the same ρ as for the boundary question. The benefits of discovery on the boundary of an area with $X > 8q$ are always smaller than for some interior distance by Lemma 4. Hence, an interior choice is optimal for $X > 8q$.

Step 4. If d^i is optimal, it must be that $d^i < 4q$ and that $X - d^i > 4q$.

For $X \in (4q, 8q)$ and $X - d < 4q$,

$$\frac{V_d(d, X)}{\sigma_d^2(d, X)} \bigg/ \frac{V(d, X)}{\sigma^2(d, X)} = \frac{2d(X - d)}{-2d^2 + 2dX - \sqrt{X}(X - 4q)^{3/2}},$$

which decreases in d with $\lim_{d \rightarrow X/2} X^2/2/(X^2/2 - \sqrt{X}(X - 4q)^{3/2})$, which, in turn, increases in X and is one for $X = 4q$. Hence, any interior solution must be such that $X - d > 4q$ for the same reasons given in Step 2 of this proof. For $X - d < 4q$, $\partial u_R(\rho, d; X)/\partial$ is always positive.

For $X > 8q$, $X - d^i > 4q$. $d^i < 4q$ follows as $V(d; X)$ decreases in d when $d > 4q$ by Lemma 4.

Summary Step 1–4. We know that (i) in areas with $X < 4q$, the researcher's distance choice on the deepening area will be d^b , (ii) in areas with $X > 8q$ the researcher's distance choice will be d^i , (iii) in areas with $X \in [4q, 8q]$ the researcher's distance choice may be d^i or d^b , but (iv) if the solution is d^i , it has to satisfy $X - d^i > 4q$ and $d^i < 4q$. The latter two imply $d^i < X/2$ in this case.

Step 5. Single Crossing of the Payoffs. With three observations, we show that the payoffs, $U_R(d^b; X)$ and $U_R(d^i; X)$, cross once assuming $\rho(d, X)$ is chosen optimally.

1. At area length X for which $U_R(d^b; X) = U_R(d^i; X)$, the payoff at the boundary solution must be decreasing faster than at the interior solution.
2. On the interval $[4q, 8q]$, the payoff of the boundary solution has a strictly lower second derivative with respect to X for all X than that of the interior solution. Hence, the two values can cross at most once on this interval.
3. $U_R(d^b; X) \leq U_R(d^i; X)$ if $X \geq 8q$.

The first observation follows because the first switch is from the boundary to the interior solution, by continuous differentiability of all terms and $d^\eta(X) = X/2$ for $X < 4q$. The third observation follows from Step 3 above.

The second observation follows from totally differentiating U_R for both local maxima. Using envelope conditions, we obtain that the payoff is concave in the boundary solution and convex in the interior solution implying the second observation. Define $\varphi(X) := \max_\rho u(d = X/2, \rho, X)$ for the boundary; we show in Lemma 31 of [Carnehl and Schneider \(2025\)](#) that $\varphi(X)$ is concave. In Lemma 32 of [Carnehl and Schneider \(2025\)](#), we show that $U_R(X) = \max_{\rho, d} u(d, \rho, X)$ is convex in X if the maximizer satisfies $d^\eta(X) < X/2$.

Step 6. Asymptotics. As $X \rightarrow \infty$, $V(d, X)$ converges to $V(d, \infty)$ and $\sigma^2(d, X)$ to $\sigma^2(d, \infty)$ and the researcher's optimization on the deepening interval converges to that on the expanding interval which has a unique maximum at $(d^\eta(\infty), \rho^\eta(\infty))$. In particular, if such an optimum exists, the envelope condition implies that $dU_R(d^i(X); X)/dX = \rho V_X(d^i, X) - \eta \tilde{c}(\rho) \sigma_X^2(d^i, X) < 0$ as $V_X(d, X) < 0$ according to Lemma 29 for $X > 4q$ and $X - d > 4q$ and $\sigma_X^2(d, X) > 0$. Hence, the payoff of any optimal interior choice decreases in X . Q.E.D.

A.4. Proof of Proposition 4

We prove the statements in Proposition 4 in reverse order. A side-product of this proof is that we show: $4q < \check{X}^\eta \leq \check{X}^0$, $\check{X} \approx 4.548q$, and $\hat{X}^\eta < \hat{X}^0$.

Step 1: Proof of Item 3. We use a series of lemmata to show that a local maximum, \check{X}^η , exists (Lemmata 14 and 15) and that it is global (Lemma 16).

LEMMA 14: Fix $d = X/2$ and assume that an interior optimum exists. Then $U_R(X|d = X/2)$ is maximal only if the total differential $dV(d = X/2; X)/dX \geq 0$.

PROOF: Under the assumption that $d = X/2$, $U_R(X)$ is defined and continuously differentiable for all $X \in [0, \infty)$ despite the indicator functions.¹⁷ Because $X = 0$ implies $U_R(X = 0) = 0$, because $U_R(X)$ declines for X large enough and because Lemma 11 holds, there is an interior X at which $U_R(X)$ is maximized.

As $U_R(X)$ is differentiable and maximized at some interior X it satisfies $\partial U_R / \partial X = 0$. By assumption, $d(\tilde{X}^\eta) = X/2$ and (FOC ^{ρ}) holds. Thus, $\rho \frac{dV(d=X/2; X)}{dX} = \frac{\eta}{4} \tilde{c}(\rho)$. The right-hand side is nonnegative, implying the desired result.¹⁸ *Q.E.D.*

LEMMA 15: $U_R(X; d^b = X/2)$ peaks at $X = \tilde{X}^\eta \in (4q, \tilde{X}^0]$.

PROOF: Define $\hat{U}_R(X) = U_R(X; d^b = X/2)$. Note that $\hat{U}'_R(X) > 0$ for $X \in [0, 4q]$. This follows because in this case $\hat{U}_R(X) = \rho X^2 / (12q) - \eta \tilde{c}(\rho) X / 4$, and hence, $\hat{U}'_R(X) = \rho X / (6q) - \eta \tilde{c}(\rho) / 4$. Using optimality of ρ via the (FOC ^{ρ}), we obtain $X / (6q) = \eta \tilde{c}_\rho(\rho) / 2$, which yields $\hat{U}'_R(X) = \tilde{c}_\rho(\rho) \rho \eta / 4 (2 - \tilde{c}(\rho) / (\rho \tilde{c}_\rho(\rho))) > 0$, where the inequality follows again from the properties of $\tilde{c}(\rho)$.

Moreover, $\hat{U}_R(X)$ is strictly concave on $[4q, 8q]$ as $\hat{V}(X) := V(d = X/2, X)$ is concave by Lemma 30 and $\text{dd}\sigma^2(d = X/2, X) / (dX dX) = 0$ implying $\hat{U}''_R(X) = \rho \hat{V}_{XX} < 0$.¹⁹ For $X > \tilde{X}^0$, $dV(d = X/2; X) / dX < 0$ by the definition of \tilde{X}^0 implying that for $X > \tilde{X}^0$ the researcher's value decreases. By Lemma 14, it follows that the value-maximizing area length $\tilde{X}^\eta \in (4q, \tilde{X}^0]$. *Q.E.D.*

LEMMA 16: $U_R(X)$ is single peaked in X with the maximum attained at \tilde{X}^η .

PROOF: The result follows from three observations: First, $\tilde{X}^\eta > \tilde{X}^0 > 4q$ by Lemmata 13 and 15. Second, $V_X(d; X) < 0$ if $X > 4q$ and $d < X/2$ by Lemma 5. Third, by the envelope theorem, if $d^\eta(X) < X/2$, then $\partial U_R(X) / \partial X = \rho^\eta(X) V_X(d^\eta(X); X) - \eta \tilde{c}(\rho^\eta(X)) \sigma_X(d^\eta(X); X) < \rho^\eta(X) V_X(d^\eta(X); X)$. Thus, the interior-solution payoff intersects the boundary-solution payoff from below while both decrease. *Q.E.D.*

Step 2. Proof of Item 2. Step 2.1 Maximum of $d^\eta(X)$ at \tilde{X}^η . By Lemma 13, $d^\eta(X)$ increases for $X < \tilde{X}^\eta$. By Step 4 in the proof of Lemma 13, we know that any interior solution d^i is such that $d^i < 4q < X - d^i$, and thus, strictly smaller than $X/2$. Thus, $d^\eta(X)$ decreases when it switches from the boundary to an interior solution.

Step 2.2 Maximum of $\rho^\eta(X)$ at \tilde{X} . We guess (and verify in Step 4) that a maximum of $\rho^\eta(X)$ exists in the range $[\hat{X}^\eta, \tilde{X}^\eta]$, that is, the region in which it is optimal to deepen knowledge and to select $d = X/2$.

LEMMA 17: Suppose $d = X/2$ is optimal for a range $[\underline{X}, \overline{X}]$ such that $d^\eta(X) = X/2$. Then the optimal $\rho^\eta(X)$ is single peaked. It is highest at $\hat{X} = 8 \cos(\pi/18) / \sqrt{3}$.

PROOF: By Lemma 14, we know that $dV(d = X/2; X) / dX \geq 0$ and, by Lemma 15, $\overline{X} > \hat{X}^0$. Moreover, recall $\sigma^2(d = X/2; X) = X/4$. The first-order condition with respect

¹⁷Note that the terms appearing in the indicator functions are of the form $\sqrt{a}(a - 4q)^{3/2}$. Taking the limit of their derivative from above to $4q$ yields zero such that the left and right derivative coincide at the point at which the indicator functions become active.

¹⁸The RHS is only 0 if $\eta = 0$, $\rho^\eta(X) = 1$ and $U_R(X) = V(X)$.

¹⁹Where $\rho'(x) = 0$ by optimality and the property of the first-order condition.

to ρ becomes $V(X/2; X)/X = \eta \tilde{c}_\rho(\rho)/4$, with $V(X/2; X)/X = X/(12q) - \mathbf{1}_{X>4q}(X - 4q)^{3/2}/(\sqrt{X}6q)$.

The latter is continuous and concave. Since $\tilde{c}(\rho)$ is an increasing, twice continuously differentiable and convex function, ρ increases in X if and only if $V(X/2; X)/X$ increases in X . By concavity of $V(X/2; X)/X$ that implies single peakedness.

Thus, \hat{X} is independent of η and given by $\hat{X} = 8 \cos(\pi/18)/\sqrt{3} \approx 4.548q$. Q.E.D.

Step 3. Proof of Item 1.

LEMMA 18: \hat{X}^η exists, $\lim_{X \searrow \hat{X}^\eta} \rho^\eta(X) > \rho^\eta(\infty)$, and \hat{X}^η decreases in η .

PROOF: As $X \rightarrow 0$, $d^\eta(X) \rightarrow 0$, and thus, $U_R(X) \rightarrow 0$. By Lemma 12, $U_R(\infty) > 0$. Thus, by continuity of $U_R(X)$, $\exists \hat{X}^\eta > 0$ such that expanding research dominates deepening research for all $X < \hat{X}^\eta$. Cost are increasing in X and by Proposition 2, $V(d; X \in (\hat{X}^0, \infty)) > V(d; \infty)$, which implies $U_R(X \in (\hat{X}^0, \infty)) > U_R(\infty)$. By Lemma 16 and again continuity of $U_R(X)$, that payoff is maximal at \hat{X} . Thus, we obtain that \hat{X}^η exists and that $\hat{X}^\eta < \hat{X}^\eta$.

We now show that $\lim_{X \searrow \hat{X}^\eta} \rho^\eta(X) > \rho^\eta(\infty)$ holds if $\hat{X}^\eta < 6q$, then we show that \hat{X}^η decreases in η which, together with the observation that $\hat{X}^0 < 6q$, proves the lemma. At \hat{X}^η , we have $U_R(\hat{X}^\eta) = U_R(\infty)$:

$$\begin{aligned} \rho(\hat{X}^\eta)V(\hat{X}^\eta/2; \hat{X}^\eta) - \eta \tilde{c}(\rho(\hat{X}^\eta))\hat{X}^\eta/4 \\ = \rho^\eta(\infty)V(d^\eta(\infty); \infty) - \eta \tilde{c}(\rho^\eta(\infty))d^\eta(\infty), \end{aligned} \quad (5)$$

where the fact that $d(\hat{X}^\eta) = \hat{X}^\eta/2$ follows from Lemmata 13, 15, and 16. Moreover, the following holds by optimality:

$$V(d^\eta(\infty); \infty) = \eta \tilde{c}_\rho(\rho^\eta(\infty))d^\eta(\infty), \quad (\text{FOC } \rho^\eta(\infty))$$

$$V(\hat{X}^\eta/2; \hat{X}^\eta) = \eta \tilde{c}_\rho(\rho(\hat{X}^\eta))\frac{\hat{X}^\eta}{4}. \quad (\text{FOC } \rho(\hat{X}^\eta))$$

CLAIM 1: $\rho^\eta(\infty) < \rho(\hat{X}^\eta)$ if $\hat{X}^\eta < 6q$.

Using (FOC $\rho^\eta(\infty)$) and (FOC $\rho(\hat{X}^\eta)$), we obtain from the properties of the error function $\rho(\hat{X}^\eta) > \rho^\eta(\infty)$ if and only if

$$4 \frac{V(\hat{X}^\eta/2; \hat{X}^\eta/2)}{\hat{X}^\eta} > \frac{V(d^\eta(\infty); \infty)}{d^\eta(\infty)}.$$

Case 1: $\hat{X}^\eta > 4q$. Substituting for the $V(\cdot)$'s, the above becomes $d^\eta(\infty) + 2\hat{X}^\eta - 4(\hat{X}^\eta - 4q)^{3/2}/\sqrt{\hat{X}^\eta} > 6q$. A sufficient condition for this to hold is $d^\eta(\infty) - 2\hat{X}^\eta + 10q > 0$. Using that $d^\eta(\infty) > 2q$ by Lemma 12, we obtain that a sufficient condition for $\rho(\hat{X}^\eta) > \rho^\eta(\infty)$ is $\hat{X}^\eta < 6q$.

Case 2: $\hat{X}^\eta \in (2q, 4q]$. Performing the same steps assuming that $\hat{X}^\eta \in [2q, 4q]$, the claim holds if and only if $\hat{X}^\eta/(3q) > 1 - d^\eta(\infty)/(6q) \Leftrightarrow 2\hat{X}^\eta > 6q - d^\eta(\infty) > 4q$ implying the result.

Case 3: $\hat{X}^\eta < 2q$. We show that Case 3 never occurs, that is, $\hat{X}^\eta > 2q$. To do so, we compare $U_R(d = 2q; \infty)$ with $U_R(d = 1q; X = 2q)$ and show that the former is always larger. Hence, $X = 2q < \hat{X}^\eta$ for any η . For $X = d = 2q$, we have $X/(3q) = 1 - d/(6q)$, and thus, $\rho(X = 2q) = \rho(d; \infty) = \rho$ (cf. Case 2). Moreover, we have $V(1q; 2q) = q/3$ and $V(2q; \infty) = 4q/3$. (FOC ρ^X) implies $4V(1q; 2q)/2q = 2/3 = \eta\tilde{c}_\rho(\rho)$. Since $\tilde{c}_\rho(\rho) > \tilde{c}(\rho)/\rho$ for any $\rho > 0$, these imply $\eta\tilde{c}(\rho)/\rho < 2/3$. Note that $U_R(d = 2q; \infty) - U_R(X = 2q) = q(\rho - 3/2\eta\tilde{c}(\rho))$, which is positive as $\eta\tilde{c}(\rho)/\rho < 2/3$. Thus, $U_R(d = 2q; \infty) > U_R(X = 2q)$ and, therefore, $\hat{X}^\eta > 2q$.

CLAIM 2: If $\rho^\eta(\infty) < \rho(\hat{X}^\eta)$, then \hat{X}^η decreases in η .

Using (FOC $\rho^\eta(\infty)$) and (FOC $\rho(\hat{X}^\eta)$) to replace $V(\cdot)$ s in equation (5), we obtain $d^\eta(\infty)(\rho^\eta(\infty)\tilde{c}_\rho(\rho^\eta(\infty)) - \tilde{c}(\rho^\eta(\infty))) = \tilde{X}^\eta/4(\rho(\hat{X}^\eta)\tilde{c}_\rho(\rho(\hat{X}^\eta)) - \tilde{c}(\rho(\hat{X}^\eta)))$, yielding

$$\hat{X}^\eta/4 = d^\eta(\infty) \frac{(\rho^\eta(\infty)\tilde{c}_\rho(\rho^\eta(\infty)) - \tilde{c}(\rho^\eta(\infty)))}{(\rho(\hat{X}^\eta)\tilde{c}_\rho(\rho(\hat{X}^\eta)) - \tilde{c}(\rho(\hat{X}^\eta)))}.$$

The envelope theorem gives

$$\frac{\partial}{\partial \eta}(U_R(\hat{X}^\eta) - U_R(\infty)) = -\tilde{c}(\rho(\hat{X}^\eta))\frac{\hat{X}^\eta}{4} + \tilde{c}(\rho^\eta(\infty))d^\eta(\infty).$$

Replacing for \hat{X}^η implies that the RHS is positive if and only if

$$(\tilde{c}(\rho^\eta(\infty))) - \tilde{c}(\rho(\hat{X}^\eta)) \frac{\rho^\eta(\infty)\tilde{c}_\rho(\rho^\eta(\infty)) - \tilde{c}(\rho^\eta(\infty))}{\rho(\hat{X}^\eta)\tilde{c}_\rho(\rho(\hat{X}^\eta)) - \tilde{c}(\rho(\hat{X}^\eta))} > 0.$$

Using that $\rho\tilde{c}_\rho(\rho) > \tilde{c}(\rho)$ by the properties of the inverse error function and factoring out the denominator of the first term, the above holds if and only if

$$\begin{aligned} & \tilde{c}(\rho^\eta(\infty))\rho^\eta(\infty)\tilde{c}_\rho(\rho(\hat{X}^\eta)) - \tilde{c}(\rho(\hat{X}^\eta))\rho^\eta(\infty)\tilde{c}_\rho(\rho^\eta(\infty)) > 0 \\ \Leftrightarrow & \frac{\rho(\hat{X}^\eta)\tilde{c}_\rho(\rho(\hat{X}^\eta))}{\tilde{c}(\rho(\hat{X}^\eta))} > \frac{\rho^\eta(\infty)\tilde{c}_\rho(\rho^\eta(\infty))}{\tilde{c}(\rho^\eta(\infty))}, \end{aligned}$$

which holds if and only if $\rho(\hat{X}^\eta) > \rho^\eta(\infty)$ by the properties of the error function. Thus, \hat{X}^η decreases if $\rho(\hat{X}^\eta) > \rho^\eta(\infty)$.

Conclusion: Since $\hat{X}^0 \in [2q, 6q]$, $\rho^\eta(\infty) < \rho(\hat{X}^\eta)$ implying that \hat{X}^η decreases in η . Q.E.D.

Step 4.

LEMMA 19: $\hat{X}^\eta < \dot{X}^\eta < \check{X}^\eta < \tilde{X}^\eta$.

PROOF: We first show that $\check{X}^\eta > \dot{X}$. By the envelope theorem, we need at $X = \check{X}^\eta$ that $\rho dV(d = \check{X}^\eta/2; \check{X}^\eta)/dX = \eta \tilde{c}(\rho)/4$. The FOC for ρ at $X = \check{X}^\eta$ gives $V/\check{X}^\eta = \eta \tilde{c}_\rho(\rho)/4$. Assume for a contradiction that $\rho(\check{X}^\eta)$ increases. Then V/\check{X}^η must be increasing, which holds if and only if $dV(d = \check{X}^\eta/2; \check{X}^\eta)/dX > V(d = \check{X}^\eta/2; \check{X}^\eta)/\check{X}^\eta$. Combining this inequality with the above properties leads to a contradiction as $\tilde{c}_\rho(\rho) > \tilde{c}(\rho)/\rho$ by the properties of $\tilde{c}(\rho)$.

Ordering. By Lemma 18, we know that $\hat{X}^\eta < \hat{X}^0$. Thus, because $\hat{X}^0 < \dot{X} \Rightarrow \hat{X}^\eta < \dot{X}$. Moreover, $\check{X}^\eta > \check{X}^\eta$ by Lemma 13, which concludes the proof. Q.E.D.

A.5. Proof of Proposition 5

By assumption, knowledge is dense initially and the base step is satisfied. We show the induction step assuming R_t chooses $d^\eta(\infty)$. We have to show that $\hat{X}^\eta > d^\eta(\infty)$ to show that knowledge is dense in $t+1$. Suppose the opposite holds, then from (FOC ^{ρ}) we know $(6q - d^\eta(\infty))/(6q) = V(d^\eta(\infty); \infty)/\sigma^2(d^\eta(\infty); \infty) = \eta \tilde{c}_\rho(\rho^\eta(\infty))$ and $d^\eta(\infty)/(6q) = V(d^\eta(\infty)/2; d^\eta(\infty))/\sigma^2(d^\eta(\infty)/2; d^\eta(\infty)) = \eta \tilde{c}_\rho(\rho(d^\eta(\infty)))$ implying

$$\rho^\eta(\infty) = \operatorname{erf} \left(\sqrt{\frac{W \left(\frac{36q^2 - 12qd^\eta(\infty) + (d^\eta(\infty))^2}{18q^2\eta^2\pi} \right)}{2}} \right) \quad \text{and}$$

$$\rho(d^\eta(\infty)) = \operatorname{erf} \left(\sqrt{\frac{W \left(\frac{(d^\eta(\infty))^2}{18q^2\eta^2\pi} \right)}{2}} \right),$$

where $W(\cdot)$ is the Lambert W function. Because $d^\eta(\infty)$ is linear in q by Lemma 33 in [Carnehl and Schneider \(2025\)](#); and $d^\eta(\infty) < 3q$ by Lemma 12, it follows that $36q^2 - 12qd^\eta(\infty) > 0$, which implies that $\rho^\eta(\infty) > \rho(d^\eta(\infty))$ by the monotonicity of the Lambert W function.

By Lemma 17, we know that $\rho^\eta(X)$ increases for $X < \dot{X} = 8\cos(\pi/18)/\sqrt{3}$. By Lemma 18, we know $\rho(\hat{X}^\eta) > \rho^\eta(\infty)$. By Lemma 19, $\hat{X}^\eta < \dot{X}$. Thus, $d^\eta(\infty) < \hat{X}^\eta$, which implies that if R_t expands knowledge, so does R_{t+1} .

The proof of the comparative statics result follows directly from the implicit function theorem and the first-order conditions; see the Supplemental Appendix ([Carnehl and Schneider \(2025\)](#)) for details.

A.6. Proof of Proposition 6

We begin with the negative benchmark results for $\eta \rightarrow \infty$ and $\eta = 0$. If $\eta \rightarrow \infty$, research becomes infinitely costly. Hence, $\rho \rightarrow 0$ and $u_R(d^\eta(X); \rho^\eta(X); X) \rightarrow 0$ for any X . Thus, absent interventions, research creates no value. Any disclosure by the designer should maximize the immediate payoff $V(d; \infty)$.

If $\eta = 0$, $\rho^0(\cdot) = 1$ and $u_R(d, \rho^0; X) = V(d; X)$. Thus, each researcher maximizes $V(d; X)$. By construction, maximizing the per-period $V(d; X)$ corresponds to maximizing the long-run objective of the decision maker.

For intermediate ranges, it suffices to show that selecting a moonshot of length $6q$ is preferred to selecting the myopically optimal $d = 3q$ for some $(\underline{\eta}, \bar{\eta})$ and $\delta(\eta) < 1$.

We restrict attention to η -levels such that $d^\eta(6q) = 3q$. These levels exist due to $\tilde{X}^0 > 6q$ by Lemma 10. A moonshot is preferred if and only if

$$\begin{aligned} & \frac{V(6q; \infty) + \delta \rho^\eta(6q)(V(3q; 6q) + \delta CV)}{1 - \delta} \\ & \geq \frac{V(3q; \infty) + \delta \rho^\eta(\infty)(V(d^\eta(\infty); \infty) + \delta CV)}{1 - \delta}, \end{aligned}$$

where CV is the (common) continuation value conditional on researchers not being stuck at $t = 3$. Because $\rho^\eta(6q) > \rho^\eta(\infty)$ by Proposition 3 a sufficient condition for the above is that

$$V(6q; \infty) + \delta \rho^\eta(6q)V(3q; 6q) \geq V(3q; \infty) + \delta \rho^\eta(\infty)V(d^\eta(\infty); \infty), \quad (6)$$

which we will now show for some η .

For the moonshot $d = 6q$, we obtain as value in $t = 1$, $V(6q; \infty) = 2/\sqrt{3}q$, and in $t = 2$, $V(3q; 6q) = (3 - 2/\sqrt{3})q$. The success probability in $t = 2$ follows from Proposition 3 for deepening research on a research area with $X = 6q$ while being on the boundary distance $d(X = 6q) = X/2$.

For the myopic optimum $d = 3q$, we obtain as value in $t = 1$, $V(3q; \infty) = 3/2q$, and in $t = 2$, $V(d^\eta(\infty); \infty) = d^\eta(\infty)(1 - d^\eta(\infty)/(6q))$. The distance and success probability in $t = 2$ follow from Proposition 3 for expanding knowledge.

Plugging into (7),

$$2/\sqrt{3}q + \delta \rho^\eta(6q)(3 - 2/\sqrt{3})q \geq 3/2q + \delta \rho^\eta(\infty)d^\eta(\infty)(1 - d^\eta(\infty)/(6q)).$$

By continuity in η and δ , it suffices to show that, for $\delta = 1$ and some $\eta > 0$ this inequality is strict. Numerically solving for $d^\eta(\infty)$, $\rho^\eta(\infty)$, $\rho^\eta(6q)$ using $(\eta = 0.25, q = 1)$ verifies strict inequality.²⁰ As $d^\eta(\infty)$ is linear in q (see Lemma 33 in the Supplemental Appendix, [Carnehl and Schneider \(2025\)](#)), $\rho^\eta(\infty)$ is constant in q . Linearity of distance and invariance of probability in the moonshot are immediate. Thus, restricting attention to $q = 1$ is without loss.

A.7. Proof of Proposition 7

PROOF: We first state a relationship between different types of knowledge for the dynamic setting. We begin by defining a (forward-looking) notion of equivalence.

DEFINITION 3—Forward Equivalence of Knowledge: Knowledge \mathcal{F}_t and \mathcal{F}'_t are considered forward equivalent at time t , if—absent a disclosure opportunity at time t —the expected future generated payoff is the same.

²⁰In this case, $\rho^\eta(6q) = 0.8457$, $\rho^\eta(\infty) = 0.7174$, $d^\eta(\infty) = 2.3988$.

Note that the current value of \mathcal{F}_t and \mathcal{F}'_t need not be the same, but are inconsequential for the designer's decisions moving on.

LEMMA 20: *Knowledge \mathcal{F} is forward equivalent to \mathcal{F}_0 if it is dense, and there have been no failures in the past.*

PROOF: If \mathcal{F} is dense, no researcher ever deepens knowledge inside areas existing under \mathcal{F} because $X < \hat{X}^\eta$. No designer would disclose a question to deepen knowledge, because $X < \hat{X}^0$. Hence, behavior and the generated payoffs are the same. *Q.E.D.*

LEMMA 21: *Suppose $\mathcal{F}_t = \{(0, y(0)), (x_1, y(x_1))\}$ with $x_1 \in [4q, \min\{\tilde{X}^\eta, 2\hat{X}^\eta\}]$ and that there was a failure at $x_1/2$. Further, assume the designer can disclose one answer at $t + 1$ for free and expects no future disclosure. Then, if (η, δ) is such that $\delta\rho^\eta(\infty) > 1/2$, the designer strictly prefers to disclose the answer to $x_1/2$ than to $2x_1$.*

PROOF: If the designer discloses the answer to $x_1/2$, her continuation payoff is

$$V(x_1/2; x_1) + \sum_{k=1}^{\infty} (\delta\rho^\eta(\infty))^k V(d^\eta(\infty); \infty)$$

Alternatively, the designer could disclose the answer to question $2x_1$ thereby starting a new research cycle. Her expected payoff from that action is $V(x_1; \infty) + \delta\rho^\eta(x_1)V(x_1/2; x_1)$ because after the new research cycle is completed, researchers revert back to trying to find the answer to question $x_1/2$ and fail due to symmetry.

The designer strictly prefers the disclosure of $x_1/2$ if

$$\begin{aligned} V\left(\frac{x_1}{2}; x_1\right) + \sum_{k=1}^{\infty} (\delta\rho^\eta(\infty))^k V(d^\eta(\infty); \infty) &> V(x_1; \infty) + \delta\rho^\eta\left(\frac{x_1}{2}; x_1\right) V\left(\frac{x_1}{2}; x_1\right) \\ \Leftrightarrow \sum_{k=1}^{\infty} (\delta\rho^\eta(\infty))^k &> \frac{V(x_1; \infty) - (1 - \delta\rho^\eta(x_1/2; x_1))V(x_1/2; x_1)}{V(d^\eta(\infty); \infty)}. \end{aligned} \quad (7)$$

Notice that $V(x_1; \infty)/V(d^\eta(\infty); \infty) < V(x_1; \infty)/V(2q; \infty) \leq 1$ because $x_1 \geq 4q$. As $\delta\rho^\eta(x_1/2; x_1) < 1$ and $V(x_1/2; x_1) > 0$, an upper bound for the RHS is 1. Inequality (7) holds if $\delta\rho^\eta(\infty)/(1 - \delta\rho^\eta(\infty)) \geq 1$, which is true if $\rho^\eta(\infty)\delta > 1/2$. *Q.E.D.*

Combining Lemmata 20 and 21 with the fact that payoffs converge to the “no future disclosure opportunity” payoffs as $\lambda \rightarrow 0$ yields the proposition.²¹ *Q.E.D.*

REFERENCES

- AGHION, PHILIPPE, AND PETER HOWITT (1992): “A Model of Growth Through Creative Destruction,” *Econometrica*, 60 (2), 323–351. [0628]
 AGHION, PHILIPPE, MATHIAS DEWATRIPONT, AND JEREMY C. STEIN (2008): “Academic Freedom, Private-Sector Focus, and the Process of Innovation,” *The Rand Journal of Economics*, 39 (3), 617–635. [0627]

²¹Use the parameters ($\eta = 0.25$, $q = 1$) from the proof of Proposition 6, assume $\delta \geq \rho^\eta(\infty)$ so the environment is promising, and observe that for $\delta = \rho^\eta(\infty)$ a moonshot of around ≈ 5.3 is optimal. Local continuity implies a non-knife-edge case.

- AKERLOF, GEORGE A., AND PASCAL MICHAILLAT (2018): "Persistence of False Paradigms in Low-Power Sciences," *Proceedings of the National Academy of Sciences of the United States of America*, 115 (52), 13228–13233. [0626,0643]
- ALEXANDER, JASON MCKENZIE, JOHANNES HIMMELREICH, AND CHRISTOPHER THOMPSON (2015): "Epistemic Landscapes, Optimal Search, and the Division of Cognitive Labor," *Philosophy of Science*, 82 (3), 424–453. [0626]
- ANDREWS, ISAIAH, AND MAXIMILIAN KASY (2019): "Identification of and Correction for Publication Bias," *American Economic Review*, 109 (8), 2766–2794. [0643]
- BARDHI, ARJADA (2024): "Attributes: Selective Learning and Influence," *Econometrica*, 92 (2), 311–353. [0627]
- BARDHI, ARJADA, AND NINA BOBKOVA (2023): "Local Evidence and Diversity in Minipublics," *Journal of Political Economy* 131 (9), 2451–2508. [0627]
- BOBTCHIEFF, CATHERINE, JÉRÔME BOLTE, AND THOMAS MARIOTTI (2017): "Researcher's Dilemma," *The Review of Economic Studies*, 84 (3), 969–1014. [0643]
- BRAMOULLÉ, YANN, AND GILLES SAINT-PAUL (2010): "Research Cycles," *Journal of Economic Theory*, 145 (5), 1890–1920. [0627]
- BROCK, WILLIAM A., AND STEVEN N. DURLAUF (1999): "A Formal Model of Theory Choice in Science," *Economic Theory*, 14, 113–130. [0626]
- BRYAN, KEVIN A., AND JORGE LEMUS (2017): "The Direction of Innovation," *Journal of Economic Theory*, 172, 247–272. [0627,0644]
- BUSH, VANNEVAR (1945): "Science—the Endless Frontier," Report to the President—Letter of Transmittal. URL: <https://www.nsf.gov/od/lpa/nsf50/vbush1945.htm>. [0623]
- CALLANDER, STEVEN (2011): "Searching and Learning by Trial and Error," *American Economic Review*, 101 (6), 2277–2308. [0627,0628]
- CALLANDER, STEVEN, AND TOM S. CLARK (2017): "Precedent and Doctrine in a Complicated World," *American Political Science Review*, 111 (1), 184–203. [0627]
- CALLANDER, STEVEN, AND PATRICK HUMMEL (2014): "Preemptive Policy Experimentation," *Econometrica*, 82 (4), 1509–1528. [0627]
- CALLANDER, STEVEN, AND NIKO MATOUSCHEK (2019): "The Risk of Failure: Trial and Error Learning and Long-Run Performance," *American Economic Journal: Microeconomics*, 11 (1), 44–78. [0627]
- CALLANDER, STEVEN, NICOLAS LAMBERT, AND NIKO MATOUSCHEK (2021): "The Power of Referential Advice," *Journal of Political Economy*, 129 (11), 3073–3140. [0627]
- (2022): "Innovation and Competition on a Rugged Technological Landscape," Report. [0627]
- CARNEHL, CHRISTOPH, AND JOHANNES SCHNEIDER (2025): "Supplement to 'A Quest for Knowledge'," *Econometrica Supplemental Material*, 93, <https://doi.org/10.3982/ECTA22144>. [0634,0642,0644–0647,0651,0655,0656]
- FORTUNATO, SANTO, CARL T. BERGSTROM, KATY BÖRNER, JAMES A. EVANS, DIRK HELBING, STAŠA MILOJEVIĆ, ALEXANDER M. PETERSEN, FILIPPO RADICCHI, ROBERTA SINATRA, BRIAN UZZI, ALESSANDRO VESPIGNANI, LUDO WALTMAN, DASHUN WANG, AND ALBERT-LÁSZLÓ BARABÁSI (2018): "Science of Science," *Science*, 359 (6379). [0626,0643]
- FOSTER, JACOB G., ANDREY RZHETSKY, AND JAMES A. EVANS (2015): "Tradition and Innovation in Scientists' Research Strategies," *American Sociological Review*, 80 (5), 875–908. [0627]
- FRANKEL, ALEXANDER P., JOSHUA L. KRIEGER, DANIELLE LI, AND DIMITRIS PAPANIKOLAOU (2024): "Evaluation and Learning in R&D Investment," NBER Working Paper 31290. [0643]
- GARFAGNINI, UMBERTO, AND BRUNO STRULOVICI (2016): "Social Experimentation With Interdependent and Expanding Technologies," *The Review of Economic Studies*, 83 (4), 1579–1613. [0627]
- GROSSMAN, GENE M., AND ELHANAN HELPMAN (1991): "Quality Ladders in the Theory of Growth," *The Review of Economic Studies*, 58 (1), 43–61. [0628]
- HECKMAN, JAMES J., AND SIDHARTH MOKTAN (2020): "Publishing and Promotion in Economics: The Tyranny of the Top Five," *Journal of Economic Literature*, 58 (2), 419–470. [0643]
- HILL, RYAN, AND CAROLYN STEIN (2025): "Scooped! Estimating Rewards for Priority in Science," *Journal of Political Economy* 133 (3), 793–845. [0624,0643,0644]
- (2024): "Race to the Bottom: Competition and Quality in Science," Report. [0624,0643,0644]
- HOPENHAYN, HUGO, AND FRANCESCO SQUINTANI (2021): "On the Direction of Innovation," *Journal of Political Economy*, 129 (7), 1991–2022. [0627,0644]
- JAFFE, ADAM B., MICHAEL S. FOGARTY, AND BRUCE A. BANKS (2003): "Evidence From Patents and Patent Citations on the Impact of NASA and Other Federal Labs on Commercial Innovation," *Journal of Industrial Economics*, 46 (2), 183–205. [0643]
- JOVANOVIĆ, BOYAN, AND RAFAEL ROB (1990): "Long Waves and Short Waves: Growth Through Intensive and Extensive Search," *Econometrica*, 1391–1409. [0627]

- KIM, SOOMI (2023): "Shortcuts to Innovation: The Use of Analogies in Knowledge Production," Report. [0644]
- KITCHER, PHILIP (1990): "The Division of Cognitive Labor," *The Journal of Philosophy*, 87 (1), 5–22. [0626]
- KLETTE, TOR JAKOB, AND SAMUEL KORTUM (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112 (5), 986–1018. [0628]
- KUHN, THOMAS S. (1962): *The Structure of Scientific Revolutions*. University of Chicago Press. [0626]
- MEDAWAR, PETER B. (1967): *The Art of the Soluble Creativity and Originality in Science*. London: Penguin Books. [0623]
- MERTON, ROBERT K. (1957): "Priorities in Scientific Discovery: A Chapter in the Sociology of Science," *American Sociological Review*, 22 (6), 635–659. [0626]
- MYERS, KYLE (2020): "The Elasticity of Science," *American Economic Journal: Applied Economics*, 12 (4), 103–134. [0627,0644]
- NAGARAJ, ABHISHEK, AND MATTEO TRANCHERO (2024): "How Does Data Access Shape Science? Evidence From the Impact of US Census's Research Data Centers on Economics Research," NBER Working Paper 31372. [0644]
- ROMER, PAUL M. (1990): "Endogenous Technological Change," *Journal of Political Economy*, 98 (5, Part 2), S71–S102. [0628]
- RZHETSKY, ANDREY, JACOB G. FOSTER, IAN T. FOSTER, AND JAMES A. EVANS (2015): "Choosing Experiments to Accelerate Collective Discovery," *Proceedings of the National Academy of Sciences of the United States of America*, 112 (47), 14569–14574. [0627,0643]
- STREVEVS, MICHAEL (2003): "The Role of the Priority Rule in Science," *The Journal of Philosophy*, 100 (2), 55–79. [0626]
- URGUN, CAN, AND LEEAT YARIV (2025): "Contiguous Search: Exploration and Ambition on Uncharted Terrain," *The Journal of Political Economy*, 133 (2), 522–567. [0627]
- WEISBERG, MICHAEL, AND RYAN MULDOON (2009): "Epistemic Landscapes and the Division of Cognitive Labor," *Philosophy of Science*, 76 (2), 225–252. [0626]

Co-editor Marina Halac handled this manuscript.

Manuscript received 15 August, 2023; final version accepted 11 December, 2024; available online 12 December, 2024.