

A Simplified Quest for Knowledge

JOSHUA S. GANS*

Rotman School of Management, University of Toronto and NBER

This paper develops a transparent, simplified version of Carnehl and Schneider (2025)’s model of knowledge creation. Our tractable framework, which yields closed-form solutions for key welfare trade-offs, preserves the essential economic mechanisms while eliminating mathematical complexity. We derive four main insights. First, in the costless two-period benchmark analysed here, the first-best planner never chooses an *expand–then–deepen* cycle in which one pushes the frontier and then returns to deepen the newly created region. This clarifies that the “moonshot” mechanism in Carnehl and Schneider is a *second-best* rationale that operates once research costs (and the associated dynamic externalities) are introduced, rather than a first-best implication of direct welfare comparisons in the costless benchmark. Second, in the same benchmark, private and social incentives coincide on the *extensive margin* of whether to expand or deepen in a given period; any divergence that appears in our extensions concerns the *intensive margin* of where within a long bounded gap deepening occurs. Third, we analyse how citation-based incentive systems affect knowledge creation trajectories. We show that systems that privilege unique contributions over shared ones align private behaviour with social welfare objectives, while those that reward shared contributions lead to excessive knowledge deepening. Fourth, our analysis provides precise characterisations of optimal knowledge creation paths under various initial conditions and offers clear guidance for science policy. By clarifying when interventions can address misalignments between researchers’ incentives and social welfare, our simplified model offers practical insights for the design of research funding mechanisms.

Keywords: knowledge creation, research policy, moonshot, multidisciplinary research, scientific exploration, citation incentives

JEL Classifications: O31, D83, H41

1 Introduction

How do researchers choose which questions to pursue? When is their choice misaligned with social welfare? And how might science policy correct these misalignments? These questions lie at the heart of the economics of science and are crucial for designing effective science policy.

In a recent contribution, Carnehl and Schneider (2025) (henceforth CS) develop a rich model of knowledge creation in which researchers select both the questions they investigate and the

*Rotman School of Management, University of Toronto and NBER. Thanks to Pierre Azoulay, Christoph Carnehl and Johannes Schneider for helpful discussions, two anonymous referees for comments and Claude Sonnet 3.7 for valuable research assistance. Responsibility for all errors remains my own.

© 2026 Joshua S. Gans. Licenced under the Creative Commons Attribution-Noncommercial 4.0 Licence (<http://creativecommons.org/licenses/by-nc/4.0/>). Available at <http://rofea.org>.

intensity of their research efforts. Their framework elegantly captures the trade-off between pursuing novel questions distant from existing knowledge and refining our understanding within established domains. While CS are often cited as providing a rationale for “moonshots”—highly novel research beyond the usual frontier—our analysis reveals a more nuanced picture that challenges simplistic policy narratives.

Our contribution is methodological: we focus on a transparent, simplified version of CS that yields closed-form solutions for key welfare trade-offs. This simplified framework preserves the essential economic mechanisms while eliminating mathematical complexity, making the insights more accessible to economists, science policy researchers, and research funders. By stripping away non-essential complexities, we expose the fundamental welfare implications of different knowledge creation strategies.

Our analysis yields four key insights that both clarify and challenge aspects of the original CS model. First, in the stripped-down costless two-period benchmark, direct welfare comparisons do not generate a first-best case for an *expand–then–deepen* “moonshot” cycle. This point is best read as a benchmark clarification: in Carnehl and Schneider, moonshots are motivated by second-best considerations that arise once one introduces research costs and the resulting intertemporal externalities. In the costless benchmark, the planner prefers either repeated expansion or deepening existing gaps (depending on their length), but does not optimally overshoot by expanding far and then returning to deepen the newly created region.

Second, regarding multidisciplinary research, we identify a genuine but moderate misalignment between private and social incentives. When researchers face large knowledge gaps between disciplines, they systematically choose deepening locations that create suboptimal knowledge structures from a social welfare perspective. Remarkably, this misalignment persists even without research costs. However, our quantitative analysis shows that the optimal policy intervention is relatively subtle; nudging researchers slightly closer to the midpoint between disciplines rather than mandating complete bridges across disciplinary divides. The social planner never chooses equal spacing across disciplines when optimising welfare, challenging naive intuitions about multidisciplinary research policy.

Third, our simplified model provides precise characterisations of optimal knowledge creation paths under various initial conditions. We derive exact thresholds for when researchers and social planners should expand versus deepen knowledge, and show how these choices depend on the initial knowledge gap and social discount factor. This will be of use to subsequent researchers building on the CS approach.

Fourth, we extend our analysis to examine how citation-based incentive systems affect knowledge creation. Scientific recognition typically depends on how researchers’ work is utilised by others, rather than its direct contribution to knowledge. We show that citation systems can be calibrated to align private and social incentives, but require careful design. Systems that strongly privilege unique contributions over shared ones encourage knowledge expansion,

while those that distribute substantial credit for shared contributions lead to excessive deepening. This provides a novel perspective on how scientific reward structures shape research trajectories.

By making the model more accessible and deriving precise welfare results, our work helps ground policy discussions in robust economic analysis that was pioneered by CS's paper. While the discussion here might be seen as critical of the Carnehl and Schneider (2025) paper, it is intended as a constructive clarification of which results follow from first-best welfare comparisons in the costless benchmark and which rely on second-best mechanisms once research costs are introduced.

1.1 Related Literature

Our paper builds directly on Carnehl and Schneider (2025), who develop a model of knowledge creation where the position of research questions on the real line determines both their novelty and the difficulty of answering them. Their analysis shows that myopic researchers tend to select questions that are too narrow and fail too often compared to what would maximise social welfare.

Carnehl and Schneider (2025) belongs to a broader literature on the economics of science and innovation. Models of cumulative innovation (Scotchmer, 1991; Hopenhayn et al., 2006) highlight how early research affects subsequent discoveries. Work on the direction of innovation (Bryan and Lemus, 2017; Hopenhayn and Squintani, 2021) examines how researchers select among potential research paths. Empirical studies (Foster et al., 2015; Myers, 2020) document how researchers balance novelty against the probability of success.

Our simplified model bridges theoretical insights with practical policy implications, contributing to a growing literature on science funding (Azoulay et al., 2019; Hill and Stein, 2024). By clearly isolating the welfare effects of different research trajectories, we provide a framework to evaluate policies aimed at correcting research incentives, complementing recent work on research funding mechanisms (Myers, 2020; Hill and Stein, 2025).

The conceptual foundation of our work relates to literature on knowledge representation and reasoning under uncertainty. The Brownian path model of knowledge, which we adopt from Carnehl and Schneider (2025), builds on work by Callander (2011) that uses Brownian motion to represent the correlation structure of answers to related questions. This approach has been extended to various settings, including political lobbying (Callander and Clark, 2017), search processes (Callander and Hummel, 2014), and innovation landscapes (Callander et al., 2022).

Our analysis of when moonshots might be justified relates to debates about the appropriate balance between high-risk, high-reward research and incremental advances (Azoulay et al., 2011; Hill and Stein, 2024). By providing precise conditions for when and why exploratory research should be incentivised, we contribute to this ongoing debate in both the academic literature and policy discussions.

Our analysis of multidisciplinary research connects to a growing literature on interdisciplinary science and the challenges of bridging distinct knowledge domains (Van Noorden, 2015; Fortunato et al., 2018). This literature highlights the institutional and cognitive barriers to interdisciplinary work, while our model provides an economic rationale for why such barriers persist and how they might be overcome through appropriate funding mechanisms.

Our examination of citation-based incentives builds on a rich literature on scientific reward structures and how they shape knowledge creation. Beginning with Merton (1957)’s seminal work on priority in scientific discovery, scholars have analysed how formal and informal rewards affect researchers’ choices. More recent empirical studies by Fortunato et al. (2018) and Wang et al. (2017) demonstrate that citation patterns significantly influence research directions, while theoretical models (Dasgupta and David, 1994; Aghion et al., 2008) explore how various reward systems affect knowledge production.

Of particular relevance to our citation model is the work on scientific attribution by Gans and Murray (2014), Bikard et al. (2015), and Gans and Murray (2023). Gans and Murray (2014) examines how attribution norms evolve over time and influence the organisation of science itself. Bikard et al. (2015) identifies a fundamental trade-off between productive efficiency in collaboration and credit allocation. Most recently, Gans and Murray (2023) analyse how formal attribution processes influence collaborative decisions between researchers, showing that different attribution regimes create varying incentives for research quality.

Our citation model complements this literature by specifically examining how different citation reward structures affect researchers’ decisions to expand versus deepen knowledge. While prior work focuses on attribution decisions within research teams, our model emphasises how citation rewards shape the broader landscape of knowledge by influencing the types of contributions researchers choose to make. By identifying precise thresholds where citation systems align or misalign with social welfare objectives, we provide guidance for the design of scientific reward structures that encourage optimal knowledge creation trajectories.

The remainder of the paper is organised as follows. Section 2 presents our simplified model of knowledge creation. Section 3 analyses the researcher’s problem. Section 4 characterises the planner’s optimal policy. Section 5 examines why moonshots are not socially desirable in our framework and discusses the conditions under which they might be justified in the original CS model. Section 6 extends the model to multidisciplinary research contexts, demonstrating how private and social incentives misalign even without research costs. Section 7 extends our model to incorporate citation-based reward structures and analyses their impact on knowledge creation trajectories. Section 8 concludes.

2 Model

Our model builds on the framework of CS with simplifications that allow us to focus on the core dynamics of knowledge creation. We consider a two-period setting where researchers choose

questions to investigate, with specific attention to the distinction between expanding knowledge beyond current frontiers and deepening knowledge within established domains.

2.1 Questions, Answers, and Knowledge

Following CS, we represent the universe of questions as the real line. Each question $x \in \mathbb{R}$ has a unique answer $y(x) \in \mathbb{R}$. The truth—the mapping from questions to answers—follows a standard Brownian motion defined over the entire real line. This structure captures the intuition that questions closer to each other are likely to have related answers.

Knowledge \mathcal{F}_k consists of a finite collection of known question-answer pairs denoted by $\{(x_i, y(x_i))\}_{i=1}^k$. We assume that these pairs are ordered such that $x_i < x_{i+1}$ for all i . Knowledge partitions the real line into $k + 1$ intervals: $(-\infty, x_1)$, $[x_1, x_2)$, \dots , $[x_{k-1}, x_k)$, and $[x_k, \infty)$. The length of a bounded interval is denoted by $X_i = x_{i+1} - x_i$.

At the beginning of our two-period model, knowledge consists of two question-answer pairs located at positions $x = 0$ and $x = x_0 > 0$, creating an interval of length $X_0 = x_0$ and two unbounded frontiers. This initial knowledge is denoted $\mathcal{F}_0 = \{(0, y(0)), (x_0, y(x_0))\}$.

Given this knowledge configuration, a researcher can contribute to knowledge in two fundamentally different ways:

- **Expanding knowledge:** The researcher discovers the answer to a question beyond either frontier ($x < 0$ or $x > x_0$), which creates a new bounded interval $[x_0, x]$.
- **Deepening knowledge:** The researcher discovers the answer to a question within the existing bounded interval ($0 < x < x_0$), which splits the original interval into two smaller intervals: $[0, x]$ and $[x, x_0]$.

This key distinction between expanding and deepening knowledge forms the fundamental strategic choice in our model. The choice determines not only the immediate benefits but also future research opportunities.

2.2 Conjectures and Value of Knowledge

Knowledge is valuable because it guides society's decision-making. For any question x , a decision-maker forms a conjecture about the answer based on existing knowledge. Given knowledge \mathcal{F}_k , this conjecture follows a normal distribution with mean $\mu_x(Y|\mathcal{F}_k)$ and variance $\sigma_x^2(Y|\mathcal{F}_k)$.

For our specific setup with knowledge $\mathcal{F}_0 = \{(0, y(0)), (x_0, y(x_0))\}$, the variance for a question $x \in [0, x_0]$ is:

$$\sigma_x^2(Y|\mathcal{F}_0) = \frac{x(x_0 - x)}{x_0} \tag{1}$$

Intuitively, the variance is lowest at the known points (where it equals zero) and highest at the midpoint of the interval (where it equals $\frac{X^2}{4}$). This captures the idea that confidence in our conjectures decreases as we move away from known facts.

More generally, for a question at distance d from the nearest known point in an interval of length X , the variance is:

$$\sigma^2(d; X) = \frac{d(X-d)}{X} \quad \text{for } 0 \leq d \leq X \quad (2)$$

For questions beyond the frontier, the variance simply equals the distance to the frontier: $\sigma^2(d; \infty) = d$. This reflects the increasing uncertainty as we venture into unexplored territory.

For each question x , the decision-maker either applies knowledge or selects an outside option with a normalised payoff of zero. When applying knowledge, the decision-maker's expected payoff is:

$$1 - \frac{\sigma_x^2(Y|\mathcal{F}_k)}{q} \quad (3)$$

where $q > 0$ is an exogenous parameter that governs the importance of precision.

The decision-maker applies knowledge only when this payoff is positive, which occurs when $\sigma_x^2(Y|\mathcal{F}_k) < q$. This creates a natural boundary: questions with variance exceeding q are too uncertain to answer confidently, so the decision-maker prefers the outside option.

The total value of knowledge to society is:

$$v(\mathcal{F}_k) := \int_{-\infty}^{\infty} \max \left\{ 1 - \frac{\sigma_x^2(Y|\mathcal{F}_k)}{q}, 0 \right\} dx \quad (4)$$

This represents the aggregate benefit of applying knowledge across all questions where doing so is preferable to the outside option. The value of knowledge increases as more questions can be addressed with sufficient precision.

2.3 Benefits of Discovery

The benefit of discovering the answer to question x is the marginal increase in the value of knowledge:

$$V(x; \mathcal{F}_k) := v(\mathcal{F}_k \cup \{(x, y(x))\}) - v(\mathcal{F}_k) \quad (5)$$

CS show that this benefit depends only on the distance d from x to the nearest known question and the length X of the interval containing x (or $X = \infty$ for expanding beyond the frontier).

The full expression for the benefits of discovery, as derived by CS, is:

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

where $\sigma^2(d; X) = \frac{d(X-d)}{X}$ and the indicator functions $\mathbf{1}_{condition}$ equal 1 when the condition is true and 0 otherwise.

This complex expression reflects different regimes that arise based on whether variances at different points exceed the threshold q . However, for some common cases, this expression simplifies considerably:

Case 1: Deepening knowledge within a short interval ($X \leq 4q$). With this short interval, as shown by CS, it is optimal for the decision-maker to take an action if a question falls within that interval; that is, at the mid-point, $X/2$, $1 > \sigma^2/q$. In this case, all indicator functions are zero, and the benefit function reduces to:

$$V(d; X) = \frac{1}{6q} \cdot 2X \cdot \frac{d(X-d)}{X} = \frac{d(X-d)}{3q} \quad (7)$$

This function is symmetric and attains its maximum at the midpoint $d = \frac{X}{2}$.

Case 2: Expanding knowledge beyond the frontier ($X = \infty$) when $d \leq 4q$. In this case, only the first term matters, and the benefit function simplifies to:

$$V(d; \infty) = d - \frac{d^2}{6q} \quad (8)$$

This function increases up to $d = 3q$ and then decreases, capturing a fundamental trade-off: more distant discoveries extend knowledge further but create longer intervals where conjectures remain imprecise.

Case 3: Deepening knowledge within a long interval ($X > 4q$). In this more complex case, the indicator functions come into play. CS show that the benefits-maximising distance depends on the length of the interval:

- For $4q < X \leq \tilde{X}_0$ (where $\tilde{X}_0 < 8q$), the midpoint $d = \frac{X}{2}$ maximises benefits.
- For $X > \tilde{X}_0$, the benefits-maximizing distance is strictly between $3q$ and the midpoint, $d_0(X) \in (3q, X/2)$, and approaches $3q$ as X becomes very large.

These formulations embody key insights from CS: benefits are maximised at intermediate distances (balancing novelty against connection to existing knowledge), and deepening knowledge is particularly valuable for longer intervals where existing conjectures are less precise. However, for very long intervals, the optimal deepening location shifts away from the midpoint and toward the “sweet spot” around $3q$ that balances costs and benefits effectively.

2.4 Research Decisions

In each period, a short-lived researcher makes a decision about which question to investigate. The researcher selects a question x to research, which determines whether to expand or deepen knowledge and the distance d from the nearest known question.

CS consider choosing a greater distance from existing knowledge to be a costly endeavour. They assume that choosing a higher d involves costs of its own and also makes it more difficult to successfully complete a research project. Specifically, they assume that the cost associated with a choice of distance, d , and a probability of success, ρ , is given by $\hat{c}(\rho, d; X) = \eta(\text{erf}^{-1}(\rho))^2 \sigma^2(d, X)$, where $\text{erf}^{-1}(\cdot)$ is the inverse error function. This has the property that as you increase d it becomes increasingly costly to achieve project success.¹

Here, however, as choosing a higher d already involves trade-offs and, conceptually at least, any particular project is investigated on its own terms and does not explicitly use past knowledge as an input, these costs are set aside. This allows us to focus on the choice of which knowledge to create. To be sure, there are still resource limitations that make this an economic choice; in particular, there is a limited number of scientists in each period (in our case, there will be a single scientist). This focus allows us to consider exclusively the value of knowledge with relatively unfettered project choice options for researchers. Below, we will explore how the answers we derive here would change if CS like costs were added to the model.

2.5 Payoffs and Social Welfare

Researcher. For any period $t \in \{1, 2\}$, the short-lived researcher chooses a distance d_t (novelty). Their expected private payoff is:

$$u_R(d_t; X) = V(d_t; X) \quad (9)$$

This objective function captures the fundamental trade-off researchers face: greater novelty (distance) can increase benefits, but too much novelty could cause some decisions to become inactive as their variance is too high.

Planner. The planner maximises the following social welfare function:

$$W = v(\mathcal{F}_1) + \delta[v(\mathcal{F}_2)] \quad (10)$$

where $\delta \in (0, 1)$ is the social discount factor. Here, $v(\mathcal{F}_t)$ represents the total value of knowledge in period t . The planner considers both the immediate and future benefits of research. Crucially,

¹CS's specification has several important properties: (1) it is multiplicatively separable in ρ and (d, X) , (2) it is increasing and convex in ρ , (3) it establishes an endogenous link between novelty (distance) and research output (probability of success), and (4) it models research as a costly search process that may fail.

the planner's problem involves not just the optimal selection of research distances d_1 and d_2 , but also the optimal sequencing of expanding versus deepening knowledge. As we will show, this sequencing decision is critical for social welfare.

3 Researcher's Problem

In each period, the researcher faces a choice between expanding knowledge beyond the frontier or deepening knowledge within the existing interval. We analyse each option separately and then determine the researcher's optimal strategy.

3.1 Optimal Expansion

For expansion beyond the frontier, the researcher's problem is to maximise the following:

$$\max_{d>0} \left[d - \frac{d^2}{6q} + \mathbf{1}_{d>4q} \frac{\sqrt{d}(d-4q)^{3/2}}{6q} \right]. \quad (11)$$

As shown by Carnehl and Schneider (2025), the maximiser of $V(d; \infty)$ is $d^* = 3q$, which lies below $4q$. Hence, the indicator term is not operative at the optimum in the costless benchmark; it is included only for completeness.

This is a simple concave maximisation problem. The first-order condition (noting that the optimal $d \leq 4q$) is:

$$1 - \frac{d}{3q} = 0 \quad (12)$$

Solving for d yields the optimal expansion distance:

$$d_E^* = 3q \quad (13)$$

Given this, the researcher's payoff from expansion is:

$$u_R(\text{Expand}) = \frac{3q}{2} \quad (14)$$

3.2 Optimal Deepening

For deepening knowledge within the interval $[0, X_0]$, the researcher's problem is to choose $d \in [0, X_0/2]$ to maximize:

$$u_R(\text{Deepen}) = V(d; X_0) \quad (15)$$

The optimal deepening location depends on the length of the interval X_0 :

Case 1: Short interval ($X_0 \leq 4q$). In this case, the benefit function simplifies to $V(d; X_0) = \frac{d(X_0-d)}{3q}$. This function is strictly concave in d and maximised at the midpoint $d = X_0/2$. At this midpoint, the variance $\sigma^2(d; X_0)$ equals $X_0/4$. Therefore, the researcher's benefit from deepening at the midpoint is:

$$u_R(\text{Deepen}) = \frac{X_0^2}{12q} \quad (16)$$

Case 2: Intermediate interval ($4q < X_0 \leq \tilde{X}_0$). CS show that there exists a threshold $\tilde{X}_0 < 8q$ such that for intervals of length $X_0 \in (4q, \tilde{X}_0]$, the optimal deepening location remains at the midpoint $d = X_0/2$, though the benefit function now includes additional terms due to the indicator functions. The researcher’s benefit is more complex but still maximised at the midpoint.

Case 3: Long interval ($X_0 > \tilde{X}_0$). For intervals longer than the threshold \tilde{X}_0 , the optimal deepening location shifts away from the midpoint. CS show that for $X_0 > \tilde{X}_0$, the optimal distance $d_0(X_0)$ lies strictly between $3q$ and the midpoint: $d_0(X_0) \in (3q, X_0/2)$. As X_0 becomes very large, the optimal distance approaches $3q$ from above.

The intuition is that for very long intervals, deepening exactly at the midpoint creates two regions where variances exceed the threshold q , making knowledge less valuable there. By moving closer to one of the known points (but still maintaining distance $d > 3q$), the researcher can ensure that more questions have variances below the critical threshold.

3.3 Researcher’s Optimal Choice

The researcher chooses expansion over deepening if $u_R(\text{Expand}) \geq u_R(\text{Deepen})$. Comparing these values gives us the following result:

Proposition 1 (Private cut-off; cf. CS Prop. 2 and Lemma 9). *The researcher:*

- expands knowledge (at distance $d_E^* = 3q$) if $X_0 < X_R$;
- deepens knowledge if $X_0 > X_R$, where:
 - for $X_0 \leq \tilde{X}_0$, deepening occurs at the midpoint $d = X_0/2$
 - for $X_0 > \tilde{X}_0$, deepening occurs at distance $d_0(X_0) \in (3q, X_0/2)$

where the threshold is given by:

$$X_R = \hat{X} \approx 4.338 q.$$

The proofs of all results are in the appendix. Parts matching CS (cutoff value and midpoint vs. interior choice) restate their results in our two-period, zero-cost environment; any deviations are noted explicitly below. Proposition 1 provides a clear characterisation of the researcher’s decision rule: when the existing interval is sufficiently small ($X_0 < X_R$), the researcher prefers to expand knowledge; when the interval is sufficiently large ($X_0 > X_R$), the researcher prefers to deepen knowledge.

The intuition is straightforward. For small intervals, conjectures within that interval already have relatively high precision, so the marginal benefit of further deepening knowledge is

limited. In contrast, expanding knowledge creates entirely new opportunities for valuable conjectures. Conversely, when the existing interval is large, conjectures within that interval have low precision, making deepening knowledge more valuable than expanding.

For very large intervals ($X_0 > \tilde{X}_0$), the optimal deepening location shifts from the midpoint toward a “sweet spot” that balances the benefits of reducing variances with the costs of creating asymmetric intervals. This reflects a sophisticated trade-off inherent in knowledge creation: balancing precision against breadth of knowledge.

4 The Planner’s Problem

The social planner solves a two–period optimisation problem. In contrast to a myopic researcher who considers only the marginal impact of a single discovery, the planner internalises the fact that knowledge created in period 1 continues to benefit society in period 2. This section characterises the planner’s objective, analyses the second–period decision in isolation and then derives the optimal sequence of actions across both periods. We conclude by showing that, in the absence of research costs, the planner’s and researcher’s incentives coincide on the *extensive margin* (expand versus deepen). When the deepening problem has a unique maximiser, the chosen location coincides as well; later we show that location choices can diverge on the *intensive margin* in multidisciplinary settings.

4.1 Social welfare

Let $v(\mathcal{F}_t)$ denote the total value of knowledge in period t . Because knowledge once discovered remains available in all subsequent periods, the planner’s lifetime utility from the stock of knowledge is $v(\mathcal{F}_1) + \delta v(\mathcal{F}_2)$, where $\delta \in (0, 1)$ is the social discount factor. Write $V_1 = v(\mathcal{F}_1) - v(\mathcal{F}_0)$ and $V_2 = v(\mathcal{F}_2) - v(\mathcal{F}_1)$ for the incremental benefits of the first and second discoveries. Rearranging yields

$$W = v(\mathcal{F}_0)(1 + \delta) + (1 + \delta)V_1 + \delta V_2,$$

highlighting that the planner places a weight of $1 + \delta$ on the first discovery, whereas a myopic researcher counts only the period–1 gain.

4.2 Second–period incentives

To determine the planner’s optimal two–period strategy it is useful first to study the second–period problem conditional on the state created in period 1. At the start of period 2 there are two intervals, and the planner may either expand the frontier or deepen within one of these intervals. Let X denote the length of the candidate interval. For a given interval length X , the value of deepening depends on the best deepening choice $d_0(X) \in \arg \max_{d \in [0, X/2]} V(d; X)$ and the associated maximised value

$$M(X) := \max_{d \in [0, X/2]} V(d; X) = V(d_0(X); X).$$

Carnehl and Schneider (2025, Lemma 9) show that $M(X)$ crosses the expansion payoff $V(3q; \infty) = 3q/2$ at a unique cutoff $\hat{X}^0 \approx 4.338q$: $M(X) < 3q/2$ for $X < \hat{X}^0$ and $M(X) > 3q/2$ for $X > \hat{X}^0$. They also show that $M(X)$ attains an interior maximum near $X \approx 6.2q$ and converges back to $3q/2$ as $X \rightarrow \infty$, so $M(X)$ is not monotone on $(4q, \infty)$. A second threshold arises when one interval has been halved by first-period deepening: then the planner compares $3q/2$ with $M(X/2)$, which exceeds $3q/2$ precisely when $X > X_2^* := 2\hat{X}^0 \approx 8.676q$.

The following proposition summarises the second–period choice.

Proposition 2 (Optimal second–period action). *Let X be the length of the interval available for deepening in period 2. The planner compares $V(3q; \infty) = 3q/2$ with $V(d_0(X); X)$ and chooses:*

1. *expand the frontier if $X < \hat{X}^0 \approx 4.338q$;*
2. *deepen the interval if $X > \hat{X}^0$.*

When X is itself the result of a period–1 deepening, the length available in period 2 is $X/2$. In this case the frontier is expanded if $X < X_2^ \approx 8.676q$ and further deepening is chosen otherwise. These thresholds reflect the fact that $M(X) = V(d_0(X); X)$ crosses $3q/2$ at $X = \hat{X}^0$, attains an interior maximum near $6.2q$, and converges back to $3q/2$ as $X \rightarrow \infty$.*

These thresholds come from comparing the constant expansion payoff $3q/2$ to the maximised deepening payoff $M(X)$. In particular, after an expansion in period 1 the new interval has length $3q < \hat{X}^0$, so deepening in that interval yields less than $3q/2$; and since period 1 expansion occurs only when $X_0 < \hat{X}^0$, deepening within $[0, X_0]$ also yields less than $3q/2$. Hence period 2 optimally expands again.

4.3 Candidate strategies

In period 1 the planner may expand or deepen. The period–2 choice then follows Proposition 2. Four possible trajectories emerge:

1. **Expand–Expand (EE)**. Expand by $3q$ in period 1 and expand by $3q$ in period 2. The period–1 benefit is $V(3q; \infty) = \frac{3q}{2}$ and the period–2 benefit is the same. Summing gives

$$W_{EE} = v(F_0)(1 + \delta) + \frac{3q}{2}(1 + 2\delta).$$

2. **Expand–Deepen (ED)**: expand by distance d_1 in period 1, then deepen the newly created interval. If commitment were possible, the planner would solve $(1 + \delta)(d_1 - d_1^2/(6q)) + \delta \frac{d_1^2}{12q}$, yielding $d_1 = \frac{6q(1+\delta)}{2+\delta}$. This "moonshot" expands further than $3q$ to create a larger interval for deepening. Without commitment, however, period–2 deepening is optimal only if $d_1 \geq \hat{X}^0$, so any expand–deepen plan must choose $d_1 \geq 4.338q$; such a plan is always dominated by one of the other trajectories.

3. **Deepen–Expand (DE):** deepen in period 1 at the static optimum $d_0(X_0) \in \arg \max_{d \in [0, X_0/2]} V(d; X_0)$, generating payoff $M(X_0) := \max_{d \in [0, X_0/2]} V(d; X_0)$. For $X_0 \leq \hat{X}^0$, $M(X_0) < V(3q; \infty)$ so the planner never deepens first. For $\hat{X}^0 < X_0 < X_2^*$, the planner deepens in period 1 and then expands by $3q$ in period 2, so total welfare is

$$W_{DE} = v(\mathcal{F}_0)(1 + \delta) + (1 + \delta)M(X_0) + \delta V(3q; \infty).$$

(When $X_0 \leq 4q$, $d_0(X_0) = X_0/2$ and $M(X_0) = X_0^2/(12q)$.)

4. **Deepen–Deepen (DD):** deepen in period 1 at $d_0(X_0)$ and deepen again in period 2 in whichever induced subinterval yields the larger deepening payoff. Let the induced lengths be $X_L = d_0(X_0)$ and $X_R = X_0 - d_0(X_0)$. Then the period–2 deepening payoff is $\max\{M(X_L), M(X_R)\}$, and total welfare is

$$W_{DD} = v(\mathcal{F}_0)(1 + \delta) + (1 + \delta)M(X_0) + \delta \max\{M(d_0(X_0)), M(X_0 - d_0(X_0))\}.$$

4.4 Optimal strategy

Combining the second–period thresholds with the candidate welfare expressions yields the planner’s optimum as a function of the initial X_0 .

Proposition 3 (Planner’s optimal strategy). *Let X_0 denote the length of the initial knowledge gap. Then*

1. *If $X_0 < \hat{X}^0 \approx 4.338q$ it is optimal to expand in both periods (EE).*
2. *If $\hat{X}^0 \leq X_0 < X_2^* \approx 8.676q$ it is optimal to deepen in period 1 and expand in period 2 (DE).*
3. *If $X_0 \geq X_2^*$ it is optimal to deepen in both periods (DD).*

The expand–deepen trajectory (ED) is never optimal, because the interval created by a socially optimal expansion in period 1 is smaller than \hat{X}^0 and therefore does not warrant deepening in period 2.

The cutoffs \hat{X}^0 and X_2^* reveal how the planner balances the trade-off between creating new knowledge and refining existing knowledge. When the initial gap is small ($X_0 < \hat{X}^0$), the interval between known points is already reasonably precise, so deepening provides little additional value. In this region, it is better to expand in both periods, as each expansion yields the same benefit and opens up new questions where conjectures can be applied. For intermediate gaps ($\hat{X}^0 \leq X_0 < X_2^*$), a single deepening is worthwhile to split the moderately large interval and reduce variances; but once the interval has been halved, the gains from further deepening are limited and the planner prefers to expand the frontier in period 2. For large gaps ($X_0 \geq X_2^*$), the

existing interval is so long that successive divisions create much more precision than venturing into unexplored territory, so two rounds of deepening dominate expansion. The expand–deepen trajectory (ED) is never optimal because a socially optimal expansion in period 1 creates an interval of length $3q$, which is smaller than \hat{X}^0 and thus does not justify deepening in period 2. Instead, any expansion is followed by another expansion, while deepening is reserved for sufficiently long intervals.

4.5 Alignment with private incentives

The absence of research costs implies that the private researcher and the social planner share the same benefit function. The only difference is that the planner discounts future payoffs. Because the ranking of strategies does not depend on the weight attached to period 1, the planner and the researcher choose the same action in each period. Specifically:

Proposition 4 (Alignment on the extensive margin). *In both periods, the researcher’s and the planner’s optimal action type coincides: they expand iff the relevant interval length is below \hat{X}^0 and deepen otherwise. In particular, the expand–vs.–deepen thresholds \hat{X}^0 (and, when applicable, X_2^*) are independent of δ , so private and social incentives are aligned on the extensive margin.*

In our two-period, zero-cost setting, the private researcher and the social planner evaluate discoveries using the same benefit function. Discounting changes the weight placed on period 1 vs. period 2, but it does not alter the expand–vs.–deepen comparison within a given period because expansion always yields the fixed payoff $V(3q; \infty) = 3q/2$, while the best deepening payoff depends only on the relevant interval length through $M(X) = \max_d V(d; X)$. Thus both agents use the same cutoff \hat{X}^0 (and, when applicable, the same X_2^*) to decide whether to expand or deepen. When deepening has a unique maximiser (e.g. for intervals where $d_0(X) = X/2$), the chosen location coincides as well; when the deepening problem has a nontrivial interior solution, the intensive margin can diverge, as shown below.

5 Why is a Moonshot *not* Socially Desirable?

Our finding that the expand–deepen (ED) strategy is never socially optimal is formally present in the model of Carnehl and Schneider (2025). They, however, dedicate significant attention to the value of “moonshots” that are followed by deepening as one of the conclusions from their model. Understanding this requires careful examination, as it potentially challenges the framing around the policy implications of their analysis.

Although the model here is a special case of theirs, it is worthwhile to consider several generalisations of their model. Each is examined in turn because understanding which factors matter can help clarify the foundation of the case for a moonshot.

5.1 Time Horizon

One of the key simplifications made in this paper is to examine a two-period time horizon for the planner rather than an infinite horizon as in CS. This allows them to analyse research trajectories over many periods, where the structure of knowledge affects generations of future researchers. In contrast, our two-period model captures only the immediate effects of research decisions. The cumulative, long-run benefits of creating specific knowledge structures through moonshots may not be fully captured in our finite analysis.

There is a good reason, however, to suppose that the two-period horizon is not responsible for this.² CS examine possible evolutions that involve a steady “step ladder” of improvements whereby knowledge is expanded at a steady rate over time. They then compare this to the possibility of “research cycles,” where expansion is followed by deepening until knowledge becomes sufficiently dense, allowing expansion to take place again, and so on.

It is, however, very apparent that the social planner would not choose a research cycle (intermittent, relatively large expansion pushes) over other dynamic strategies. In the two-period model, expansion followed by deepening was dominated by the other two-period strategies. This will remain true even if there were more periods that a moonshot would influence. Recall that even with a single period, the expansion distance chosen by the planner in period 1 is insufficient to encourage deepening in period 2. Increasing the number of periods would enhance the novelty associated with the initial expansion, but at best, this would result in a single round of deepening in future periods. The alternative, engaging in a steady expansion, would ultimately generate the same outcome but without the “sacrifice” associated with a persistent but costly expansion in the first period. Therefore, research cycles will not be chosen by the social planner for the same reason. In other words, it is the nature of discovery itself, rather than a greater scope for intertemporal options, that drives the social preference against moonshots.

5.2 Cost Structure and Success Probabilities

CS (Proposition 6) demonstrate that a moonshot that relies on the planner choosing the first period expansion followed by researchers choosing later deepening would not be optimal for the planner if research was costless; as it is here in this paper. In their baseline model, research

²As an example, consider the following: normalise $q = 1$ and consider three periods with $\delta = 1$, so that the planner’s objective weights the three marginal discoveries by $(3, 2, 1)$. The “baseline” path of expanding by 3 in each period yields marginal values $(V_1, V_2, V_3) = (1.5, 1.5, 1.5)$ and total incremental welfare $3 \cdot 1.5 + 2 \cdot 1.5 + 1 \cdot 1.5 = 9$.

In contrast, consider the most favourable moonshot-style deviation: expand by some $d > 4$, deepen that newly created interval in the next period (yielding at most $\max_L M(L) \approx 1.85$), and then expand again. Even at the deepening peak ($d \approx 6.2$), one obtains $V(6.2; \infty) \approx 1.15$ in the first period and $M(6.2) \approx 1.85$ in the second, giving total incremental welfare $3 \cdot 1.15 + 2 \cdot 1.85 + 1 \cdot 1.5 \approx 8.64 < 9$. Thus, longer horizons do not mechanically overturn the first-best benchmark conclusion: moonshots remain a second-best mechanism that requires research costs (or other constraints) to be welfare-relevant.

is costly in that a researcher faces a payoff:

$$\rho V(d; X_0) - c(\rho, d) \quad (17)$$

where $\rho \in [0, 1]$ is the success probability of the research. When there are no costs the researcher sets $\rho = 1$ as in the model here. Otherwise, the researcher will set $\rho < 1$ and, moreover, the researcher will, when expanding knowledge, set $d < 3q$.

Critically, when researchers face costs, this drives a wedge between the social planner's and researcher's interests. First, the social planner places additional weight on the value of immediate improvements in knowledge as that knowledge persists through both periods with a weight of $1 + \delta > 1$, implying that the planner would set ρ higher than the level set by the researcher. Second, the planner takes into account the impact of today's knowledge expansion on the future cost of research. As noted earlier, these costs depend on $\sigma^2(d, X) = \frac{d(X-d)}{X}$ where X is the length of the research area. This means that greater knowledge expansion in period 1 can reduce the cost of research in period 2. If deepening is anticipated in the future, a researcher there will have a higher probability of success if the area in which deepening takes place is closer to existing points of knowledge.

Herein lies CS's result regarding moonshots. When research is costly, a planner would prefer to reduce the costs of that research in future periods and so may expand knowledge in period 1 by correspondingly more. To be sure, that is doubly costly for the planner. Not only is the period 1 moonshot expansion suboptimal in the present, but that suboptimality persists. Nonetheless, they show that for intermediate levels of research cost, a moonshot may be optimal in these circumstances. However, if research is too costly, the probability of success in period 2 is so low that a moonshot is no longer desirable.

The point here is that a moonshot is decidedly a second-best optimal outcome only. No planner actually wants to conduct a moonshot. As we have shown, it is too costly relative to simply expanding in a step ladder approach. However, when research is costly, there are intertemporal externalities impacting the productive efficiency of research. This enables there to be some rationales for a moonshot-type approach. However, it must be noted that even in this case, the strategy being implemented is more 'shot' than 'moon.'

Therefore, it would certainly be reasonable to conclude that, in fact, the CS knowledge framework does not imply that moonshots are optimal. Indeed, from a strict perspective, it provides a clear argument that they are not unless certain other distortions and/or additional intertemporal effects are added to the model. However, even with these, such effects mitigate the first-order implication of the framework against a moonshot policy.

6 Fostering Multidisciplinary Research

While Carnehl and Schneider (2025) showed that private and social incentives for knowledge creation are aligned with respect to moonshots when there are no research costs, here we iden-

tify a related case where there is misalignment: namely, in multidisciplinary research contexts with large knowledge gaps. This is a central challenge in science policy: to advance knowledge in multidisciplinary domains—areas that lie between established fields with significant knowledge gaps. In this section, we extend our model to analyse the strategic considerations in multidisciplinary research, where the initial knowledge gap X_0 represents the distance between two separate disciplines or research domains.

6.1 The Multidisciplinary Research Problem

Consider our model with a large knowledge gap $X_0 > \tilde{X}_0 \in (6q, 8q)$ between two established knowledge points at positions 0 and X_0 . This gap represents the distance between two distinct disciplines, each with its own established research traditions. The fundamental question is: how should researchers bridge this gap to maximise social welfare?

When analysing this problem, we need to determine whether researchers should deepen knowledge near existing discipline boundaries or target the central region between disciplines. We prove that even without research costs, there exists a fundamental misalignment between private researcher incentives and social welfare in multidisciplinary contexts.

6.2 Researcher Behaviour in Multidisciplinary Settings

First, we characterise how individual researchers approach bridging disciplines when faced with large knowledge gaps.

Proposition 5 (Researcher’s Multidisciplinary Choice). *For a knowledge gap $X_0 > 6q$, the myopic researcher deepens in period 1. There exists a threshold $\tilde{X}_0 \in (6q, 8q)$ (identified by CS) such that the optimal deepening location switches from the midpoint to an interior point closer to a discipline boundary as the gap grows. Then:*

1. For $6q < X_0 \leq \tilde{X}_0$, the researcher chooses $d_1^R = X_0/2$.
2. For $X_0 > \tilde{X}_0$, the researcher chooses $d_1^R \in (3q, X_0/2)$.
3. As $X_0 \rightarrow \infty$, the researcher’s optimal distance approaches $d_1^R \rightarrow 3q$ from above.

In period 2, the researcher deepens only if at least one interval exceeds $4.338q$; otherwise they expand. When deepening is feasible in both subintervals, the choice is not “always the longer one”: the argmax of $V(d; X)$ is interior and approaches $3q$ for long X , so the researcher selects the interval that yields the larger $V(\cdot)$, which need not be the longer interval.

The researcher’s strategy creates an unbalanced knowledge structure that favours one side of the interdisciplinary space. Once $X_0 > \tilde{X}_0$, by choosing to deepen knowledge closer to one discipline boundary than the midpoint, the researcher maximises immediate benefits but creates an inefficient trajectory for bridging the disciplines.

6.3 Socially Optimal Multidisciplinary Bridging

In contrast to the myopic researcher, a social planner with a positive discount factor would choose a different bridging strategy.

Proposition 6 (Planner’s Optimal Multidisciplinary Strategy). *Fix $X_0 > \tilde{X}_0$ and $\delta \in (0, 1)$. Let $d_1^R \in \arg \max_{d \in [0, X_0/2]} V(d; X_0)$ denote the myopic researcher’s first-period deepening location, and define $M(X) := \max_{d \in [0, X/2]} V(d; X)$.*

If the planner deepens in period 1, the planner’s optimal first-period location d_1^P solves

$$d_1^P \in \arg \max_{d \in [0, X_0/2]} \left\{ (1 + \delta)V(d; X_0) + \delta \max \{V(3q; \infty), M(d), M(X_0 - d)\} \right\}.$$

Moreover, $d_1^P \rightarrow d_1^R$ as $\delta \rightarrow 0$, and for any fixed δ we have $d_1^P \rightarrow 3q$ as $X_0 \rightarrow \infty$. In general $d_1^P \neq d_1^R$ for intermediate gap lengths because the continuation value depends on how period 1 splits the interval, and $M(\cdot)$ is not monotone.

The social planner chooses a deepening location that balances two competing considerations: (1) maximising immediate benefit in period 1, and (2) creating an optimal knowledge structure for period 2 research. The planner’s first-period location reflects an intertemporal tradeoff: it balances the immediate payoff $V(d; X_0)$ against how the choice splits the interval and thereby changes the best continuation value in period 2. Because $M(\cdot)$ is not monotone, the planner’s optimal adjustment relative to the myopic location need not move monotonically toward the midpoint. In the numerical example below ($X_0 = 11q$, $\delta = 1$), the planner chooses a location slightly closer to $X_0/2$ than the myopic researcher.

6.4 Welfare Loss from Misaligned Incentives

The misalignment between researcher behaviour and social welfare leads to significant welfare losses in multidisciplinary contexts.

Proposition 7 (Welfare Loss in Multidisciplinary Research). *Fix $X_0 > \tilde{X}_0$ and $\delta \in (0, 1)$. Let W_P denote the planner’s maximal welfare and let W_R denote welfare under the myopic researcher’s choices. Then $\Delta W := W_P - W_R \geq 0$, with strict inequality for a nonempty set of parameter values. Moreover,*

$$\lim_{\delta \rightarrow 0} \Delta W = 0 \quad \text{and} \quad \lim_{X_0 \rightarrow \infty} \Delta W = 0.$$

This non-monotonic behaviour of the welfare loss with respect to X_0 is intuitive: for moderate gap sizes, the planner’s more balanced approach creates significant advantages, but as the gap becomes extremely large, the relative differences between strategies become negligible compared to the total gap size.

6.5 Illustrative Example: Bridging a large Knowledge Gap

To illustrate these principles concretely, consider a knowledge gap of $X_0 = 11q$ with discount factor $\delta = 1$. We calculate and compare the welfare outcomes for three distinct strategies: the myopic researcher’s choice, the social planner’s optimal choice, and a seemingly intuitive equal-spacing approach. Throughout this example, discovery positions are reported as absolute locations along the gap $[0, X_0]$ (measured from the left endpoint).

Table 1 summarises the three strategies for bridging the knowledge gap and their outcomes.

Strategy	Researcher	Planner	Equal-Spacing
First discovery position (d_1)	$3.05q$	$3.42q$	$3.67q$
Second discovery position	$6.33q$	$7.21q$	$7.33q$
Resulting intervals	$\{3.05q, 3.28q, 4.67q\}$	$\{3.42q, 3.79q, 3.79q\}$	$\{3.67q, 3.67q, 3.67q\}$
First period benefit (V_1)	$1.534q$	$1.512q$	$1.472q$
Second period benefit (V_2)	$1.614q$	$1.681q$	$1.735q$
Total welfare (W)	$4.683q$	$4.704q$	$4.679q$
Relative performance	-0.45%	<i>Optimal</i>	-0.52%
Knowledge space activated ³	83.9%	100%	100%

Table 1: Comparison of strategies for bridging a knowledge gap of $X_0 = 11q$. All entries are computed from the full benefits function in (6) using the case-appropriate indicator terms; see Appendix A.8 for derivations. The ordering between the *researcher* and *equal-spacing* strategies is not generic and can flip when X_0 is near multiples of $3q$ (e.g. $X_0 \approx 9q$).

Each strategy creates a different knowledge structure:

- Researcher strategy: Deepens at $3.05q$ in period 1, then deepens at $6.33q$ in period 2, creating intervals $\{3.05q, 3.28q, 4.67q\}$.
- Planner strategy: Deepens at $3.42q$ in period 1, then deepens at $7.21q$ in period 2, creating intervals $\{3.42q, 3.79q, 3.79q\}$.
- Equal spacing: Deepens at $11q/3 \approx 3.67q$ in period 1, then deepens at $22q/3 \approx 7.33q$ in period 2, creating symmetric intervals $\{3.67q, 3.67q, 3.67q\}$.

The welfare calculations are reported in Appendix A.8. Since a period 1 discovery is useful in both periods while a period 2 discovery affects only the second period, total welfare takes the

³Activation counts points with $\sigma^2 < q$. Equality cases $\sigma^2 = q$ occur on a set of measure zero and are ignored.

form

$$W = (1 + \delta)V_1 + \delta V_2.$$

With $\delta = 1$, this reduces to $W = 2V_1 + V_2$.

As shown in Table 1, the planner’s strategy achieves the highest welfare at $W = 4.704q$, followed by the researcher’s $W = 4.683q$, and equal spacing at $W = 4.679q$. The differences are small in welfare terms (on the order of half a percent), but the induced interval structure differs sharply.

To see what drives the ranking, it is helpful to decompose the tradeoffs using $W = 2V_1 + V_2$:

- Relative to the *researcher*, the planner accepts a lower first–period benefit (V_1 falls from $1.534q$ to $1.512q$) in exchange for a higher second–period benefit (V_2 rises from $1.614q$ to $1.681q$). Because V_1 is weighted twice, the planner’s gain comes from making the period 2 opportunity substantially better without sacrificing too much in period 1.
- Relative to *equal spacing*, the planner obtains a higher first–period benefit ($1.512q$ vs. $1.472q$) but a lower second–period benefit ($1.681q$ vs. $1.735q$). The planner still wins overall because the period 1 term is double–counted when $\delta = 1$.

These calculations reveal three takeaways:

1. **The planner’s advantage is an intertemporal tradeoff, not ”symmetry for its own sake.”** In this example the planner places the first discovery at $d_1 = 3.42q$, closer to the large–interval deepening optimum near $3q$ than equal spacing ($3.67q$), while still leaving a remaining interval that can be split almost evenly in period 2.
2. **Equal spacing does very well on coverage and on V_2 , but it underweights the value of the first discovery.** Equal spacing produces three equal subintervals below $4q$ (hence full activation) and delivers the highest second–period benefit in the table. It nevertheless performs slightly worse overall because its first discovery at $X_0/3$ is too far from the large–interval optimum, and that first–period shortfall is magnified by the $(1 + \delta) = 2$ weight.
3. **The myopic researcher places the first discovery near the static optimum, but leaves an overly long residual interval.** The researcher’s first discovery at $3.05q$ yields the highest V_1 , but the resulting partition includes a $4.67q$ subinterval, which exceeds the $4q$ dead–zone threshold. This lowers both the second–period payoff and the activated share of the gap (83.9% in Table 1).

In short, equal spacing underperforms here not because balance is bad, but because the benefit function (and the welfare weights) make the *timing* of variance reduction matter: the planner prefers to keep the first discovery closer to the high–value region near $3q$ while using the second discovery to eliminate dead zones and improve the continuation payoff.

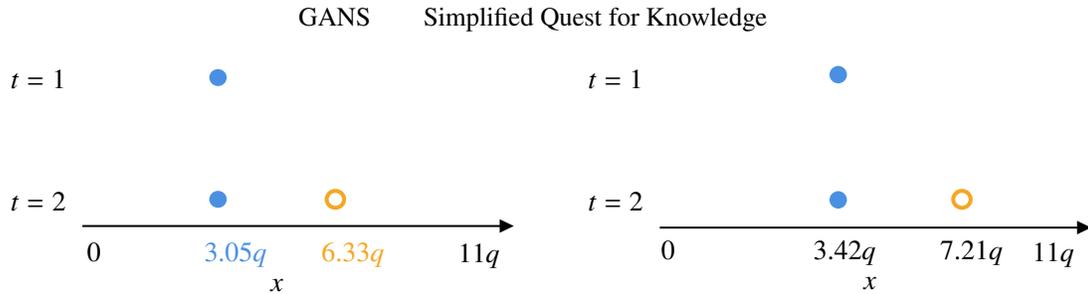


Figure 1: **Fostering Multidisciplinary Research with $X_0 = 11q$.** The dots show which questions have a known answer at each time t . The left side shows the researcher’s strategy (creating a more asymmetric structure), whereas the right-hand side shows the planner’s strategy (creating a more balanced structure).

6.6 Knowledge Activation and Bridging Efficiency

To understand the practical implications of different bridging strategies, we introduce the concept of “knowledge activation.” A point in the knowledge space is “activated” when conjectures at that point have variance below q , making them useful for decision-making.

After two periods, the researcher’s strategy creates activated regions concentrated near one discipline boundary, with a substantial portion of the interdisciplinary space remaining “unactivated.” In contrast, the planner’s strategy creates a more balanced activation pattern across the interdisciplinary space, efficiently bridging the gap between disciplines. This is depicted in Figure 1. The planner’s strategy activates significantly more of the interdisciplinary space than the researcher’s strategy after two periods, whereas the intuitively appealing equally spaced strategy falls short of the optimal approach.

The relevant percentages are contained in Table 1. The planner and equal-spacing strategies activate 100% of the gap, while the researcher’s strategy activates only 83.9% due to a remaining dead zone in the largest ($4.67q$) subinterval. Thus, the researcher’s short-run focus can leave a persistent region where knowledge cannot be applied.

Full activation, however, is not sufficient for maximal welfare. Even when the entire gap is activated (as under equal spacing), welfare still depends on average variance: the planner’s placement of the first deepening closer to the large-interval optimum near $3q$ raises the (heavily weighted) first-period payoff while still producing a well-balanced structure for the second period.

This analysis demonstrates that the optimal strategy for bridging multidisciplinary knowledge gaps involves nuanced placement decisions that balance several competing factors: reducing variance across the knowledge space, creating manageable interval lengths, and strategically positioning discoveries to maximise future research value.

The strategic considerations in multidisciplinary research become even more pronounced when we extend our analysis to three periods. Suppose the planner has a three-period hori-

zon and chooses two deepening discoveries followed by an expansion. Using the same two-period bridging configurations from Table 1 with $X_0 = 11q$ and $\delta = 1$, the researcher’s strategy yields knowledge points approximately $\{0, 3.05q, 6.33q, 11q, 14q\}$, while the planner yields approximately $\{0, 3.42q, 7.21q, 11q, 14q\}$. The activation gap within the original $[0, 11q]$ interval persists (about 83.9% vs. 100%) even after the subsequent expansion.

6.7 Policy Implications for Multidisciplinary Research

Our analysis yields several important policy implications for funding multidisciplinary research:

1. **Strategic Placement:** Funding agencies should incentivise research positions that differ from those that would emerge naturally from researchers’ myopic incentives, particularly positions that create more balanced knowledge structures. However, the optimal positions are not necessarily those that create perfectly equal intervals.
2. **Long-Term Perspective:** The case for interventions in multidisciplinary research strengthens with the planner’s patience (higher δ). Long-term funding initiatives should direct research toward strategic positions that may not yield the highest immediate payoffs.
3. **Coordinated Programs:** Multidisciplinary initiatives should be designed as coordinated programs rather than independent projects, allowing for strategic sequencing of knowledge creation.
4. **Balance Metrics:** Evaluating multidisciplinary programs should include metrics that assess how completely and evenly they “activate” the knowledge space between disciplines.
5. **Beyond Bridging:** Our three-period analysis shows that after establishing a balanced multidisciplinary foundation, expansion beyond the disciplinary boundaries becomes optimal. Funding strategies should account for this natural evolution rather than indefinitely focusing on deepening the interdisciplinary space.

The misalignment between private and social incentives in multidisciplinary research represents a distinct market failure not addressed in Carnehl and Schneider (2025). While their model implies alignment between private and social incentives without research costs, we demonstrate that large knowledge gaps between disciplines create a context where misalignment persists even without costs.

This finding underscores the importance of strategic intervention in multidisciplinary research. By directing research toward positions that create more balanced and efficient knowledge structures over time, funding agencies can accelerate the bridging of disciplinary divides and maximise the value of interdisciplinary knowledge creation. However, our analysis also highlights that the optimal intervention is relatively mild—nudging researchers toward more

balanced structures rather than imposing perfectly symmetric ones—and should evolve as the knowledge landscape changes.

7 Citation-Based Incentives in Knowledge Creation

Existing models of research behaviour often assume that scientists aim to maximise the increase in social welfare their discoveries generate (the marginal value V). In reality, academic careers are built on citations: papers are rewarded when others use them, and credit is shared across multiple contributors. This suggests researchers may instead prioritise the total value of the knowledge that relies on their contribution, subject to credit-sharing norms. To capture this contrast between marginal social value and attributed recognition, we extend our baseline model by introducing a formal citation-based reward system.

A formal citation mechanism. We formalise citation rewards based on how a researcher’s discovery is utilised by decision makers. We introduce a credit-sharing parameter $\alpha \in (0, 1)$. When a decision maker uses the researcher’s newly discovered point alongside an existing point (i.e., it bounds a finite interval), the researcher receives a fraction α of the value created in that interval.

In contrast, when the new point is the sole basis for answering a question (i.e., it bounds an unbounded frontier), the researcher receives full credit (100%) for the resulting value. The parameter α thus measures the degree of credit sharing in the scientific community: low values place significant weight on unique contributions at the frontier, while high values spread credit more evenly across collaborators who refine existing knowledge.

Throughout this section, we focus on initial intervals satisfying $X_0 \leq 4q$, ensuring there are no “dead zones.” In this regime, the value of an interval of length L is $v(L) = L - L^2/(6q)$, and the value generated beyond an unbounded frontier is $v_U = q/2$.

7.1 The researcher’s objective under citation rewards

We first derive the citation rewards for expansion and deepening.

Reward for expansion. When a researcher expands knowledge by a distance d , they create a new point at $X_0 + d$. This discovery generates shared value in the new interval $[X_0, X_0 + d]$ and unique value beyond the new frontier. The total citation reward for expansion, $C_E(d)$, is:

$$C_E(d) = \alpha v(d) + v_U = \alpha \left(d - \frac{d^2}{6q} \right) + \frac{q}{2}. \quad (18)$$

To maximise this reward, the researcher chooses the distance d that maximises $v(d)$, which is $d_E^* = 3q$. This location coincides with the socially optimal expansion distance. The optimal

citation reward for expansion is:

$$C_E^* = C_E(3q) = \alpha \left(\frac{3q}{2} \right) + \frac{q}{2} = \frac{q}{2}(3\alpha + 1). \quad (19)$$

Reward for deepening. When a researcher deepens knowledge by adding a point at distance d within $[0, X_0]$, they create two new intervals: $[0, d]$ and $[d, X_0]$. Both intervals are bounded and thus shared with existing knowledge. The total citation reward for deepening, $C_D(d)$, is:

$$C_D(d) = \alpha v(d) + \alpha v(X_0 - d) = \alpha \left[X_0 - \frac{d^2 + (X_0 - d)^2}{6q} \right]. \quad (20)$$

This reward is maximised when the interval is split at its midpoint, $d_D^* = X_0/2$, again aligning with the socially optimal location. The optimal citation reward for deepening is:

$$C_D^* = C_D(X_0/2) = \alpha \left(X_0 - \frac{X_0^2}{12q} \right). \quad (21)$$

Crucially, citation rewards do not distort *where* researchers add knowledge, as the optimal locations remain unchanged. Instead, they affect *whether* researchers choose to expand or deepen by altering the relative payoffs.

7.2 How citation incentives affect research choices

The trade-off between expanding and deepening depends on the comparison between C_E^* and C_D^* . The expansion reward is constant, while the deepening reward increases with X_0 (for $X_0 \leq 4q$).

We determine if expansion can be dominant regardless of X_0 by comparing C_E^* with the maximum possible deepening reward, which occurs at $X_0 = 4q$:

$$\max_{X_0 \leq 4q} C_D^*(X_0) = C_D^*(4q) = \alpha \left(4q - \frac{16q^2}{12q} \right) = \frac{8\alpha q}{3}. \quad (22)$$

The researcher always prefers expansion if $C_E^* > \max C_D^*$:

$$\frac{q}{2}(3\alpha + 1) > \frac{8\alpha q}{3} \implies 3(3\alpha + 1) > 16\alpha \implies 3 > 7\alpha \implies \alpha < \frac{3}{7}. \quad (23)$$

If the credit-sharing parameter is sufficiently small ($\alpha < 3/7$), the premium placed on the unique contribution at the frontier ($q/2$) outweighs the potential gains from sharing credit on even the longest permissible interval.

When $\alpha \geq 3/7$, the choice depends on X_0 . The researcher expands if $C_E^* \geq C_D^*$. This defines a threshold $X_R^C(\alpha)$ such that the researcher expands if $X_0 < X_R^C(\alpha)$ and deepens otherwise. This threshold, derived by solving $C_E^* = C_D^*(X_0)$, is decreasing in α . As credit sharing becomes more generous (high α), deepening becomes relatively more attractive.

A citation-maximising researcher's strategy depends critically on the credit-sharing regime. If rewards strongly favour unique contributions ($\alpha < 3/7$), the researcher always chooses to expand. If credit is shared more generously ($\alpha \geq 3/7$), the researcher expands only if the initial interval is short ($X_0 < X_R^C(\alpha)$) and deepens otherwise.

7.3 Second-period behaviour and dynamic implications

A second-period researcher inherits the knowledge generated in period 1 and faces a similar trade-off.

After an expansion. If the first researcher expanded, the knowledge structure consists of intervals of length X_0 and $3q$. The second researcher chooses between expanding again (C_E^*), deepening the original interval ($C_D^*(X_0)$), or deepening the new interval ($C_D^*(3q)$).

Since the first researcher chose to expand, we know $C_E^* > C_D^*(X_0)$. Therefore, deepening the original interval is never optimal in the second period. The choice is between expanding again or deepening the new interval (of length $3q$). We compare C_E^* with $C_D^*(3q) = 9\alpha q/4$.

$$C_E^* > C_D^*(3q) \iff \frac{q}{2}(3\alpha + 1) > \frac{9\alpha q}{4} \iff 2(3\alpha + 1) > 9\alpha \iff \alpha < \frac{2}{3}. \quad (24)$$

If $\alpha < 2/3$, the second researcher expands again. If $\alpha > 2/3$, they deepen the newly created interval. This stark threshold highlights how changes in credit sharing can alter the trajectory of knowledge creation.

After a deepening. If the first researcher deepened, the knowledge structure consists of two intervals of length $X_0/2$. The second researcher chooses between expanding (C_E^*) or deepening one of the subintervals ($C_D^*(X_0/2)$).

Given the constraint $X_0 \leq 4q$, the subintervals have length at most $2q$. The deepening reward $C_D^*(2q) = \alpha(2q - 4q^2/12q) = 5\alpha q/3$. Comparing this to the expansion reward:

$$C_E^* > C_D^*(2q) \iff \frac{q}{2}(3\alpha + 1) > \frac{5\alpha q}{3} \iff 3(3\alpha + 1) > 10\alpha \iff 3 > \alpha. \quad (25)$$

Since $\alpha < 1$, this condition is always satisfied. Therefore, when the first researcher deepens, the second researcher always chooses to expand. The initial deepening sufficiently reduces the gap such that the frontier becomes the most attractive option.

7.4 Alignment with the social optimum

Section 4 established that for $X_0 \leq 4q$, the marginal social value of expansion ($3q/2$) always exceeds the marginal social value of deepening (at most $4q/3$). Therefore, the socially optimal strategy is always to expand in both periods (EE). Do citation incentives support this outcome?

We analyse the conditions under which private incentives yield the EE path.

- **Low credit sharing** ($\alpha < 3/7$). The first researcher always expands. The second researcher also expands (since $\alpha < 3/7 < 2/3$). Private incentives perfectly align with the social optimum, resulting in the EE path.
- **Intermediate credit sharing** ($3/7 \leq \alpha < 2/3$). The first researcher may deepen if X_0 is large (leading to DE). Alignment (EE) occurs only if X_0 is sufficiently small.
- **High credit sharing** ($\alpha \geq 2/3$). Even if the first researcher expands (because X_0 is small), the second researcher will deepen (leading to ED). If X_0 is large, the first researcher deepens (leading to DE). Misalignment is pervasive in this regime.

The analysis reveals that citation incentives align with social welfare only under strong rewards for unique contributions ($\alpha < 3/7$). For larger values of α , the citation system over-incentivises deepening.

The misalignment stems from the fundamental difference between the researcher’s objective (attributed share of *total* value) and the social objective (*marginal* increase in value). We can quantify this distortion by examining the ratio of the citation reward to the marginal social value created ($R = C/V$).

- For expansion, $R_E = C_E^*/V_E^* = \alpha + 1/3$.
- For deepening, $R_D = C_D^*/V_D^* = \alpha(12q/X_0 - 1)$.

When $R_D > R_E$, the citation system provides a relatively stronger incentive for deepening than its social value warrants. Considering the case where deepening is most attractive ($X_0 = 4q$), $R_D = 2\alpha$. $R_D > R_E$ when $2\alpha > \alpha + 1/3$, or $\alpha > 1/3$.

Thus, when $\alpha > 1/3$, the citation system inherently favours deepening large intervals over expansion, relative to the social optimum. When this bias is strong enough (specifically, when $\alpha \geq 3/7$), it leads researchers to make socially suboptimal choices.

7.5 Implications

Two key lessons emerge from this formal analysis. First, citation systems that strongly privilege novel, independent contributions (low α) encourage researchers to push the frontier outward. This aligns private incentives with social welfare when knowledge gaps are modest ($X_0 \leq 4q$). Second, systems that grant substantial credit for shared contributions (high α) may induce excessive deepening, leading researchers to refine existing domains even when society would benefit more from expansion. These findings underline the importance of designing scientific reward systems that carefully balance recognition for collaborative refinement against the imperative to discover and explore new questions.

8 Conclusion

This paper presents a simplified version of the Carnehl and Schneider (2025) model of knowledge creation, making their rich framework more accessible and yielding closed-form solutions for key welfare trade-offs. Our analysis has generated three main insights that both complement and challenge aspects of the original model.

First, we have demonstrated that the “moonshot” approach - expanding knowledge beyond the frontier with the intention of later deepening - is never the socially optimal research trajectory in direct welfare comparisons. Under our transparent welfare calculations, the social planner always prefers either consistent expansion or strategies that begin with deepening knowledge, depending on the initial knowledge gap. This finding clarifies an important subtlety in CS: their case for moonshots depends crucially on research costs and second-best considerations, not on the direct welfare effects of different knowledge structures. The moonshot strategy becomes desirable only when research is costly and researchers’ choices deviate from social optimality due to cost-related distortions. Our analysis helps ground policy discussions about high-risk, high-reward research in precise welfare economics rather than intuitive but potentially misleading narratives.

Second, we have identified a novel misalignment between private and social incentives in multidisciplinary research contexts. When bridging significant knowledge gaps between disciplines, researchers systematically choose locations that deepen understanding near disciplinary boundaries, creating knowledge structures that are inefficient relative to the social optimum. Crucially, this misalignment persists even without research costs. This finding has significant implications for funding agencies seeking to promote interdisciplinary collaboration. It suggests that targeted interventions may be explicitly warranted for multidisciplinary research, even when other distortions are minimal.

Third, our simplified model has provided precise characterisations of optimal knowledge creation paths under various initial conditions. We have derived exact thresholds for when researchers and social planners should expand versus deepen knowledge, and shown how these choices depend on the initial knowledge gap and social discount factor. These clear, quantitative insights offer practical guidance for science policy design. They allow policymakers to identify precisely when intervention is warranted and what form it should take.

The simplified framework developed in this paper offers several advantages for future research and policy applications. By stripping away non-essential complexities, we have created a transparent model that directly exposes the fundamental trade-offs in knowledge creation. This clarity makes it easier to extend the model to address other vital questions in the economics of science, such as the impact of different funding mechanisms, the role of researcher heterogeneity, or the effects of competition in the research process.

For science policy, our model provides clear guidance on when intervention is warranted and what form it should take. The precise thresholds derived here can inform funding decisions

about whether to support the expansion of knowledge or the deepening of knowledge in different contexts. Our analysis of multidisciplinary research offers a rationale for special funding schemes targeted at bridging disciplinary divides. And our examination of the conditions under which moonshots may or may not be justified provides a framework for evaluating high-risk, high-reward research initiatives.

Looking ahead, several promising directions for future research emerge from our analysis. One avenue is to explore how different funding mechanisms - such as grants, prizes, or research fellowships - might address the misalignments identified here. Another is to examine how researcher heterogeneity affects optimal knowledge creation paths and policy interventions. A third is to investigate how competition among researchers shapes knowledge creation dynamics and whether it mitigates or exacerbates the inefficiencies we have identified.

In conclusion, our simplified model of knowledge creation offers both theoretical clarity and practical insights for science policy. By making the complex framework of CS more accessible and deriving precise welfare results, we have contributed to a deeper understanding of how knowledge evolves and how policy can shape that evolution. While we have focused primarily on analytical simplification, our findings have substantial implications for the design of research funding mechanisms and the governance of scientific institutions.

References

- Aghion, P., Dewatripont, M. and Stein, J. C. (2008), Academic freedom, private-sector focus, and the process of innovation, *The RAND Journal of Economics* **39**(3), 617–635.
- Azoulay, P., Fuchs, E., Goldstein, A. P. and Kearney, M. (2019), Funding Breakthrough Research: Promises and Challenges of the “ARPA Model”, *Innovation Policy and the Economy* **19**, 69–96.
- Azoulay, P., Graff Zivin, J. S. and Manso, G. (2011), Incentives and Creativity: Evidence from the Academic Life Sciences, *RAND Journal of Economics* **42**(3), 527–554.
- Bikard, M., Murray, F. and Gans, J. S. (2015), Exploring Trade-offs in the Organization of Scientific Work: Collaboration and Scientific Reward, *Management Science* **61**(7), 1473–1495.
- Bryan, K. A. and Lemus, J. (2017), The Direction of Innovation, *Journal of Economic Theory* **172**, 247–272.
- Callander, S. (2011), Searching and Learning by Trial and Error, *American Economic Review* **101**(6), 2277–2308.
- Callander, S. and Clark, T. S. (2017), Precedent and Doctrine in a Complicated World, *American Political Science Review* **111**(1), 184–203.
- Callander, S. and Hummel, P. (2014), Preemptive Policy Experimentation, *Econometrica* **82**(4), 1509–1528.

- Callander, S., Lambert, N. and Matouschek, N. (2022), Innovation and Competition on a Rugged Technological Landscape, *Management Science* **68**(10), 7413–7431.
- Carnehl, C. and Schneider, J. (2025), A Quest for Knowledge, *Econometrica* **93**(2), 623–659.
- Dasgupta, P. and David, P. A. (1994), Toward a new economics of science, *Research Policy* **23**(5), 487–521.
- Fortunato, S., Bergstrom, C. T., Börner, K., Evans, J. A., Helbing, D., Milojević, S., Petersen, A. M., Radicchi, F., Sinatra, R., Uzzi, B., Vespignani, A., Waltman, L., Wang, D. and Barabási, A.-L. (2018), Science of Science, *Science* **359**(6379), eaao0185.
- Foster, J. G., Rzhetsky, A. and Evans, J. A. (2015), Tradition and Innovation in Scientists’ Research Strategies, *American Sociological Review* **80**(5), 875–908.
- Gans, J. S. and Murray, F. (2014), Credit History: The Changing Nature of Scientific Credit, in A. B. Jaffe and B. F. Jones (eds), *The Changing Frontier: Rethinking Science and Innovation Policy*, University of Chicago Press, Chicago, chapter 4, pp. 107–131.
- Gans, J. S. and Murray, F. E. (2023), Markets for Scientific Attribution, *Journal of Law, Economics, and Organization* **39**(3), 828–846.
- Hill, R. and Stein, C. (2024), Race to the Bottom: Competition and Quality in Science, *Review of Economic Studies* **91**(2), 807–835.
- Hill, R. and Stein, C. (2025), Scooped! Estimating Rewards for Priority in Science, *Journal of Political Economy* **133**(3), 793–845.
- Hopenhayn, H. and Squintani, F. (2021), On the Direction of Innovation, *Journal of Political Economy* **129**(7), 1991–2022.
- Hopenhayn, H., Llobet, G. and Mitchell, M. (2006), Patents and Innovation: Evidence from an Inventor Level Analysis, *Journal of Economic Literature* **44**(3), 795–838.
- Merton, R. K. (1957), Priorities in scientific discovery: A chapter in the sociology of science, *American Sociological Review* **22**(6), 635–659.
- Myers, K. (2020), The Elasticity of Science, *American Economic Journal: Applied Economics* **12**(4), 103–134.
- Scotchmer, S. (1991), Standing on the Shoulders of Giants: Cumulative Research and the Patent Law, *Journal of Economic Perspectives* **5**(1), 29–41.
- Van Noorden, R. (2015), Interdisciplinary Research by the Numbers, *Nature* **525**(7569), 306–307.
- Wang, J., Veugelers, R. and Stephan, P. (2017), The structure of scientific revolutions in the age of big data science and team science, *Journal of Big Data* **4**(1), 1–21.

Appendix

A Formal Proofs

A.1 Proof of Proposition 1

To prove Proposition 1 we need to characterise the researcher’s optimal actions. The proof proceeds in two steps. We first establish the optimal distance for expansion beyond the knowledge frontier. We then determine the preferred location for a deepening move and compare the resulting payoffs to derive the cut-off at which deepening is chosen over expansion.

Optimal distance for expansion. When a researcher expands the frontier to a point at distance $d > 0$ beyond X_0 , the change in the value function is

$$V(d; \infty) = d - \frac{d^2}{6q} + \mathbf{1}_{(d > 4q)} \frac{\sqrt{d}(d - 4q)^{3/2}}{6q}.$$

Carnehl and Schneider (2025) show that the maximiser is $d_E^* = 3q < 4q$, so the indicator term is inactive at the optimum. Restricting attention to $d \leq 4q$ therefore yields $V(d; \infty) = d - \frac{d^2}{6q}$, which is concave and maximised at $d = 3q$. Substituting d_E^* back into $V(d; \infty)$ yields the benefit of expansion, $V(3q; \infty) = 3q/2$.

Preferred location for deepening. When the researcher chooses to deepen knowledge within the existing interval $[0, X_0]$, the optimal location depends on the length of that interval. Carnehl and Schneider’s results show that there is a continuous function $d_0(X_0)$, symmetric around the midpoint, that maximises the value of a deepening move. The key features of $d_0(X_0)$ are as follows:

- For short intervals, $X_0 \leq 4q$, the benefit of deepening takes the simple form $V(d; X_0) = d(X_0 - d)/(3q)$, which is maximised at the midpoint; hence $d_0(X_0) = X_0/2$ and the gain from deepening is $X_0^2/(12q)$.
- For intermediate intervals, $4q < X_0 \leq 6q$, the midpoint remains optimal. The full benefit function includes indicator terms that account for the existence of a dead zone beyond $4q$, but the derivative with respect to d still changes sign only at $d = X_0/2$; see Lemma 2 of Carnehl and Schneider (2025) for details.
- For longer intervals, $6q < X_0 < 8q$, the optimal deepening point moves inwards from the midpoint towards $3q$. There exists $\tilde{X}_0 \in (6q, 8q)$ such that $d_0(X_0) = X_0/2$ for $X_0 \leq \tilde{X}_0$ and $d_0(X_0) \in (3q, X_0/2)$ for $X_0 > \tilde{X}_0$.
- When $X_0 \geq 8q$, the optimal deepening location remains in the range $(3q, 4q)$ and converges to $3q$ from above as X_0 grows. Deepening at exactly $3q$ is never optimal for finite

X_0 because the marginal gain from splitting an interval longer than $8q$ is strictly concave on either side.

These properties imply that the maximised deepening value $M(X_0) = V(d_0(X_0); X_0)$ is continuous, has an interior maximum near $X_0 \approx 6.2q$, and satisfies $M(X_0) > \frac{3q}{2}$ iff $X_0 > \hat{X}^0$; moreover $M(X_0) \downarrow \frac{3q}{2}$ as $X_0 \rightarrow \infty$.

Comparing expansion and deepening. To determine when the researcher prefers expansion over deepening, we compare the payoff from expanding the frontier by $3q$ with the payoff from a deepening move at the optimal interior location $d_0(X_0)$. When $X_0 \leq 4q$ the deepening benefit reduces to $X_0^2/(12q)$; equating this with $3q/2$ leads to

$$3q/2 = X_0^2/(12q) \quad \Leftrightarrow \quad X_0 = \sqrt{18}q \approx 4.243q.$$

However, this calculation ignores the dead-zone correction which becomes relevant once $X_0 > 4q$. Using the full expression for $V(d; X_0)$ (with indicator terms), Carnehl and Schneider (2025, Lemma 9) show that there is a unique value $\hat{X}^0 \approx 4.338q$ such that $V(d_E^*; \infty) = V(d_0(X_0); X_0)$. For $X_0 < \hat{X}^0$, the value of deepening is less than $3q/2$, so expansion is preferred; for $X_0 > \hat{X}^0$ the reverse holds. Since $\hat{X}^0 > 4q$, the midpoint characterisation remains valid at the cut-off.

Collecting these observations, we have established that the researcher expands the knowledge frontier by $3q$ whenever $X_0 < \hat{X}^0 \approx 4.338q$, and deepens at the interior point $d_0(X_0)$ otherwise.

A.2 Proof of Proposition 2

Let X denote the length of the interval in which the period-2 decision maker considers a deepening move. Expansion yields the constant payoff $V(3q; \infty) = 3q/2$. Deepening yields at most

$$M(X) := \max_{d \in [0, X/2]} V(d; X) = V(d_0(X); X).$$

By Carnehl and Schneider (2025, Lemma 9), there exists a unique cutoff $\hat{X}^0 \approx 4.338q$ such that $M(X) < 3q/2$ for $X < \hat{X}^0$ and $M(X) > 3q/2$ for $X > \hat{X}^0$. Hence the period-2 choice is to deepen iff $X > \hat{X}^0$, and to expand iff $X < \hat{X}^0$.

When the relevant interval arises from a first-period deepening that halves an interval of length X , the available length for the second-period deepening comparison is $X/2$. The period-2 decision is then to deepen iff $M(X/2) > 3q/2$, which holds iff $X/2 > \hat{X}^0$, i.e. iff $X > X_2^* := 2\hat{X}^0$. \square

A.3 Proof of Proposition 3

The planner maximises the discounted sum of knowledge values over two periods, choosing between four strategies: expand-expand (EE), expand-deepen (ED), deepen-expand (DE) and

deepen–deepen (DD). Denote the discount factor by $\delta \in (0, 1)$ and the initial knowledge value by $v(\mathcal{F}_0)$. We compare the welfare associated with each strategy and derive cut–offs on X_0 that define the planner’s optimal strategy.

Step 1: EE versus DE. Under EE, the planner expands by $3q$ in both periods. Under DE, the planner deepens the initial interval in period 1 and then expands by $3q$ in period 2. Ignoring the constant $v(\mathcal{F}_0)$, the welfare difference is

$$W_{EE} - W_{DE} = (1 + \delta)\frac{3q}{2} - (1 + \delta)V(d_0(X_0); X_0).$$

When $X_0 \leq 4q$, we have $V(d_0(X_0); X_0) = X_0^2/(12q)$, so the sign of $W_{EE} - W_{DE}$ is positive if and only if $X_0 < \hat{X}^0$, where \hat{X}^0 solves $X_0^2/(12q) = 3q/2$ after adjusting for the dead-zone correction. As argued above, $\hat{X}^0 \approx 4.338q$. For $X_0 > 4q$, we must use the full benefit function for deepening and the difference becomes

$$(1 + \delta)\frac{3q}{2} - (1 + \delta)\left[\frac{X_0^2}{12q} - \frac{\sqrt{X_0}(X_0 - 4q)^{3/2}}{6q}\right].$$

A straightforward numerical calculation shows that this expression remains positive for $X_0 < \hat{X}^0$ and negative thereafter. Hence, EE is preferred to DE if and only if $X_0 < \hat{X}^0$.

Step 2: DE versus DD. In DE, the planner deepens in period 1 and expands in period 2; in DD, the planner deepens in both periods. Write $d_1 = d_0(X_0)$ for the optimal first-period deepening location and let $X_1 = X_0 - d_1$ denote the length of the larger subinterval created by deepening. The welfare difference is

$$W_{DE} - W_{DD} = \delta[V(3q; \infty) - V(d_0(X_1); X_1)].$$

Again, when $X_0 \leq 4q$, deepening splits X_0 into two intervals of length $X_0/2$. A second deepening would split one of these further and yield $X_0^2/(48q)$. Comparing $3q/2$ with $X_0^2/(48q)$ gives the condition $X_0 < 6\sqrt{2}q \approx 8.485q$. Adjusting for dead-zone terms in $V(d; X)$ shifts this threshold upward slightly. Using the full expression for deepening, Carnehl and Schneider’s numerical analysis (Lemma 7 and subsequent discussion) shows that the exact cut–off is $X_2^* \approx 8.676q$. For $X_0 < X_2^*$ the planner prefers DE over DD, whereas for $X_0 > X_2^*$ the planner deepens in both periods.

Step 3: EE versus ED. The ED strategy entails an expansion in period 1 followed by a deepening in period 2. Let d_1 denote the first-period expansion distance and d_2 the second-period deepening location. Anticipating the period–2 deepening, the planner chooses d_1 to maximise the sum of a current expansion (with benefit $V(d_1; \infty)$) and the discounted deepening

benefit $V(d_2; d_1)$. The optimal d_1 satisfies $d_1 = 6q(1 + \delta)/(2 + \delta) > 3q$. If $d_1 \leq 4q$ the benefit of expanding beyond $3q$ is strictly smaller than that of $3q$, so ED is dominated by EE. If $d_1 > 4q$ the first-period expansion immediately leaps over a dead zone, but the cost of doing so outweighs any gain from deepening in period 2. In all cases, $W_{EE} > W_{ED}$.

Step 4: Summary. Combining the above comparisons, we conclude that the planner's optimal strategy depends on the length of the initial knowledge gap:

- If $X_0 < \hat{X}^0 \approx 4.338q$, the planner expands in both periods (EE), as deepening yields less value than expansion in period 1 and repeating the expansion in period 2 maximises value.
- If $\hat{X}^0 < X_0 < X_2^* \approx 8.676q$, the planner deepens in the first period and expands in the second (DE); the initial deepening creates more value than expansion, but the new interval is not so long that deepening again is desirable.
- If $X_0 > X_2^* \approx 8.676q$, the planner deepens in both periods (DD), as the initial gap is sufficiently large that splitting it twice delivers the greatest benefit.

A.4 Proof of Proposition 4

Proposition 4 states that private and social incentives are aligned: the researcher's cut-off for choosing between expansion and deepening coincides with the planner's threshold in both periods. The proof follows directly from the previous results.

Alignment on the extensive margin. Proposition 1 characterises the researcher's expand-vs.-deepen rule in period 1 using the cutoff \hat{X}^0 . Proposition 2 characterises the period-2 expand-vs.-deepen comparison using the same cutoff applied to the relevant interval length. Since expansion yields the constant payoff $V(3q; \infty) = 3q/2$ and the best deepening payoff is $M(X) = \max_d V(d; X)$, the cutoff comparisons do not depend on δ . Hence the researcher and planner use the same thresholds to decide whether to expand or deepen in each period. \square

Alignment in period 2. When period 1 action is *expansion*, the state is $\{0, X_0, X_0 + 3q\}$. Both researcher and planner compare: expand again; deepen in $[0, X_0]$; or deepen in $[X_0, X_0 + 3q]$. Because period 1 expansion implies $X_0 < \hat{X}^0$, we have $V(d_0(X_0); X_0) < \frac{3q}{2}$. For the new interval of length $3q$, Case 1 of (7) yields

$$V\left(\frac{3q}{2}; 3q\right) = \frac{\left(\frac{3q}{2}\right)\left(3q - \frac{3q}{2}\right)}{3q} = \frac{3q}{4} < \frac{3q}{2}.$$

Hence expansion strictly dominates both deepening options in period 2. When period 1 action is *deepening*, the state is $\{0, d_0(X_0), X_0\}$. The period 2 decision uses the same threshold X_2^* : expand if $X_0 < X_2^*$ and deepen again if $X_0 \geq X_2^*$. Thus period 2 rules coincide.

We conclude that, in the benchmark model, private and social incentives coincide on the *extensive margin*: the researcher and planner use the same cutoffs \hat{X}^0 and X_2^* to decide whether to expand or deepen. When the deepening maximiser is unique (e.g. when $d_0(X) = X/2$), the location choice coincides as well.

A.5 Proof of Proposition 5

Proposition 5 describes the behaviour of a myopic researcher when the knowledge gap is large and the universe of questions is divided into multiple disciplines. In this setting the value of deepening depends on the location of the new point relative to both disciplinary boundaries. Carnehl and Schneider (2025) show that there is a threshold $\tilde{X}_0 \in (6q, 8q)$ such that the structure of the value function changes at this point. We summarise their results and use them to characterise the researcher’s choices.

Optimal location in period 1. Let $X_0 > 6q$ denote the length of the initial gap and $d \in (0, X_0)$ the distance from the left boundary to the new point. When $X_0 > 4q$ the full benefit function for deepening is

$$\begin{aligned} V(d; X_0) = & \frac{1}{6q} \left(2X_0\sigma^2(d; X_0) + \mathbf{1}_{\{d>4q\}} \sqrt{d} (d - 4q)^{3/2} \right. \\ & + \mathbf{1}_{\{X_0-d>4q\}} \sqrt{X_0-d} (X_0 - d - 4q)^{3/2} \\ & \left. - \mathbf{1}_{\{X_0>4q\}} \sqrt{X_0} (X_0 - 4q)^{3/2} \right), \end{aligned}$$

where $\sigma^2(d; X_0) = d(X_0 - d)/X_0$ is the posterior variance. For $X_0 \leq \tilde{X}_0$ the value function is quasi-concave and symmetric around $d = X_0/2$, so the researcher’s optimal choice is the midpoint; see Lemma 2 of Carnehl and Schneider (2025). For $X_0 > \tilde{X}_0$ the shape of the value function changes. Lemma 4 of Carnehl and Schneider (2025) establishes that for $X_0 > 8q$ the optimal deepening point lies between $3q$ and $4q$, never at the midpoint. By continuity (Lemma 6 of the same paper) the unique maximiser $d_1^R(X_0)$ moves continuously from $X_0/2$ at $X_0 = \tilde{X}_0$ towards $3q$ as X_0 grows, always remaining in the interval $(3q, X_0/2)$. Moreover, Lemma 7 shows that

$$\lim_{X_0 \rightarrow \infty} d_1^R(X_0) = 3q,$$

so the optimal deepening location converges to the optimal expansion distance as the gap becomes arbitrarily large. These results prove the three statements in Proposition 5 concerning the period–1 decision.

Choosing the second-period action. After period 1 there are two intervals: a smaller one of length d_1^R and a larger one of length $X_0 - d_1^R$. The researcher compares three options: (i) expanding beyond one frontier, generating value $V(3q; \infty) = 3q/2$; (ii) deepening in the small interval, generating $V(d_1^R/2; d_1^R)$; and (iii) deepening in the large interval, generating $V(d_2^R; X_0 - d_1^R)$ where d_2^R is its optimal deepening location. As X_0 grows large, the small interval remains bounded while the large interval approaches length X_0 . The benefit from deepening in the small interval tends to zero, while the benefit from expansion is constant at $3q/2$. The maximised deepening benefit $M(X)$ exceeds $3q/2$ whenever $X > \hat{X}^0$ (and converges back down to $3q/2$ as $X \rightarrow \infty$). Consequently, in period 2 the researcher deepens in the larger subinterval whenever its length exceeds \hat{X}^0 ; otherwise the researcher expands. When deepening in the larger subinterval occurs and that subinterval is long, the optimal location approaches $3q$ from its boundary.

A.6 Proof of Proposition 6

Let the planner choose a first-period deepening location $d_1 \in (0, X_0/2]$. The incremental welfare (relative to the constant baseline $v(\mathcal{F}_0)(1 + \delta)$) can be written as

$$W(d_1) = (1 + \delta)V(d_1; X_0) + \delta V_2(d_1),$$

where the continuation term is the best period-2 payoff available after splitting $[0, X_0]$ into subintervals of lengths d_1 and $X_0 - d_1$:

$$V_2(d_1) = \max \{V(3q; \infty), M(d_1), M(X_0 - d_1)\}, \quad M(X) := \max_{d \in [0, X/2]} V(d; X).$$

Since $V(\cdot; X_0)$ and $M(\cdot)$ are continuous, $W(\cdot)$ is continuous on $[0, X_0/2]$ and a maximiser d_1^P exists. When $\delta \rightarrow 0$, $W(d_1)$ converges uniformly to $V(d_1; X_0)$, so any accumulation point of d_1^P is a maximiser of $V(\cdot; X_0)$, i.e. $d_1^P \rightarrow d_1^R$ when the myopic maximiser is unique.

Finally, as $X_0 \rightarrow \infty$, Carnehl and Schneider (2025) show $M(X) \downarrow 3q/2$ and $V(d; X_0) \rightarrow V(d; \infty)$ for fixed d . Hence for large X_0 the planner's objective is arbitrarily close to $(1 + \delta)V(d; \infty) + \delta \cdot (3q/2)$, which is maximised at $d = 3q$, implying $d_1^P \rightarrow 3q$.

A.7 Proof of Proposition 7

Let $W_P := \max_{d_1 \in [0, X_0/2]} W(d_1)$ denote the planner's maximal welfare, where $W(\cdot)$ is defined as in the proof of Proposition 6. Let d_1^R denote the myopic researcher's first-period deepening location and define $W_R := W(d_1^R)$. Since the planner optimises over the same choice set (and hence can always imitate the researcher), we have $W_P \geq W_R$, so $\Delta W = W_P - W_R \geq 0$, with strict inequality whenever d_1^R is not itself a maximiser of $W(\cdot)$.

The limit $\lim_{\delta \rightarrow 0} \Delta W = 0$ follows because $W(\cdot)$ converges uniformly to $V(\cdot; X_0)$ as $\delta \rightarrow 0$, so the planner's and researcher's objectives coincide. The limit $\lim_{X_0 \rightarrow \infty} \Delta W = 0$ follows from Proposition 6, since both the planner and researcher choose locations converging to $3q$ and continuation values converge to $3q/2$, making the welfare difference vanish.

A.8 Welfare comparisons for $X_0 = 11q$

We compute the entries in Table 1 for a gap length $X_0 = 11q$ and discount factor $\delta = 1$. All values use the full benefits function in (6).

Researcher. The myopic researcher chooses a first-period deepening location d_1^R that maximises $V(d; 11q)$, yielding $d_1^R \approx 3.05q$ and $V_1^R \approx 1.534q$. This creates subintervals of lengths $3.05q$ and $7.95q$. In period 2 the researcher deepens in the longer subinterval, choosing an interior location at distance $\approx 3.28q$ from its boundary, yielding $V_2^R \approx 1.614q$. Total welfare is

$$W_R = 2V_1^R + V_2^R \approx 2(1.534q) + 1.614q = 4.683q.$$

Planner. The planner chooses d_1^P to maximise $(1 + \delta)V(d_1; 11q) + \delta V_2(d_1)$, where $V_2(d_1)$ is the best period-2 payoff given the induced subinterval lengths. For $\delta = 1$, the planner chooses $d_1^P \approx 3.42q$, yielding $V_1^P \approx 1.512q$ and leaving a remaining interval of length $11q - d_1^P \approx 7.58q$. In period 2, the planner deepens at (approximately) the midpoint of this remaining interval, yielding $V_2^P \approx 1.681q$ and producing subintervals $\{3.42q, 3.79q, 3.79q\}$. Total welfare is

$$W_P = 2V_1^P + V_2^P \approx 2(1.512q) + 1.681q = 4.704q.$$

Equal spacing. Equal spacing places discoveries at $11q/3$ and $22q/3$, producing three equal subintervals of length $11q/3 \approx 3.67q$. The first-period payoff is $V_1^E \approx 1.472q$ and the second-period payoff is $V_2^E \approx 1.735q$, so

$$W_E = 2V_1^E + V_2^E \approx 2(1.472q) + 1.735q = 4.679q.$$

Relative performance and activation. Relative to the planner, the researcher achieves about $W_R/W_P - 1 \approx -0.45\%$, and equal spacing achieves $W_E/W_P - 1 \approx -0.52\%$. The induced interval structure under the planner and equal spacing has all subinterval lengths below $4q$ and therefore activates 100% of the gap, while the researcher's structure leaves a dead zone, activating about 83.9% of the gap.