A Model of the Babbage Firm

Joshua S. Gans*

October 28, 2025

Abstract

This paper develops a dynamic model of the cognitive division of labour within firms. A continuum of questions arrive over time, and each must be assigned either to an expert or to a routine worker. Routine workers exploit codified knowledge by interpolating between previously solved problems, while experts possess tacit knowledge and can either address the current question or conduct research to create new knowledge point. The manager chooses whether to use experts to handle current issues (*exploitation*) or to invest in knowledge creation (*exploration*), taking into account the endogeneity of the knowledge base and the evolving organisational structure. The analysis shows that organisations initially expand by adding distant knowledge points, then transition to a refinement phase in which they fill gaps between existing anchors. We derive conditions under which organisational performance exhibits an S-shaped trajectory versus concave growth. Finally, we consider the adoption of artificial intelligence, distinguishing *interpolation AI*, which improves routine workers' precision, from *application AI*, which lowers the cost of expert operations. The two technologies have opposite comparative statics: improving interpolation raises the value of exploration and hastens routinisation, whereas lowering expert costs crowds out research. These results offer new predictions about the allocation of cognitive tasks and the impact of AI on organisational learning. *Journal of Economic Literature* Classification Numbers: D83, J24, L23, L25.

Keywords: division of labour, knowledge codification, Charles Babbage, cognitive processes, organisational dynamics, knowledge work, exploration-exploitation.

*Rotman School of Management, University of Toronto and NBER. Thanks to Refine.ink, Claude Opus 4, Gemini 2.5 Pro and ChatGPT-5-pro for excellent research assistance and seminar participants at the University of Toronto and Philippe Aghion and Ajay Agrawal for useful discussions. Responsibility for all errors remains my own.

1 Introduction

The division of labour is a fundamental organising principle of the modern firm. In knowledge-intensive environments, it extends beyond the physical separation of tasks to encompass differences in cognition: some workers apply existing knowledge to routine problems, while others generate new solutions to novel challenges. This *cognitive hierarchy* is pervasive in professional services, research laboratories and digital platforms. Previous analyses have focused on the broad macro-level choices between exploration versus exploitation, balancing organisational learning with production needs (March, 1991). The canonical model is that of Garicano (2000), who demonstrates how firms endogenously sort problems by difficulty and route them to workers of increasing expertise. This paper builds on that tradition by considering a cognitive division of labour not based on sorting workers into roles modulated by communication frictions, but rather on how problems are handled using distinct types of knowledge processing.[1]

This paper develops a dynamic model that integrates these strands. We characterise an organisation that faces a sequence of questions drawn from a one-dimensional continuum. Each question has a unique correct answer, and proximity along the continuum captures the idea that nearby questions are related. Following Carnehl and Schneider (2025), we model the answer process as a Brownian motion whereby questions close to one another likely have similarly close answers.[2] Two types of workers process these questions. A *routine worker* is a knowledge user: they interpolate between codified points and deliver an action whose expected accuracy falls with the variance of the conjecture. An *expert* is a knowledge creator: they possess tacit expertise that allows them to solve a question perfectly, but this ability is localised and costly. The expert can either operate on the current problem, incurring a quadratic cost increasing in the distance from the nearest codified point, or conduct research to codify a new point for future use. A manager, observing the current question and the existing knowledge base, chooses between exploitation and exploration.

Our first contribution is to derive the static allocation of cognitive labour when the

[1] By contrast, the theory of knowledge-based hierarchies emphasises that organisations must manage the acquisition, communication and use of knowledge because knowledge is an indispensable input embedded in individuals with limited time. Hierarchical structures allow less knowledgeable workers to handle routine tasks while experts direct attention to more complex questions, thereby economising on the time of experts. Mainstream economic models often abstract from this organisational problem, jumping directly to production functions that depend on a fixed set of inputs. Recent research has sought to incorporate organisations into equilibrium frameworks by modelling agents' knowledge, communication and span of control (e.g., Calvo and Wellisz, 1978; Lucas, 1978; Rosen, 1982; Garicano, 2000; Garicano and Rossi-Hansberg, 2004; Garicano and Hubbard, 2006).

[2] This approach represents the end of a line of research into the use of knowledge in decision-making and its application to organisational design and other economic and political phenomena beginning with Callander (2011) and Callander and Clark (2017). See Bardhi and Callander (2026) for a survey.

manager is myopic. We show that the advantage of expert operation relative to routine interpolation is strictly concave between codified points. With a single knowledge anchor, the expert dominates everywhere if their application cost is below a threshold, and we assume that to be the case. With two anchors, the advantage exhibits an "M" shape: the expert is favoured near the anchors and in the tails, while the routine worker covers the midpoint when the gap is sufficiently wide. These comparative statics clarify why cognitive allocations can be fragmented rather than contiguous, in contrast to models of management by exception. They also reveal how the density of the knowledge base affects the tipping point between expertise and routine.

Our second contribution is to embed these static insights into an infinite-horizon dynamic programme. The manager discounts future payoffs and must decide when to invest the expert's time in research. Focussing on a case where experts perform strongly in localised contexts, we show that optimal research is incremental: new anchors are placed at a fixed distance from the current knowledge frontier until the tails are brought within the expert's reach, after which the organisation switches to splitting interior gaps. Using a *stationary value approximation* to keep the dynamic problem tractable, we characterise the evolution of exploration—the probability that the expert engages in research—and derive conditions under which the performance trajectory is S-shaped. When the return to research is high and routine interpolation is sufficiently precise, early expansions of the knowledge base accelerate subsequent learning, and the organisation exhibits an initial phase of accelerating growth followed by deceleration. When the return is low, growth is concave throughout. In mature stages, the growth rate decays and the organisation routinises: gaps shrink, the expert is seldom deployed, and the routine worker handles most questions.

Finally, we connect our framework to the emerging economics of artificial intelligence. Generative AI technologies can both augment experts and empower generalists. We distinguish two polar cases. *Interpolation AI* improves the precision with which routine workers can synthesise codified information, effectively raising the tolerance for variance. *Application AI* reduces the cognitive cost of expert operation by automating tacit reasoning. We show that these technologies have opposite effects on organisational learning. Increasing interpolation quality makes research more attractive, and during the acceleration phase, it steepens the S-curve and shifts the organisation towards codification. Lowering expert costs makes exploitation cheaper and delays routinisation. These results rationalise empirical observations that AI-native start-ups scale rapidly by leveraging broad improvements in generalist productivity, whereas AI tools focused on expert augmentation may entrench reliance on specialists.

The model of the "Babbage Firm" presented here offers a conceptual innovation by shift-

ing attention from skills and task complexity to cognitive processes in the organisation of knowledge work. It also contributes technically by integrating learning and uncertainty into a model of task assignment and knowledge creation. The framework relates to research on specialisation and human capital (e.g., Becker, 1964; Rosen, 1983), on knowledge-based hierarchies and span of control (Garicano, 2000; Garicano and Rossi-Hansberg, 2004; Garicano and Hubbard, 2006; Calvo and Wellisz, 1978; Lucas, 1978; Rosen, 1982), and on optimal task assignment and matching (Roy, 1951; Acemoglu et al., 2007), but it differs by explicitly modelling how information is used in decision-making and how knowledge codification endogenously evolves over time. It also distinguishes itself from "management by exception" results by recognising that knowledge is not ordered by complexity or frequency; rather, experts and routine workers are distinguished by their cognitive technology—creating versus using knowledge. In sum, the paper offers a unified framework in which the location of knowledge points, the variance of routine interpolation, and the dynamic creation of new knowledge jointly determine the division of cognitive labour.

The remainder of the paper is organised as follows. Section 2 formalises the environment and the cognitive technologies. Section 3 studies the static allocation of tasks and the dynamics of knowledge accumulation, first in a two-period setting and then in the infinite horizon. Section 4 analyses the shape of the organisational learning curve and identifies regimes of concave, linear and S-shaped growth. Section 5 investigates the adoption of artificial intelligence and derives comparative statics with respect to interpolation and application technologies. Section 6 concludes and outlines directions for future research.

2 Model Setup

The baseline model is dynamic, handling incoming problems with different knowledge worker types, who may be assigned roles focusing on the organisation's short-term versus long-term investments. The goal is to characterise the cognitive division of labour and its drivers, as well as how these evolve over the lifecycle of the organisation that is endogenous to the model.

2.1 Questions, decisions and knowledge

There are an infinite number of discrete time periods. In each period, the organisation faces a decision in response to a question indexed by a real number x. Following the framework of Carnehl and Schneider (2025), a decision is the choice of an action whose correctness depends on the underlying state. Formally, for each question x the organisation selects an

action $a \in \mathbb{R}$. There is a unique correct answer $y(x)$ prescribing the action that maximises the payoff. The payoff function is normalised so that

$$u(a,x) = \begin{cases} 1, & \text{if } a = y(x), \\ 0, & \text{otherwise.} \end{cases}$$

Thus, choosing the correct action yields value one, while choosing any other action or abstaining yields value zero.[3] Throughout, if a worker is not asked to actively work on knowledge, they perform an alternative task whose value is normalised to zero.

Questions are drawn (uniformly) from a one-dimensional continuum, $x \in [x_L, x_H]$, and proximity along this continuum captures the idea that nearby questions tend to have related answers. At time t, the organisation has a finite set of codified question–answer pairs $\mathcal{F}_t = \{(x_i, y(x_i))\}_{i=1}^{k}$ with $x_1 < x_2 < \cdots < x_k$. When a question x lies between two adjacent codified points, the best conjecture based on existing knowledge is obtained by linearly interpolating between the two known answers. Conjectures are *normally distributed* with mean

$$\mu_x(\mathcal{F}_t) = y(x_i) + \frac{x - x_i}{x_{i+1} - x_i}\big(y(x_{i+1}) - y(x_i)\big)$$

and variance

$$\sigma_x^2(\mathcal{F}_t) = \frac{(x_{i+1} - x)(x - x_i)}{x_{i+1} - x_i}\,.$$

These expressions follow from modelling the answer process $y(\cdot)$ as a *Brownian motion* indexed by the question domain. A Brownian motion is a continuous stochastic process with stationary, independent increments and normally distributed increments whose variance equals the increment length. In the present context, this means that answers to questions further apart evolve like a random walk: the change in the answer between x_i and x_{i+1} is normally distributed with mean zero and variance $(x_{i+1} - x_i)$. Conditional on the values at two points, the process between them forms a *Brownian bridge*, yielding the linear conditional mean and the concave conditional variance above. In particular, uncertainty is highest at the midpoint and falls to zero at the known endpoints. Outside the interval $[x_1, x_k]$, the variance increases linearly with distance from the nearest known point:

$$\sigma_x^2(\mathcal{F}_t) = \begin{cases} x_1 - x, & x < x_1, \\ x - x_k, & x > x_k. \end{cases}$$

[3] An equivalent continuous specification that underlies the value functions below assumes quadratic loss $u(a,x) = 1 - \alpha\,(a - y(x))^2$ for $\alpha > 0$. Both formulations lead to the same linear value functions for routine workers, defined below.

The organisation has a tolerance for imprecision represented by $q > 0$. If the variance associated with a conjecture exceeds q then the probability of correctly identifying the right action is so low that deferral is optimal. This threshold distinguishes routine problems, for which extrapolation yields useful guidance, from frontier problems that require the creation of new knowledge. While it is the case that if $x_{i+1} - x_i \leq 4q$, routine conjectures can cover the interior region, we do not impose this assumption.

2.2 Workers, actions and cost structures

Time is discrete and indexed by $t = 1, 2, \ldots$. In each period, a single question $x(t) \in [x_L, x_H]$ arrives. Two workers are available: a routine worker (R) and an expert (E). The manager discounts future payoffs at rate $\delta \in [0, 1]$.

The two types of workers reflect fundamentally different ways of processing information. A routine worker is a *knowledge user*. Routine workers exploit codified, commonly available knowledge within the organisation (that is, $\mathcal{F}_t$) and form conjectures by interpolating between known points. They apply this public knowledge to high-frequency tasks, but they do not attempt to generate new insights beyond what is already codified. In contrast, the expert corresponds to a *knowledge creator*. The expert can do this in two ways. First, they can engage in *research* that creates new knowledge points (i.e., exploration), codified for use later. Second, they can be *operators* and directly make decisions (i.e., exploitation).

The *routine worker* applies codified knowledge to choose an action for the current question. Given a conjecture distribution with variance $\sigma_x^2(\mathcal{F}_t)$, the expected payoff from the routine worker's decision equals the probability of selecting the correct action. Given the normal distribution, this probability is linear in the variance up to the threshold q. Thus, the value created by the routine worker is

$$V_R(x) = \begin{cases} 1 - \frac{\sigma_x^2(\mathcal{F}_t)}{q}, & \sigma_x^2(\mathcal{F}_t) \leq q, \\ 0, & \sigma_x^2(\mathcal{F}_t) > q. \end{cases}$$

If the routine worker is not assigned to a question, they undertake a different task with a payoff normalised to 0. The formal decision rule for the routine worker is to set the action equal to the mean of the conjecture, $a_R(x) = \mu_x(\mathcal{F}_t)$. Because the conjectured answer is normally distributed around $y(x)$ with variance $\sigma_x^2(\mathcal{F}_t)$, the probability that $a_R(x)$ coincides with the correct answer is $1 - \sigma_x^2(\mathcal{F}_t)/q$ when $\sigma_x^2(\mathcal{F}_t) \leq q$ and zero otherwise. The value $V_R(x)$ above is exactly this expected payoff.

The *expert* can either decide on the current question or engage in research. When they

are operators, experts rely on tacit knowledge acquired through personal experience and judgment. In the spirit of Polanyi's paradox—summarised by Autor (2014) as "we know more than we can tell"—tacit expertise cannot be codified or easily communicated. As a result, an expert's operational ability to deliver answers is localised and does not spill over to other workers. Building tacit knowledge requires costly experience, meaning that an expert and a routine worker are exogenously fixed in their knowledge processing capabilities. This distinction between knowledge users and knowledge creators underpins the different value and cost structures set out below. When the expert acts as an operator, they anchor at the nearest codified point x_i and deploy tacit reasoning to select the correct action with certainty. This reasoning incurs a cognitive cost that grows with the distance between the question and the anchor.

Operator experts, like routine workers, directly answer questions. Specifically, if the question is x and the nearest codified point is x_i, the cost of operating is

$$C_E(x) = \eta \left(x - x_i\right)^2,$$

where $\eta > 0$ captures the breadth of an expert's expertise. The net value of expert decision-making is therefore

$$V_E(x) = 1 - C_E(x) = 1 - \eta \left(x - x_i\right)^2.$$

This cost structure reflects that making perfectly accurate decisions becomes disproportionately harder as the question moves further from known territory.

A natural question is why an expert would not augment tacit reasoning with codified knowledge beyond the nearest reference point. In the present framework, the expert pays a cognitive (or effort) cost to construct a perfectly accurate answer around a focal question. Once this tacit knowledge is created, consulting public knowledge cannot improve an answer that is already exact. Indeed, tacit expertise is inherently localised and personal—it is built through experience, difficult to articulate and cannot be seamlessly combined with codified information. This assumption captures the idea that experts "know more than they can tell," as noted above, and it clarifies why the expert's performance depends solely on the distance from the nearest tacit anchor rather than on other codified points. The expert's decision rule is simply to choose the action equal to the correct answer, $a_E(x) = y(x)$. Because the expert knows or discovers the correct answer, their payoff from making a decision is 1 minus the cognitive cost $C_E(x)$.

Alternatively, in a period, the expert can undertake *research* to expand the set of codified points. The expert's role in any given period is chosen by an agent that I will refer to as "the manager" who optimally and centrally decides on the allocation of agents and also

dictates research direction if it is to be carried out. Thus, if an expert is allocated to research, the manager chooses a target location $x_{\text{new}} \in [x_L, x_H]$. It is assumed here that the only cost of research is the opportunity cost (if any) of having the expert not engage in an operational role.[4] With a lag of one period, research produces new knowledge and the new pair $(x_{\text{new}}, y(x_{\text{new}}))$ is added to $\mathcal{F}_{t+1}$, thereby reducing future variance.

2.3 Managerial decision rule

The manager observes the current question $x(t)$ and the existing knowledge base $\mathcal{F}_t$ and decides whether to allocate the expert to exploitation or exploration. If the manager assigns $x(t)$ to the expert, the routine worker performs their outside task (value 0); if the manager assigns $x(t)$ to the routine worker, the expert engages in research. Let $W(\mathcal{F}_t)$ denote the expected continuation value of the organisational outcome given the knowledge set $\mathcal{F}_t$. The expert's immediate advantage in making a decision at $x(t)$ is

$$A(x(t) \mid \mathcal{F}_t) = V_E(x(t)) - V_R(x(t)).$$

The expected net present value of the best research project is summarised by the *organisational hurdle rate*

$$H^*(\mathcal{F}_t) = \max_{x_{\text{new}}} \Big\{ \delta \left[W(\mathcal{F}_t \cup \{(x_{\text{new}}, y(x_{\text{new}}))\}) - W(\mathcal{F}_t) \right] \Big\}.$$

This hurdle rate measures the increment in the net present discounted value (NPV) generated by adding a knowledge point in the current period.

Let $V(\mathcal{F})$ be the per-period value, $W(\mathcal{F})$ be the continuation value and $\rho(\mathcal{F})$ be the probability of exploration when knowledge is $\mathcal{F}$. The Bellman equation is

$$W(\mathcal{F}) = V(\mathcal{F}) + \delta \big[\rho(\mathcal{F})\, W(\mathcal{F}') + (1 - \rho(\mathcal{F}))\, W(\mathcal{F}) \big],$$

[4] Carnehl and Schneider (2025) model knowledge creation in a similar way in terms of a directed project, but they also include specific costs associated with research. They consider existing knowledge points as bearing (negatively) on that cost (research is more demanding when the target location is far from existing knowledge because uncertainty is higher), while research outcomes are uncertain, with more certainty requiring exponentially greater search effort. This is captured by the convex function $\tilde{c}(\rho) = [\text{erf}^{-1}(\rho)]^2$, where erf^{-1} is the inverse of the error function. Their total cost of research at (x_{new}, ρ) is

$$C_{\text{res}}(x_{\text{new}}, \rho) = \kappa\, \tilde{c}(\rho)\, \sigma^2_{x_{\text{new}}}(\mathcal{F}_t),$$

where $\kappa > 0$ scales the overall difficulty of creating knowledge. As Gans (2025b) argues, many of the key results can be captured without the additional complexity of research costs, and that approach is followed here.

so $(1-\delta)W(\mathcal{F}) = V(\mathcal{F}) + \delta\,\rho(\mathcal{F})\,\Delta W(\mathcal{F})$, where $\Delta W(\mathcal{F}) = W(\mathcal{F}') - W(\mathcal{F})$. The exact hurdle is $H^*(\mathcal{F}) = \delta\,\Delta W(\mathcal{F})$. Calculating $W(\mathcal{F}_t)$ requires solving the full dynamic programming problem (the Bellman equation), which captures how today's research affects not only immediate payoffs but also the value of future research opportunities. As the knowledge base evolves non-stationarily, it is not possible to solve this exactly. We will address this complexity and the necessary analytical approach in Section 3.3.[5]

Thus, the managerial decision rule involves comparing the current advantage of the expert in handling a question to the hurdle rate generated by their engaging in exploration in the current period instead. If $A(x(t) \mid \mathcal{F}_t) > H^*(\mathcal{F}_t)$, the manager exploits by assigning the expert to decide on the current question; if $A(x(t) \mid \mathcal{F}_t) \leq H^*(\mathcal{F}_t)$, the manager explores by having the routine worker decide while the expert conducts research.[6] This rule balances the immediate gain from accurate decision making against the value of expanding the knowledge base in that period (as measured by bringing the marginal benefit of additional knowledge on the net present value of the organisation into the present period), and it fully determines the dynamic allocation of cognitive labour within the firm. The probability that the expert engages in research is $\rho(\mathcal{F}_t) = P(A(x|\mathcal{F}_t) \leq H^*(\mathcal{F}_t))$.

3 The Dynamic Evolution of Knowledge

Having described the building blocks of the model, we now turn to the dynamic allocation of tasks within the firm. This section analyses how the roles of routine workers and experts evolve over time as the knowledge base expands. Recall that when an expert operates on a question, the answer is delivered perfectly, but no new knowledge is codified, whereas when the expert conducts research, the outcome may be the addition of a new codified point that reduces future uncertainty for routine workers.

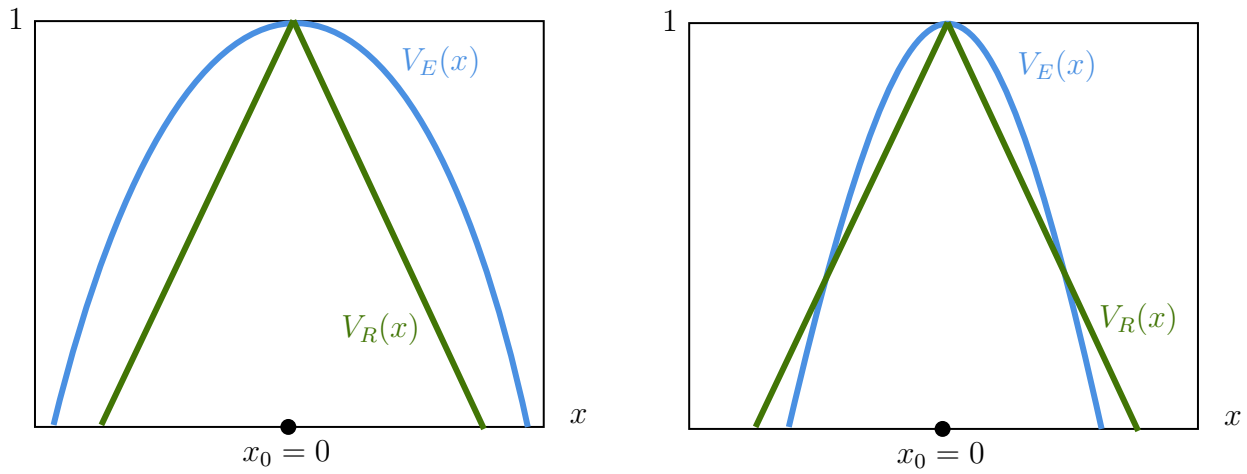
To analyse the dynamics of organisational learning from inception, we focus on a standardised initial condition: a single codified knowledge point x_0 located at the midpoint of the domain $[x_L, x_H]$. We assume a uniform distribution of questions $x(t)$. To keep notation simple, we normalise the domain to $[-L/2, L/2]$, so $x_0 = 0$ for the remainder of the paper.

3.1 The Myopic Case ($\delta = 0$)

A simple case of interest arises when the discount factor δ is zero. The manager values only the immediate payoff and assigns no weight to the future. Because research yields benefits

[5] We therefore use the Stationary Value Approximation (SVA), yielding $H_t = \frac{\delta}{1-\delta}\frac{\Delta^*(t)}{L}$ and the bias signs reported later. (Details and proof of bias signs are in Appendix E).

[6] Note that in the case of indifference it is assumed here that the manager chooses the routine worker.



(a) Low $\eta \leq 1/q^2$ (ED)

(b) High $\eta > 1/q^2$ (MA)

Figure 1: Myopic Allocation with a Single Knowledge Point.

only in future periods, the organisational hurdle rate is $H^*(\mathcal{F}_t) = 0$ for all $\mathcal{F}_t$.

The dynamic decision rule reduces to the static efficiency condition: the expert handles the question $x(t)$ if and only if their immediate advantage over the routine worker is positive, $A(x(t) \mid \mathcal{F}_t) > 0$. Otherwise, the routine worker handles the task. In this myopic scenario, the knowledge base $\mathcal{F}_t$ remains constant, and the organisational structure is fixed over time.

3.1.1 Single Knowledge Point

Under myopia, an expert does not engage in research. Recall that there is a single knowledge point, $x_0 = 0$. The distance of another question, x, from that point is $s = |x|$. Thus, $\sigma^2(x) = s$ (extrapolation). In the region where the routine worker is viable ($s < q$), the advantage is:

$$A(s) = V_E(s) - V_R(s) = (1 - \eta s^2) - (1 - s/q) = s/q - \eta s^2.$$

The routine worker is utilised if $A(s) \leq 0$, which occurs if $s \geq 1/(\eta q)$. For this region to exist within the routine worker's scope ($s < q$), we require $1/(\eta q) < q$, or $\eta > 1/q^2$. Figure 1 illustrates these different outcomes with an expert dominance (or ED) outcome where $A(s) \geq 0$ for all s and a mixed allocation (or MA) where $A(s) < 0$ for some s.

To focus the dynamic analysis on the evolution of an organisation starting from a state

entirely reliant on expertise, we introduce the following condition.

Assumption 1 (Expert Dominance)**.** *We assume the expert's application cost is sufficiently low relative to the routine worker's precision:* $\eta \leq 1/q^2$.

This is a natural assumption that implies that the expert is productively superior to a routine worker when answering questions based on a single knowledge point.[7] Under Assumption 1, $A(s) = s/q - \eta s^2 > 0$ on $s \in [0, \min\{q,\, X_E\}]$, where $X_E = 1/\sqrt{\eta}$. Hence, the expert handles all *served* questions in the myopic single-anchor case. For $s > X_E$, both $V_E(s)$ and $V_R(s)$ are zero, so far tails are unserved.

3.1.2 Two Knowledge Points

While it does not arise in the myopic case, as this situation will play an important role below, it is instructive at this point to consider the allocation of worker types when there are two knowledge points such that $\mathcal{F} = \{x_0, x_1\}$ with interior gap $X = x_1 - x_0$ and midpoint $x_m = (x_0 + x_1)/2$. $A(x(t) \mid \{x_0, x_1\})$ possesses a distinct structure resulting from the different cognitive technologies of the agents. It will be assumed here that $X \leq 4q$ so that $V_R(x_m) > 0$.[8]

1. **Concavity:** The expert's value as an operator $V_E(x)$ is strictly concave due to the convex costs of applying tacit knowledge ($\eta > 0$). The routine worker's value $V_R(x)$ is derived from the variance of the Brownian motion. In the tails (extrapolation), the variance is linear. In the interior (interpolation), the variance of the Brownian bridge is strictly concave. Since $V_R(x)$ is 1 minus the scaled variance (where positive), $V_R(x)$ is strictly convex over the interior gap X. Therefore, $A(x) = V_E(x) - V_R(x)$ is strictly concave within each region of the knowledge space..

2. **Local Minima at Anchors:** At the knowledge points x_i, the routine worker has perfect precision ($V_R(x_i) = 1$), matching the expert. Thus, $A(x_i) = 0$. Analysing the derivatives shows the slope of $A(x)$ switches from negative to positive at these points, confirming they are local minima.

3. **The "M" Shape (Interior Peaks and Midpoint Minimum):** As we move from an anchor (e.g., x_0) towards the midpoint x_m, $A(x)$ initially increases. To understand

[7] Relaxing this assumption adds undue complexity in what follows and does not foreclose that the routine worker is productively superior when interpolating more than a single knowledge point.

[8] As Carnehl and Schneider (2025) show, when $X > 4q$, at the mid-point, $x_m = \frac{1}{2}(x_1 + x_0)$, $V_R(x_m)$ would be negative if the decision-maker relied on that the resulting conjecture for questions at that point. Thus, it is better for no action to be taken. In addition, as will be down below, when additional knowledge points are added, the distance between one knowledge point and its nearest neighbour will imply that $X < 4q$.

the behaviour at the midpoint, we examine the derivatives. The routine worker's value is locally flat ($V_R'(x_m) = 0$). The expert anchors at the nearest point, so $V_E(x)$ has a kink at x_m. The left derivative of $A(x)$ at x_m is $A'(x_m)^- = -\eta X$. The right derivative of $A(x)$ at x_m is $A'(x_m)^+ = \eta X$. Since the derivative switches from strictly negative to strictly positive, x_m is a local minimum (a downward kink) for $\eta > 0$.

4. **Two Tail Peaks:** This structure—concavity on segments, combined with three local minima at x_0, x_m, x_1—implies the existence of four local maxima: one in each tail (Tail Peaks) and one strictly between each anchor and the midpoint (Interior Peaks).

The organisational landscape thus features the "M" shape shown in Figure 2.

With $\delta = 0$, the allocation threshold is $H^* = 0$, the expert handles questions where $A(x) > 0$, and the routine worker handles questions where $A(x) \leq 0$. The resulting division of labour depends critically on the expert's application cost parameter, η. The allocation balances the routine worker's uncertainty ($\sigma^2(x)/q$) against the expert's cost of adaptation ($\eta(x - x_i)^2$).

The structure of $A(x)$ implies that the allocation is generally fragmented. However, the specific pattern depends on whether $A(x)$ remains positive across the domain.

1. **The Midpoint Condition.** The sign of the advantage at the midpoint x_m determines whether the interior is fully covered by the expert. A direct calculation shows (assuming $X \leq 4q$):

$$A(x_m) = V_E(x_m) - V_R(x_m) = 1 - \eta\frac{X^2}{4} - \left(1 - \frac{X}{4q}\right) = \frac{X}{4}\left(\frac{1}{q} - \eta X\right).$$

   $A(x_m) > 0$ if and only if the gap is relatively small compared to the costs: $X < 1/(\eta q)$.

2. **The Tail Condition.** In the extreme tails (Regions 1 and 5) where $V_R(x) = 0$, the expert operates only if $V_E(x) > 0$. This requires the distance to the anchor to be less than the expert's maximum reach, $1/\sqrt{\eta}$.

These conditions define distinct organisational regimes:

Case A: Broad Expertise (Low η). When the cost of applying tacit knowledge is low, the expert's advantage is broad. If $X < 1/(\eta q)$, $A(x_m) > 0$. Note that if it were the case that $H^* > 0$, the expert would cover the tails and the entire interior, except for small regions immediately surrounding the anchors x_i where $A(x) \leq H^*$.

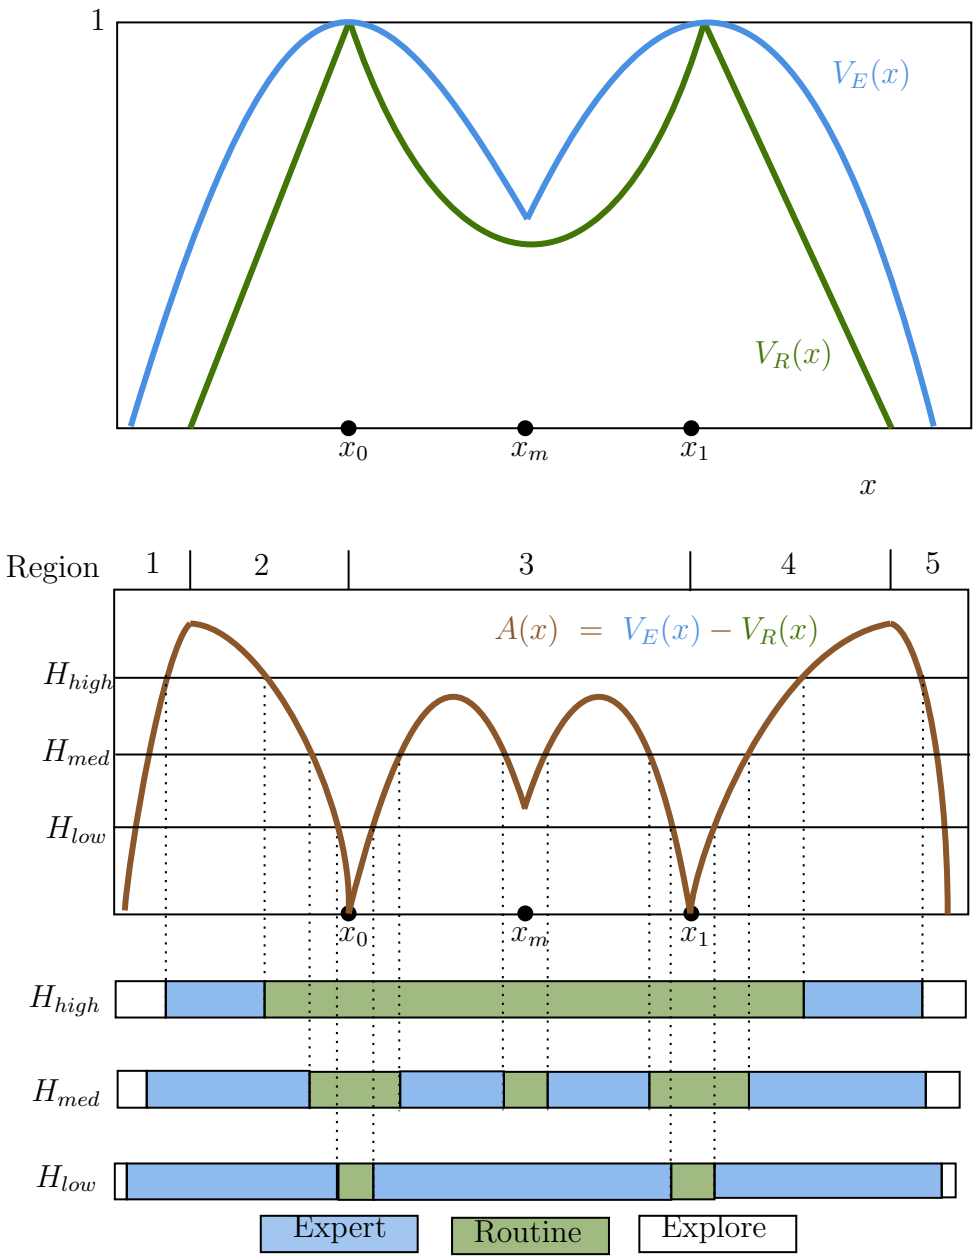


Figure 2: Myopic Allocation with Two Knowledge Points (low X). Routine allocations also result in exploration.

Case B: Localised Expertise (High η). When the cost of applying tacit knowledge is high, the expert's advantage is sharply localised.

- **Interior Fragmentation:** If $X \geq 1/(\eta q)$, $A(x_m) \leq 0$. Again, if $H^* > 0$, the routine worker handles the midpoint in addition to the anchors. The expert focuses only on the localised interior peaks.
- **Limited reach and unserved tails.** The expert's reach is $X_E = 1/\sqrt{\eta}$. In the tails, $V_R(x) = 0$ and $V_E(x) > 0$ only for distances $< X_E$. Beyond X_E, both agents yield zero value, so those questions are unserved. This motivates the expansion phase that shortens tails.

With localised expertise, the allocation becomes highly fragmented.

This analysis highlights that the cognitive division of labour is highly contingent on the nature of expertise. High costs of application (η) lead to complex, fragmented organisations where expertise is deployed selectively to localised pockets of high advantage, while routine processes cover the intervening gaps. This contrasts with models based on "management by exception" (e.g., Garicano, 2000), where allocation is typically contiguous based on difficulty.

The cognitive division of labour is highly sensitive to the structure of the knowledge base $\mathcal{F}$. As the density of codified knowledge increases (e.g., the gap X decreases), the variance $\sigma^2(x)$ falls, improving the routine worker's performance $V_R(x)$. As X falls, at its mid-point, the expert is closer to an existing knowledge point which improves $V_E(x_m)$. Note therefore, that $A(x_m) > 0$ if and only if $X < \frac{1}{\eta q}$ (i.e., in Case A) while $\frac{\partial A(x_m)}{\partial X} > 0$ if and only if $X < \frac{1}{2\eta q}$ which never holds in Case B. Thus, whether, as the knowledge base densifies, the organisation shifts to relying on routine workers depends critically on whether the expert has broad expertise or localised expertise. As will be shown below, eventually, X becomes sufficiently low for any two knowledge points, $\frac{\partial A(x_m)}{\partial X} > 0$, such that the absolute advantage to the expert contracts as knowledge becomes denser. This formalises the process of *routinisation*: tasks previously requiring expertise become standardised procedures.

3.2 Two-Period Case

When $\delta > 0$, the manager will want to consider the balance of exploitation and exploration when allocating the expert to handle questions or not. To build intuition, here it is supposed that there are only two periods ($t = 1, 2$). As before, the domain of questions is the interval $[-L/2, L/2]$ and, at the outset, the organisation holds a single codified anchor at 0, denoted $\mathcal{F}_1 = \{0\}$. In period 1 the manager observes a question $x(1)$ and either directs the expert to answer it (exploitation) or instructs the expert to undertake research (exploration) and

create a new knowledge point at distance $s > 0$ from 0. In period 2, the organisation uses the updated knowledge set $\mathcal{F}_2 = \{0, s\}$ and allocates routine and expert effort optimally. The discount factor between the two periods is $\delta > 0$. We solve the problem by backward induction.

In what follows, we focus on the case $\frac{1}{4q^2} \leq \eta$; which implies that expertise is sufficiently *localised* to give a non-trivial role to routine workers. Combined with the earlier condition on expert dominance (for a single point), this implies that we explore situations where $\eta \in [\frac{1}{4q^2}, \frac{1}{q^2}]$ as it generates a situation where experts initially dominate but can be superseded by routine workers near the centre, while constraining the second anchor to remain close to the original point. As we shall see, this simplifies the analysis in the infinite horizon and drives an outcome where organisations grow by expanding from what they know rather than making large leaps into unknown domains.[9]

Static allocation. Working backwards, consider the second period after an anchor at $s > 0$ has been created. Points on $[0, s]$ are characterised by their distance $\min\{x, s - x\}$ from the nearest anchor, 0 or s. A routine worker interpolates these points using a Brownian bridge: its conditional variance is $\sigma^2(x) = x(s - x)/s$ and the associated value is $V_R(x) = \max\{1 - \sigma^2(x)/q, 0\}$. The expert's value is $V_E(x) = \max\{1 - \eta(\min\{x, s - x\})^2, 0\}$. Recall from our above analysis of two knowledge points that $A(x)$ has a local minimum at $s/2$ over $x \in (0, s)$. At this point,

$$A(s/2) = V_E(s/2) - V_R(s/2) = \left(\frac{1}{qs} - \eta\right)\frac{s^2}{4}.$$

Let $X_c \equiv 1/(\eta q)$ be the *critical gap size*: if $s \leq X_c$ then $A(x) \geq 0$ for all $x \in (0, s/2]$ and the expert is used everywhere; i.e., the domain is *expert-dominated* (ED). When $s > X_c$, there is a central region where $A(\tilde{x}) < 0$ and routine work is preferred, i.e., the domain has a *mixed allocation* (MA). The extent of expert reach is $X_E = 1/\sqrt{\eta}$; for localised expertise we have $X_E \leq 2q$ because $\eta \geq 1/(4q^2)$.

Two further thresholds determine the allocation. The *routinisation limit* $X_R = 4q$ is the largest gap length for which routine interpolation delivers positive value at the midpoint; for $s > X_R$ the centre would be unserved by both workers. The *expert coverage limit* is $S_E = 2X_E = 2/\sqrt{\eta} \leq X_R$. Consequently, for all admissible s, the centre of the gap remains covered by the routine worker. The static allocation on $[0, s]$ therefore falls into two regimes:

(i) For $s \leq X_c$ (expert-dominated), the expert covers the entire gap and $C_{\text{Gap}}(s) =$

[9]The impact of relaxing this assumption is discussed in Appendix B.

$\eta s^3/12$.[10]

(ii) For $X_c < s \leq X_R$ (mixed allocation), the expert serves near the anchors while the routine worker interpolates the centre. The integrated cost is (derived in Appendix A):

$$C_{\text{Gap}}(s) = \frac{s^2}{6q}\left[1 - \frac{2}{(1+\eta q s)^2}\right].$$

No central unserved region occurs in this parameter range because $S_E \leq X_R$ and the optimum will always choose $s \leq X_R$ (see Proposition 1 below). In summary, we have the following thresholds of interest:

$$X_E = \frac{1}{\sqrt{\eta}} \quad \text{(expert reach)} < 2q, \qquad X_c = \frac{1}{\eta q} \quad \text{(ED/MA)} < 4q, \qquad X_R = 4q \quad \text{(routinisation limit)}.$$

where the inequalities are implied by the localised expertise condition.

Second–period optimisation. We now consider the optimal choice of new knowledge point in the second period. Let $C_{\text{Tail}}(Z)$ denote the cost of a tail of length Z attached to a single anchor. Under Assumption 1 ($\eta \leq 1/q^2$), the routine worker never strictly dominates the expert in extrapolation segments. A simple integration shows that:

$$C_{\text{Tail}}(Z) = \int_0^Z \min\{\eta z^2, 1\}\, dz = \begin{cases} \frac{\eta Z^3}{3}, & 0 \leq Z \leq \frac{1}{\sqrt{\eta}}, \\ Z - \frac{2}{3\sqrt{\eta}}, & Z > \frac{1}{\sqrt{\eta}}. \end{cases}$$

Given s, the second-period total cost is $C_2(s) = C_{\text{Gap}}(s) + C_{\text{Tail}}\big(L/2 - s\big) + C_{\text{Tail}}\big(L/2\big)$; the final term does not depend on s can be ignored for the purposes of finding the optimal location of s. Because $C_{\text{Gap}}(\cdot)$ and $C_{\text{Tail}}(\cdot)$ are convex, their sum $f(s) = C_{\text{Gap}}(s) + C_{\text{Tail}}\big(L/2 - s\big)$ is convex. The manager chooses s to minimise f.

Proposition 1 (Optimal second anchor under localised expertise)**.** *Suppose* $\frac{1}{4q^2} \leq \eta \leq \frac{1}{q^2}$. *Let* $X_c = \frac{1}{\eta q}$, $X_R = 4q$, *and let* s^* *be the minimizer of* $f(s)$.

1. *If* $L \leq 3X_c$, *then* $s^* = L/3$ *and both the gap and the tails are expert-dominated.*
2. *If* $3X_c < L \leq 3X_R$, *then* $s^* \in (X_c,\, X_R]$ *is the unique solution to* $C'_{\text{Gap}}(s) = C'_{\text{Tail}}(L/2 - s)$. *This interior optimum lies in the mixed-allocation regime.*

[10] We ignore here the non-generic routine coverage at 0 and s. We break indifference in favor of exploration when $A(x) = H$, so isolated points with $V_R = V_E$ do not affect the integrals.

3. *If $L > 3X_R$, then $C'_{\text{Tail}}(Z) = 1$ for all $Z \geq L/2 - X_R > X_R/2 = 2q \geq X_E \equiv 1/\sqrt{\eta}$. Since $C'_{\text{Gap}}(s)$ is strictly increasing and $C'_{\text{Gap}}(X_R) \geq 1$, the optimum is $s^* = \bar{s}$, where $\bar{s} \in (0, X_R]$ is the unique solution to $C'_{\text{Gap}}(\bar{s}) = 1$.*

In particular, $s^* \leq X_R = 4q$ for all admissible L; the manager never chooses an additional anchor beyond the routinisation limit.

Proof. The objective is to minimise the convex function $f(s) = C_{\text{Gap}}(s) + C_{\text{Tail}}\big(L/2 - s\big)$ over $s \in [0, L/2]$. The optimum s^* satisfies the first-order condition (FOC) $C'_{\text{Gap}}(s) = C'_{\text{Tail}}(L/2 - s)$, provided an interior solution exists. Let $Z = L/2 - s$.

Part 1 ($L \leq 3X_c$). We look for a solution where the gap is Expert-Dominated (ED), $s \leq X_c$, so $C'_{\text{Gap}}(s) = \eta s^2/4$. We also check if the tail is ED, $Z \leq X_E$, where $C'_{\text{Tail}}(Z) = \eta Z^2$. The FOC becomes $\eta s^2/4 = \eta(L/2 - s)^2$. This yields $s/2 = L/2 - s$, so $s^* = L/3$. We must verify the regime assumptions. Since $L \leq 3X_c$, $s^* = L/3 \leq X_c$, confirming the gap is ED. The tail length is $Z^* = L/2 - L/3 = L/6$. We must verify $Z^* \leq X_E$. Since $L \leq 3X_c$, $Z^* \leq X_c/2$. We verify that $X_c/2 \leq X_E$ under localised expertise. This inequality is equivalent to $1/(2\eta q) \leq 1/\sqrt{\eta}$, or $1 \leq 4\eta q^2$. This holds because $\eta \geq 1/(4q^2)$. Thus, the tail is ED, validating the solution $s^* = L/3$.

Part 2 ($3X_c < L \leq 3X_R$). Since $L > 3X_c$, the optimum s^* must be greater than X_c (otherwise we would be in the regime of Part 1). The gap is in the Mixed Allocation (MA) regime. $C_{\text{Gap}}(s)$ is strictly convex, so $C'_{\text{Gap}}(s)$ is strictly increasing. $C'_{\text{Tail}}(Z)$ is non-increasing in s and lies in $(0, 1]$. Therefore, there is a unique solution s^* to the FOC. The bound $s^* \leq X_R$ is established in Part 3.

Part 3 ($L > 3X_R$). We seek the optimum s^* in the domain $[0, L/2]$. Since $L > 3X_R = 12q$, we have $L/2 > 6q$. Consider the range $s \in (0, X_R]$, where $X_R = 4q$. The tail length is $Z = L/2 - s$. In this range, $Z > 6q - 4q = 2q$. Under localised expertise ($\eta \geq 1/(4q^2)$), the expert reach $X_E = 1/\sqrt{\eta} \leq 2q$. Thus, $Z > X_E$, implying the tail is marginally unserved: $C'_{\text{Tail}}(Z) = 1$.

The FOC simplifies to $C'_{\text{Gap}}(s) = 1$. We must show a solution exists in $(0, X_R]$. For $X_c < s \leq X_R$, the derivative is

$$C'_{\text{Gap}}(s) = \frac{s}{3q}\Big[1 - \frac{2}{(1 + \eta q s)^3}\Big].$$

$C'_{\text{Gap}}(s)$ is continuous and strictly increasing on $(0, 4q]$. We evaluate it at the boundaries. $C'_{\text{Gap}}(0) = 0$. At $s = 4q$:

$$C'_{\text{Gap}}(4q) = \frac{4}{3}\Big[1 - \frac{2}{(1 + 4\eta q^2)^3}\Big].$$

Since $\eta \geq 1/(4q^2)$, we have $1 + 4\eta q^2 \geq 2$. Thus $(1 + 4\eta q^2)^3 \geq 8$, and $2/(1 + 4\eta q^2)^3 \leq 1/4$. Therefore, $C'_{\text{Gap}}(4q) \geq \frac{4}{3}(1 - 1/4) = 1$.

Since $C'_{\text{Gap}}(0) < 1$ and $C'_{\text{Gap}}(4q) \geq 1$, by the Intermediate Value Theorem, there exists a unique $\bar{s} \in (0, 4q]$ solving $C'_{\text{Gap}}(\bar{s}) = 1$. Since $\bar{s} \leq 4q < L/2$, $\bar{s}$ is an interior point. By convexity of $f(s)$, this unique critical point is the global minimum. Thus, $s^* = \bar{s}$.

The concluding bound $s^* \leq 4q$ follows immediately from all three parts. □

The second-period equilibrium value from investing in research at distance s is $W_2(s) = 1 - f(s)/L$. Let $C_{\text{Total}}(\mathcal{F}_1)$ be the cost if the organisation does not research (so $\mathcal{F}_1 = \{0\}$). The net gain from exploring is $\Delta W(s) = [C_{\text{Total}}(\mathcal{F}_1) - f(s)]/L$. The optimal gain is $\Delta W^* = \Delta W(s^*)$. The discounted hurdle rate for investing in research is $H^* = \delta \Delta W^*$; it is increasing in η and L under our assumptions regarding η.

The choice of new knowledge balances an *anchor effect* (allowing the expert another point of operation) and an *infrastructure effect* (allowing routine workers to extrapolate over two points). Localised expertise balances these while constraining exploration to local improvements. Because the optimal research distance s^* is bounded above by the routinisation limit, knowledge accumulates in small increments, which is a feature that greatly simplifies the infinite horizon problem analysed in Section 3.3. In contrast, when $\eta < 1/(4q^2)$ (broad expertise) it may be optimal to leap across larger gaps, complicating dynamic aggregation; details are relegated to Appendix B.

First-period allocation. In period 1, the manager observes a question at distance $|x| = s \in [0, L/2]$ and must decide whether to exploit or explore. The immediate advantage of answering the question with the expert is

$$A_1(s) = \max\{0, V_E(s) - V_R(s)\} = \max\Big\{0, \frac{s}{q} - \eta s^2\Big\}.$$

This function is unimodal in s with peak value $A_1\big(1/(2\eta)\big) = 1/(4\eta)$ when q has been normalised to 1. Research is undertaken if and only if $A_1(s) \leq H^*$. For small H^* (pure exploitation) the expert handles intermediate questions where $A_1(s) > H^*$ and the routine worker handles the remainder; for large H^* (pure research) the expert always explores.

Let $r_1 < r_2$ be the two distances satisfying $A_1(r_1) = A_1(r_2) = H^*$ when $H^* < \frac{1}{4\eta}$; see Figure 3. The first-period exploitation region is (r_1, r_2) and the probability that the expert conducts research is

$$\rho = 1 - \frac{2\big(\min\{r_2, L/2\} - r_1\big)}{L}.$$

As η increases within the localised interval, X_c and X_E shrink, the expert's static advantage

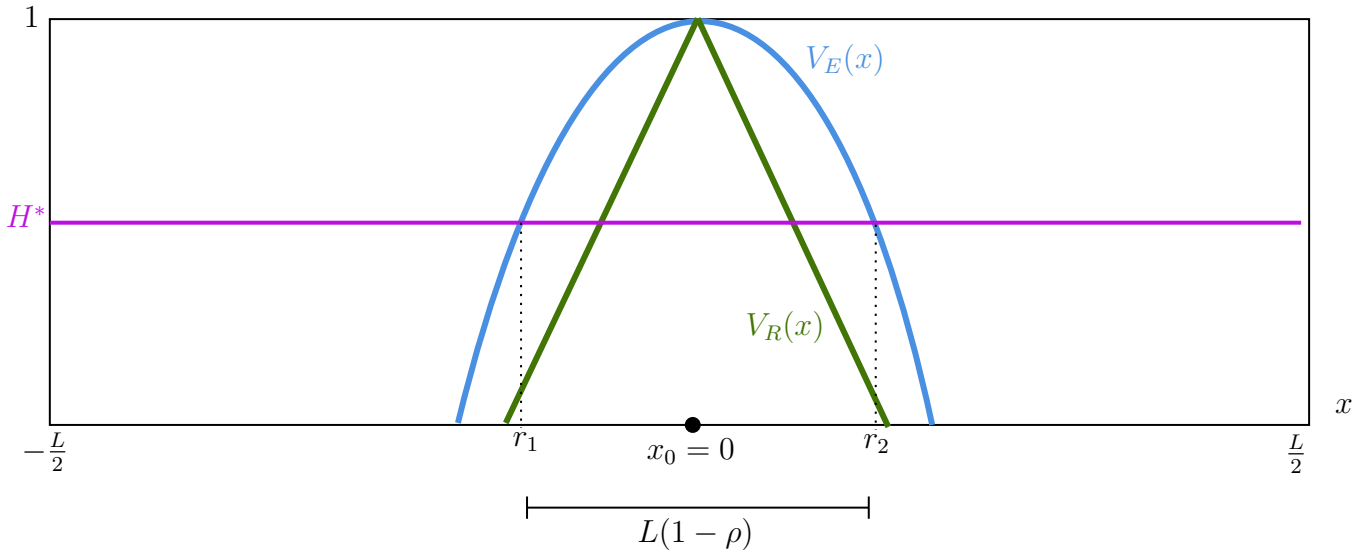


Figure 3: First period allocation for a given hurdle rate, H^*

$A_1(s)$ falls and s^* moves closer to the origin. This reduces (r_1, r_2) and increases exploration. Thus, a higher cost of expertise pushes the organisation towards local, incremental research. Figure 3 shows the resulting allocation and determinants of the probability of exploration.

3.3 Infinite Horizon Case

The two-period case shows how the anticipation of future knowledge points impacts the cognitive division of labour in a period through the interplay of $A(x|\mathcal{F}_t)$ and $H^*(\mathcal{F}_t)$. Specifically, a lower H^* or $A(.)$ leads to a reduction in ρ the probability of exploration. As we will see, in general, having more knowledge points leads to a reduction in H^* while the impact on $A(x)$ is more nuanced. These impacts will determine how much exploration occurs as more knowledge is accumulated.

We now turn to consider an organisation operating with an infinite horizon ($t = 1, 2, \dots$) to analyse these dynamics. All other aspects of the model, including one question per period arriving over the domain $[-L/2, L/2]$ are maintained starting from an initial knowledge set of $\mathcal{F}_0 = \{0\}$ and $\eta \in [\frac{1}{4q^2}, \frac{1}{q^2}]$ so that the expert's reach $X_E := \eta^{-1/2}$ satisfies $X_E \leq 2q$ and the allocation crossover $X_c := (\eta q)^{-1}$ satisfies $X_c \leq 4q$. Under these bounds, routine interpolation saturates before expert coverage does, enabling local, incremental innovations while ensuring that long tails can be served by the routine worker.

At date t the knowledge set is $\mathcal{F}_t = \{a_0, a_1, \dots, a_{n_t+1}\}$, with $a_0 \geq -L/2$ and $a_{n_t+1} \leq L/2$.

The domain is partitioned into two tails of lengths $Z_\ell(t) = a_1 - a_0$ and $Z_r(t) = a_{n_t+1} - a_{n_t}$, and n_t interior gaps $X_i(t) = a_{i+1} - a_i$ for $i = 1, \ldots, n_t$. The total expected cost is

$$C(\mathcal{F}_t) = C_{\text{Tail}}\big(Z_\ell(t)\big) + \sum_{i=1}^{n_t} C_{\text{Gap}}\big(X_i(t)\big) + C_{\text{Tail}}\big(Z_r(t)\big),$$

where C_{Gap} and C_{Tail} are the integrated costs for interior segments and tails derived in Section 3.2. Under Assumption 1, the gap cost C_{Gap} is strictly convex on $(0, 4q)$, while the tail cost C_{Tail} is strictly convex on $[0, X_E]$ and linear thereafter. Research in period t amounts to inserting one new anchor $x_{t+1} \in (-L/2, L/2)$, thereby splitting exactly one segment. The resulting knowledge set is denoted $\mathcal{F}_{t+1}$.

3.3.1 New knowledge point gains: gaps versus tails

The first task is to characterise where new knowledge points are added under exploration. In particular, we consider whether a new point lies in a gap between two existing points or at the tails that enlarge the boundaries of the organisation's knowledge set.[11] We distinguish interior gaps from tails because the cost functions differ at the endpoints.

Gap: Inserting a new anchor between two existing points divides the interval between them into two subsegments. The *split benefit* on that interval is the reduction in static cost from this division. The following result derives the optimal split location.

Lemma 1 (Midpoint optimality for interior gaps)**.** *Let $X > 0$ be the length of an interior gap. Splitting the gap at its midpoint $y = X/2$ maximises the split benefit*

$$\Delta_{\text{Gap}}(X) \;=\; C_{\text{Gap}}(X) \;-\; 2\,C_{\text{Gap}}\big(X/2\big).$$

Moreover $\Delta_{\text{Gap}}(X)$ is strictly increasing in X for $0 < X < 4q$.

Tail: Turning now to points on the tail, consider a tail of length Z. The post–split costs are $C_{\text{Gap}}(d)$ for the interior portion of length d and $C_{\text{Tail}}(Z - d)$ for the residual tail. The split benefit is

$$\Delta_{\text{Tail}}(Z; d) \;=\; C_{\text{Tail}}(Z) \;-\; \Big[C_{\text{Gap}}(d) \;+\; C_{\text{Tail}}(Z - d)\Big], \qquad 0 \le d \le \min\{Z, \bar{s}\},$$

where $\bar{s}$ denotes an upper bound on the length of new interior increments (defined precisely below). Because C_{Gap} and C_{Tail} differ in curvature, the optimal split distance generally

[11] Carnehl and Schneider (2025) termed these knowledge deepening and knowledge expansion respectively.

diverges from the midpoint.

Lemma 2 (Tail optimum and maximal tail dividend)**.** *Assume localised expertise. Let* $X_E = 1/\sqrt{\eta}$ *be the expert's reach. There is a unique* $\bar{s} \in (X_c, 4q]$ *satisfying*

$$C'_{\text{Gap}}(\bar{s}) \;=\; 1. \tag{1}$$

The tail cost derivative $C'_{\text{Tail}}(u)$ *increases on* $[0, X_E]$ *and is constant at* 1 *on* $[X_E, \infty)$*.*

(i) For any tail of length Z*, the split benefit* $\Delta_{\text{Tail}}(Z; d)$ *is maximised at* $d^*(Z)$*:*

- *If* $Z \geq \bar{s} + X_E$ *(Long Tail), then* $d^*(Z) = \bar{s}$*.*
- *If* $Z < X_E$ *(Short Tail), then* $d^*(Z) = 2Z/3$*.*
- *If* $X_E \leq Z < \bar{s} + X_E$ *(Intermediate Tail), then* $d^*(Z)$ *is the unique solution to* $C'_{\text{Gap}}(d) = \eta(Z - d)^2$*, and it satisfies* $Z - X_E < d^*(Z) < \bar{s}$*.*

(ii) Define the maximal one–step tail dividend by

$$\bar{\Delta}_{\text{Tail}} \;:=\; \sup_{Z \geq 0} \max_{0 \leq d \leq Z} \Delta_{\text{Tail}}(Z; d).$$

Then $\bar{\Delta}_{\text{Tail}} = \bar{s} - C_{\text{Gap}}(\bar{s})$*. This value is attained whenever* $Z \geq \bar{s} + X_E$*, and the dividend is strictly increasing for* $Z < \bar{s} + X_E$*.*

The proofs of these results are in Appendix C. The gap and tail results generally show that when L is sufficiently large, the tail points are extended at increments of $\bar{s}$ while gaps are split evenly. Any other complexity arises when knowledge reaches the edge of the domain of questions and relevant knowledge.

Gap vs tail. Having characterised the optimal gap and tail knowledge points, we now compare the two. Define the *gap dividend* $\Delta_{\text{Gap}}(X)$ from Lemma 1, the maximal tail dividend $\bar{\Delta}_{\text{Tail}}$ from Lemma 2, and the *crossover threshold*

$$X^\star \;:=\; \inf\bigl\{X > 0 :\; \Delta_{\text{Gap}}(X) \geq \bar{\Delta}_{\text{Tail}}\bigr\}.$$

This is the smallest interior gap length at which refining that gap yields a split dividend at least as large as splitting a sufficiently long tail. A key property is that $X^\star > 4q$ for all η in the localised expertise range. Appendix D proves that $\bar{\Delta}_{\text{Tail}} > \Delta_{\text{Gap}}(4q)$ for all $\eta \in [1/(4q^2),\, 1/q^2]$, implying that interior gaps up to length $4q$ never beat tail splits for a single knowledge point.

Proposition 2 (Expansion before refinement)**.** *Under localised expertise, let* $X_{\max}(t) = \max_i X_i(t)$ *be the length of the largest interior gap and* $Z_{\max}(t) = \max\{Z_\ell(t), Z_r(t)\}$ *the length of the longer tail. Then:*

(i) $X^\star > 4q$*. Hence* $\bar{\Delta}_{\text{Tail}} > \Delta_{\text{Gap}}(4q)$ *uniformly over the localised expertise interval.*

(ii) If there exists a tail with $Z_{\max}(t) \geq \bar{s} + X_E$ *(Pure Expansion), splitting that tail by* $\bar{s}$ *yields the maximal dividend* $\bar{\Delta}_{\text{Tail}}$ *and strictly dominates splitting any interior gap.*

(iii) If all tails satisfy $Z_{\max}(t) < \bar{s} + X_E$ *(Intermediate Expansion or Refinement), the greedy rule selects the segment (tail or gap) offering the largest dividend:* $\Delta^*(t) = \max\{\Delta_{\text{Tail}}(Z_{\max}(t)), \Delta_{\text{Gap}}(X_{\max}(t))\}$*. The organisation switches to refining interior gaps only when* $\Delta_{\text{Gap}}(X_{\max}(t)) \geq \Delta_{\text{Tail}}(Z_{\max}(t))$*.*

Sketch. Part (i) is established in Appendix D. Part (ii) follows from Lemmas 1 and 2: if any tail is long enough ($Z \geq \bar{s}+X_E$) to generate $\bar{\Delta}_{\text{Tail}}$, that dividend exceeds the gap dividend for all relevant interior gaps (by Part (i)). Part (iii) addresses the transition. Once the tails are shortened below $\bar{s}+X_E$, the tail dividend $\Delta_{\text{Tail}}(Z)$ is strictly less than $\bar{\Delta}_{\text{Tail}}$ (by Lemma 2(ii)). The greedy rule compares this value with the largest gap dividend, $\Delta_{\text{Gap}}(X_{\max})$, and selects the maximum of the two. □

Proposition 2 captures the core structural implication of localised expertise: *expansion precedes refinement.*[12] Because the maximal tail dividend exceeds the largest possible gap dividend for all interior gaps up to $4q$, the greedy rule marches outward, placing anchors at interval $\bar{s}$ along the tails until both tails have length strictly less than $\bar{s} + X_E$. Only then does the algorithm switch to splitting interior gaps.

3.3.2 The Hurdle Rate and Dynamic Policy Approximation

We now embed the static research benefits into the infinite horizon dynamic framework to determine the Organisational Hurdle Rate H_t. The Hurdle Rate is defined as the discounted change in the expected continuation value (NPV) of the organisation, $H_t^* = \delta[W(\mathcal{F}_{new}) - W(\mathcal{F}_t)]$. As noted in Section 2.3, calculating this value exactly requires solving the Bellman equation, which is analytically intractable due to the non-stationarity of the environment (see Appendix E).

To maintain analytical tractability, we employ a *Stationary Value Approximation* (SVA). The SVA calculates the continuation value by assuming that the expected per-period benefit

[12] Carnehl and Schneider (2025) have a related result characterised by the density of the knowledge set showing that refinement proceeds with expansion occuring only once knowledge has become sufficiently dense; see Gans (2025b) for examples with zero research cost.

generated by the new knowledge point persists indefinitely, ignoring the impact of future knowledge evolution on the value of the current research.

The expected per-period benefit is the integrated cost reduction $\Delta^*(t)$ normalised by the domain length L.[13]

Lemma 3 (The Hurdle Sequence under SVA)**.** *Under the Stationary Value Approximation (SVA), the Organisational Hurdle Rate at time t is given by:*

$$H_t^{SVA} \;=\; \frac{\delta}{1-\delta}\,\frac{\Delta^*(t)}{L}. \tag{2}$$

The sequence of SVA hurdle rates H_t^{SVA} is non-increasing over time. During the initial Pure Expansion phase (while $Z_{\max}(t) \geq \bar{s}+X_E$), it is constant at a plateau value H_P. Subsequently, during Intermediate Expansion and Refinement, the maximal dividend $\Delta^*(t)$ declines, and H_t^{SVA} is strictly decreasing.

Proof. Under the SVA, the change in continuation value is the present value of the perpetual stream of expected per-period gains $\Delta V = \Delta^*(t)/L$. This value is $\sum_{k=0}^{\infty} \delta^k \Delta V = \frac{1}{1-\delta}\Delta V$. The hurdle rate is the discounted value of this stream starting from the next period: $H_t^{SVA} = \delta\frac{1}{1-\delta}\Delta V$ (and, by definition, $H_t = \delta\Delta W_t$). The dynamics of H_t^{SVA} follow directly from the dynamics of $\Delta^*(t)$ established previously. □

The dynamic decision rule we analyse is the *SVA Policy*: Explore if $A(x(t)) \leq H_t^{SVA}$. This policy is inherently *greedy* regarding the choice of research project, as it always selects the project maximising $\Delta^*(t)$. Throughout the remainder of the analysis, we refer to the SVA hurdle rate simply as H_t.

Interpreting the SVA Policy. The SVA policy provides a tractable approximation of the optimal dynamic strategy. While it correctly captures the primary drivers of the research decision (the static gain Δ^*, patience δ, and domain size L), it introduces specific biases relative to the true optimal policy (H_t^*).

[13]The formulation of the Organisational Hurdle Rate incorporates a normalisation by L, the total length of the question domain. This normalisation is essential for maintaining dimensional consistency within the dynamic decision framework. The managerial decision rule hinges on a comparison between the immediate advantage of exploitation, $A(x(t))$, which is a *per-question* value, and the hurdle rate, H_t. In contrast, the benefit of research Δ^* is an *integrated* measure of improvement across the entire domain L. To establish a valid comparison, we must determine the expected value of the new knowledge realised in a single future period. Given that questions are drawn uniformly from the domain L (with probability density $1/L$), the expected per-period gain is the aggregate benefit averaged over the distribution of questions, yielding Δ^*/L. The Hurdle Rate is calculated based on this expected per-period value, ensuring a consistent comparison with the immediate opportunity cost of exploration, $A(x(t))$.

- **Refinement Phase (Overestimation):** During Refinement, the value of future research opportunities declines as the knowledge base matures. By treating the current gain as perpetual, the SVA ignores this deceleration. Consequently, $H_t^{SVA} > H_t^*$. The SVA policy leads to excessive exploration during the Refinement phase compared to the optimal policy.

- **Expansion Phase (Ambiguous Bias):** During Expansion, the static gain Δ^* is constant, but the probability of exploration $\rho(t)$ evolves. As we show in Section 4, this evolution depends on the model's exogenous parameters. If exploration is decreasing, $H_t^{SVA} > H_t^*$. However, if exploration is increasing (the S-curve case), $H_t^{SVA} < H_t^*$, because the SVA ignores the positive feedback loop where today's research accelerates future learning by increasing exploration.

Noting the direction of these biases, the SVA policy allows us to characterise the qualitative features of the organisational life cycle. The fundamental sequence of Expansion preceding Refinement (Proposition 2) is robust, as it depends only on the comparison of static gains (Δ^*).

3.3.3 Dynamics of exploration and exploitation

The evolution of the organisation is governed by the interplay between the immediate advantage of expert operation and the long-term value of research. While a full solution to the dynamic path is intractable, we can nonetheless characterise its broad shape by examining how two endogenous variables—the organisational hurdle rate H_t and the exploration probability $\rho(t)$—behave along the simple value-added (SVA) path.

At any point in time, the static advantage of using an expert instead of a routine worker depends on where a question falls relative to existing anchors. In an interior gap of length X, if the question lies at a distance $s \in [0, X/2]$ from the nearest anchor, the expert's advantage can be described by

$$A_X(s) \;=\; \frac{s}{q} \;-\; \beta(X)\, s^2, \qquad \beta(X) = \eta + \frac{1}{qX}.$$

This function captures the trade-off between the routine worker's interpolation error and the expert's localised application cost. It is unimodal: near the anchor, the routine worker interpolates well, and the expert's advantage is small; it increases towards the gap's midpoint, then declines as questions approach the far anchor. For tail segments (where the routine worker extrapolates rather than interpolates), the same quadratic form applies, except that

the coefficient of the linear term from the routine worker's continuation value is constant (i.e., independent of gap length), yielding a similar unimodal structure.

A new anchor is created when $A(s) \leq H_t$, meaning that the expert sacrifices current advantage for future knowledge if that sacrifice is less than the hurdle rate. We define exploration as the ex-ante probability that this inequality holds:

$$\rho(t) \;=\; P\big(A(s \mid \mathcal{F}_t) \;\leq H_t\big).$$

It measures how often the organisation prioritises research over immediate production.

To understand how exploration responds to changes in the knowledge structure, it is useful to fix a hurdle rate H and vary the gap length X. The exploration region is given by the set of distances s for which $A_X(s) \leq H$. As X grows from zero, this region initially shrinks: in very short gaps the routine worker's precision is high and expert advantage is nearly zero near the anchor, so exploration occurs in small neighbourhoods; increasing X widens the region where the routine worker remains precise, squeezing the exploration region. Beyond a critical gap length X_c, however, further increases in X expand the central part of the gap where the routine worker performs poorly, making it more likely that $A_X(s)$ falls below H. Thus, exploration probability $\rho(X)$ falls for $0 < X \leq X_c$ and rises for $X \geq X_c$; it is U-shaped in the gap length.

Along the dynamic SVA path, both H_t and the distribution of gaps evolve endogenously. During the *Pure Expansion* phase, when long tails are being converted to interior gaps of fixed length $\bar{s}$, the value of a new anchor is constant, so the hurdle rate stays at a plateau H_P. Because the knowledge structure consists of a mix of long tails and newly created interior segments, exploration $\rho(t)$ is bounded away from zero and one. As the organisation enters the *Intermediate Expansion* phase, both tails shorten and the marginal benefit of new anchors begins to decline, causing H_t to fall. Nevertheless, exploration remains significant while tails are still the primary source of improvement.

Eventually, the largest remaining interior gap becomes the most attractive location to place the next anchor. At this point the organisation transitions to the *Refinement* phase. Now the largest gap length $X_{\max}(t)$ and the hurdle rate H_t both decrease monotonically. Although, statically, shrinking the gap would shift the organisation into the Expert-Dominated regime (where $\rho(X)$ tends to rise as X shrinks below X_c), the dynamic decline in the hurdle rate dominates: because each new anchor yields a diminishing incremental benefit, H_t falls faster than the gains from reducing X can boost exploration. Consequently, organisational exploration decays to zero ($\rho(t) \to 0$) and the expert increasingly exploits existing knowledge rather than creating new anchors.

In summary, despite the intractability of a closed-form solution, we can trace a clear narrative: as knowledge accumulates, the organisation initially maintains considerable exploration to build a broad knowledge base. Once tails have been sufficiently trimmed, the trade-off shifts towards refining the remaining gaps. Eventually, when marginal returns to research become too small, exploration ceases and the organisation fully exploits its accrued expertise.

3.3.4 Long-run outcomes

In the Refinement phase, the organisation's remaining interior gaps are systematically subdivided by placing new anchors at their midpoints. Under the SVA policy, this "greedy" splitting implies that the largest gap length $X_{\max}(t)$ declines monotonically: each time it is chosen for refinement, it is halved. Over time, this geometric decay in the largest gap size determines the asymptotic behaviour of other key variables.

The benefit of adding an anchor to a small gap shrinks rapidly as the gap length itself becomes small. Once the organisation enters the Expert-Dominated regime—when the largest gap $X_{\max}(t)$ falls below the critical threshold X_c—the integrated cost in that gap scales as $C_{\text{Gap}}(X) \approx \frac{\eta}{12}X^3$. Splitting such a gap then yields a benefit

$$\Delta_{\text{Gap}}(X) \;=\; C_{\text{Gap}}(X) - 2C_{\text{Gap}}(X/2) \;\approx\; \frac{\eta X^3}{16},$$

so the advantage of further refinement is proportional to the cube of the gap length. Because the organisational hurdle rate is the discounted value of the next incremental gain, it falls proportionally: $H_t = \frac{\delta}{1-\delta}\,\frac{\Delta_{\text{Gap}}(X_{\max}(t))}{L}$ and hence $H_t \propto X_{\max}(t)^3$.

A shrinking hurdle rate reduces the incentive for exploration. Although, in static terms, exploration tends to become more likely in shorter gaps, the decline in H_t dominates. By explicitly comparing the advantage function $A_X(s)$ and the hurdle in the Expert-Dominated regime, one finds that the probability that the expert explores in a gap of length X decays quadratically: as $X \to 0$,

$$\rho(X; H_t) \;\approx\; \frac{1 - \sqrt{1 - 4q^2\beta(X)H_t}}{1 + \eta q X} \;\sim\; 2q\,K\,X^2,$$

where K is the constant of proportionality in the relationship $H_t = KX^3$ and $\beta(X) \approx 1/(qX)$. Thus organisational exploration during the Refinement phase satisfies $\rho(t) \propto X_{\max}(t)^2$.

Because the expected improvement at time t is the product of the probability of explo-

ration and the benefit from splitting the largest gap,

$$\mathbb{E}\big[W_{t+1} - W_t\big] \;=\; \rho(t)\,\frac{\Delta_{\mathrm{Gap}}(X_{\max}(t))}{L},$$

the improvement rate decays as the fifth power of the gap length. Combining the quadratic fall in $\rho(t)$ with the cubic decay in $\Delta_{\mathrm{Gap}}(X)$ yields $\mathbb{E}[W_{t+1} - W_t] \propto X_{\max}(t)^5$. Since the largest gap is halved each time it is selected for refinement, $X_{\max}(t)$ converges to zero, and so does the organisation's incremental growth.

This quintic deceleration underscores why performance gains flatten sharply as the organisation approaches its end-state: both the marginal return to adding new anchors and the frequency with which the expert is willing to do so collapse. In the limit, the firm's knowledge base is almost complete, the expert's role in research becomes negligible, and the routine worker handles nearly all questions using a highly refined, codified infrastructure.

3.3.5 The Organisational Life Cycle

The analysis of the SVA dynamics reveals a distinct life cycle for the organisation, characterised by shifting strategic priorities and the endogenous evolution of the knowledge base. Under the assumption of localised expertise, the organisation evolves through three sequential phases:

1. **Phase I: Pure Expansion (Building Breadth).** The organisation prioritises expanding the breadth of its knowledge domain. As long as a tail length $Z_j(t)$ exceeds the threshold $\bar{s} + X_E$, the maximal split benefit is achieved by splitting the tail (Proposition 2). The organisation marches outward in increments of $\bar{s}$.

    - *Dynamics:* The hurdle rate H_t is constant at its plateau value $H_P = \frac{\delta}{1-\delta}\,\frac{\bar{\Delta}_{\mathrm{Tail}}}{L}$. Exploration is prevalent, driven by the high returns to establishing foundational infrastructure in unexplored territory.

2. **Phase II: Intermediate Expansion (Transition).** Once both tails are shorter than $\bar{s} + X_E$, the maximal tail dividend $\Delta_{\mathrm{Tail}}(Z)$ begins to decline (Lemma 2(ii)).

    - *Dynamics:* The hurdle rate H_t starts to fall. The organisation continues to expand, but the increments $d^*(Z)$ become smaller than $\bar{s}$. The greedy rule compares the declining tail dividend with the largest gap dividend.

3. **Phase III: Refinement (Building Depth).** Eventually, the dividend from splitting the largest interior gap exceeds the largest tail dividend. The strategic priority shifts

from expansion to refinement (knowledge deepening). The organisation systematically fills the gaps between existing anchors.

- *Dynamics:* The largest gap length $X_{\max}(t)$ decays geometrically. The hurdle rate H_t continues to fall, scaling as $X_{\max}(t)^3$. Organisational exploration $\rho(t)$ declines rapidly towards zero. The rate of improvement decelerates sharply.

The life cycle describes a transition from an organisation focused on exploration and breadth, necessary to cover the domain of potential questions, to a mature organisation focused on exploitation and depth. In the mature stage, the organisation is highly routinised: the dense infrastructure allows the routine worker to handle most tasks efficiently.

It is worth noting that this characterisation relies on the SVA policy. As established in Appendix E, the SVA overestimates the value of research during the Refinement phase because it ignores the declining value of future research opportunities. Consequently, the true optimal policy would likely transition towards exploitation more rapidly in the mature stage than predicted by the SVA analysis.

4 The Dynamics of Organisational Performance

The analysis in Section 3.3 characterised the evolution of the firm's knowledge base and the resulting division of labour, establishing a three-phase life cycle: Pure expansion followed by continued expansion and then refinement. We now turn to the macroscopic implications of these micro-dynamics: the trajectory of the organisation's overall performance over time. A central question in the literature on organisational learning (e.g., March (1991)) is whether performance evolves along an S-shaped curve: growth is initially slow, then accelerates as foundational knowledge is established, and finally decelerates as the knowledge base matures and the returns to further codification diminish.

This section formally investigates whether and when the Babbage Firm exhibits this S-shaped dynamic. We demonstrate that while the refinement phase always exhibits deceleration (concave growth), the dynamics during the expansion phase are complex. The shape of the initial growth trajectory depends critically on the value of research relative to the immediate gains from exploitation.

4.1 Measuring Performance and Growth

We define organisational performance at time t as the expected value generated per period, $W_t = W(\mathcal{F}_t)$. The trajectory of performance is determined by the expected growth rate,

$G(t)$:

$$G(t) = \mathbb{E}[W_{t+1} - W_t].$$

Under the SVA, the expected growth rate is the probability of exploration multiplied by the expected gain from the optimal research project:

$$G(t) = \rho(t) \cdot \frac{\Delta^*(t)}{L}.$$

An S-shaped trajectory requires an initial phase where $G(t)$ is increasing (convex growth or acceleration), followed by a phase where $G(t)$ is decreasing (concave growth or deceleration).

4.2 The Phases of Growth

We analyse the evolution of the growth rate $G(t)$ during the distinct phases of the life cycle under localised expertise.

4.2.1 Deceleration Phases: Intermediate Expansion and Refinement (Concave Growth)

Once the organisation has moved beyond the Pure Expansion stage, the growth dynamics lose their ambiguity. Both in the Intermediate Expansion and the ensuing Refinement phase, the forces driving performance gains are unequivocally downward.

During Intermediate Expansion, the tails of the knowledge base are already short enough that the benefit from pushing farther outward is declining. Lemma 2(ii) shows that once a tail's length falls below $\bar{s} + X_E$, the incremental dividend from expanding it diminishes; as this occurs simultaneously on both sides of the knowledge base, the maximum available split benefit $\Delta^*(t)$ falls over time. Likewise, when the organisation switches to refining interior gaps, the largest gap length $X_{\max}(t)$ shrinks with each new anchor, and as noted earlier, the associated split benefit drops proportionally to $X_{\max}(t)^3$. In both cases, the probability that the expert chooses to explore $\rho(t)$ is also declining. (See Appendix F for more details).

The expected growth rate during the last two phases is, therefore, the product of two positive factors that are each contracting over time. As a result, $G(t)$ itself is strictly decreasing through the Intermediate Expansion and Refinement phases, and the performance path W_t traced out by cumulative gains is strictly concave in these later stages. This confirmation of the deceleration half of the S-curve intuition underscores that any period of accelerating growth must be confined to the organisation's early, Pure Expansion phase.

4.2.2 Phase 1: Pure Expansion (The Potential for Acceleration)

During the Pure Expansion phase (when $Z_{\max}(t) \geq \bar{s} + X_E$), the organisation expands its breadth by expanding the tails in increments of $\bar{s}$. The benefit of research $\Delta^*(t)$ is constant ($\bar{\Delta}_{\text{Tail}}$). Under the SVA policy, the hurdle rate H_t is constant at the plateau value:

$$H_P = \frac{\delta}{1-\delta}\frac{\bar{\Delta}_{\text{Tail}}}{L}.$$

The growth rate is $G(t) = \rho(t) \cdot (\bar{\Delta}_{\text{Tail}}/L)$. Therefore, the dynamics of $G(t)$ depend entirely on the evolution of the organisational exploration $\rho(t)$. Acceleration occurs if and only if $\rho(t)$ increases over time.

4.3 Exploration Dynamics during Pure Expansion

During the *Pure Expansion* phase, the organisation steadily adds new anchors at fixed intervals of length $\bar{s}$. Each expansion step replaces a *tail segment*—a long extrapolation from the last anchor—by an *interior gap* between two anchors. Whether this transition raises or lowers the organisation's propensity to explore depends on two forces.

Expert advantage in tails and gaps. When the expert answers a question at distance s from the nearest anchor, she enjoys an advantage of $\frac{s}{q} - \eta\, s^2$ in a tail. The first term $\frac{s}{q}$ reflects the extrapolation error of the routine worker. The second term ηs^2 captures the cost of applying tacit knowledge. The advantage is unimodal, peaking at a moderate distance and falling to zero at the expert's reach $X_E = 1/\sqrt{\eta}$. Under localised expertise, the expansion step $\bar{s}$ is longer than X_E. Beyond X_E, the expert's operational value is zero, so questions in the region $(X_E, \bar{s}]$ are *unserved*, and the advantage is zero.

Inside a gap, the variance of routine interpolation is lower because each point lies between two anchors. This improves the routine worker's performance and steepens the expert's advantage curve: the advantage in a gap takes the same form but with a larger curvature term. We denote the curvatures as $a_{\text{tail}} = \eta$ and $a_{\text{gap}} = \eta + 1/(q\bar{s})$. As a result, the expert's advantage in a gap is always below that in a tail at any given distance s.

Two opposing forces. Converting a tail segment into a gap therefore has two effects on exploration at a given hurdle rate H:

1. **Infrastructure effect.** Because the expert's advantage is lower in the gap than in the tail, many points that were previously exploited now become worthwhile to explore. This "infrastructure" improvement grows more important as the organisation raises its

hurdle H. Because the advantage curve in the gap is more compressed, the exploited region inside a gap shrinks faster than that inside a tail as H rises.

2. **Unserved–frontier effect.** The tail contains an unserved frontier of length $\bar{s} - X_E$. Since the advantage here is zero, any positive hurdle implies that all those points are explored. In a gap, every point is served. Thus, replacing a tail with a gap removes this block of guaranteed exploration. This loss matters most when H is low and the organisation is inclined to exploit.

These forces push exploration in opposite directions. Let $\Delta\rho(H)$ denote the net change in the measure of exploration over the segment. At a very low hurdle ($H = 0$), the organisation only explores when the expert's advantage is non-positive. In the gap, this is a narrow band in the centre, whereas in the tail, it is the entire unserved frontier. Since the frontier is larger than the central band, $\Delta\rho(0) < 0$; exploration falls when H is too low. At a high H, the infrastructure effect dominates. These observations imply that $\Delta\rho(H)$ is continuous and strictly increasing in H and must cross zero exactly once.

The critical hurdle. We identify the unique threshold H_{crit} at which the total measure of exploration is unchanged ($\Delta\rho(H_{\text{crit}}) = 0$). This occurs when the length of the *exploited region*—where the expert's advantage exceeds H—is equal in the tail and the gap configurations.

With $A(s) = \frac{s}{q} - as^2$, the exploited region on one side of the segment is the interval between two roots of $A(s) = H$; its length is

$$\text{exploited length}(a, H) = \frac{1}{a}\sqrt{\frac{1}{q^2} - 4aH}.$$

A tail has a single exploited region with curvature a_{tail}; its length is $\ell_{\text{tail}}(H)$. A gap has two symmetric exploited regions with curvature a_{gap}; its total exploited length is $2\,\ell_{\text{gap}}(H)$.

At the critical hurdle, the exploited lengths are equal, $\ell_{\text{tail}}(H_{\text{crit}}) = 2\,\ell_{\text{gap}}(H_{\text{crit}})$. Solving this equation yields the closed form for H_{crit}. It is useful to express this threshold relative to the maximum advantage in the tail, $A_{Tail}^{max} = 1/(4q^2\eta)$. Let R be the ratio of the curvatures, $R = a_{\text{gap}}/a_{\text{tail}} = 1 + 1/(\eta q\bar{s})$. This ratio captures how much steeper the advantage curve is in the gap relative to the tail. The critical hurdle is:

$$H_{\text{crit}} = A_{Tail}^{max} \cdot \frac{4 - R^2}{R(4 - R)} = \frac{1}{4q^2\eta} \cdot \frac{(\eta q\bar{s} - 1)(3\eta q\bar{s} + 1)}{(\eta q\bar{s} + 1)(3\eta q\bar{s} - 1)}.$$

When the gap is very similar to the tail ($R \approx 1$, meaning $\bar{s}$ is very large), the infrastruc-

ture difference is minimal, and H_{crit} is high. When the gap is significantly steeper, the infrastructure effect is strong even at low hurdles, and H_{crit} is low.

In sum, during Pure Expansion exploration will *decrease* when the organisation places little value on research ($H < H_{\text{crit}}$) because the loss of the unserved frontier outweighs the infrastructure gain, and it will *increase* when research is sufficiently valuable ($H > H_{\text{crit}}$). The unique threshold H_{crit} depends on the precision of interpolation q and the localisation of expertise η (which jointly determine $\bar{s}$ and R).

Comparative statics and the SVA plateau H_P. The phase outcome is governed by the inequality $H_P \gtrless H_{\text{crit}}$. We examine how the exogenous parameters affect both the organisation's valuation of research (H_P) and the threshold required for acceleration (H_{crit}). This analysis must account for the fact that the expansion step $\bar{s}$ is endogenous, optimally chosen based on η and q.

During Pure Expansion, the SVA hurdle is $H_P = \frac{\delta}{1-\delta} \cdot \frac{\bar{\Delta}_{\text{Tail}}}{L}$. The value of research $\bar{\Delta}_{\text{Tail}}$ increases with better interpolation ($q \uparrow$) and with lower expert cost ($\eta \downarrow$). Hence, H_P increases with q and decreases with η.

The critical hurdle H_{crit} also depends on these parameters, but through their impact on the balance between the two opposing forces:

- **Higher interpolation tolerance ($q \uparrow$).** An increase in q disproportionately improves routine performance in the tail configuration (where uncertainty is higher). This strengthens the relative impact of the Unserved Frontier Effect, meaning a higher hurdle is required to favour acceleration. Thus, H_{crit} increases with q.

- **More localised/expensive expertise ($\eta \uparrow$).** An increase in η reduces the expert's advantage $A(s)$ everywhere. Because the advantage profile in the gap is more sensitive (more compressed) than in the tail, this shift strengthens the relative impact of the Infrastructure Effect, making acceleration easier to achieve. Thus, H_{crit} decreases with η.

The net impact of the technology parameters on the likelihood of acceleration ($H_P > H_{\text{crit}}$) is analytically ambiguous because they move both thresholds in the same direction:

- **Impact of q.** An improvement in interpolation technology ($q \uparrow$) generates two conflicting dynamics. It increases the value of research ($H_P \uparrow$), promoting acceleration. Simultaneously, it raises the critical threshold ($H_{\text{crit}} \uparrow$), resisting acceleration.

- **Impact of η.** An increase in the cost of expertise ($\eta \uparrow$) also creates conflicting dynamics. It decreases the value of research ($H_P \downarrow$), resisting acceleration. Simultaneously, it lowers the critical threshold ($H_{\text{crit}} \downarrow$), promoting acceleration.

Although we have analytical expressions for both H_P and H_{crit} under the SVA, we cannot definitively sign the net effect. This ambiguity arises because the complexity of the expressions—particularly the implicit definition of $\bar{s}$—prevents an analytical comparison of the magnitudes of these shifts. The relative sensitivity of H_P and H_{crit} depends on the specific configuration of the parameters.

In contrast, factors that solely affect the valuation of research have clear effects. Higher patience ($\delta \uparrow$) increases H_P without affecting H_{crit}, making acceleration (S-shaped growth) more likely, while a larger domain ($L \uparrow$) decreases H_P, making deceleration more likely.

4.4 The Shape of the Organisational Trajectory

Based on the analysis of how $\rho(t)$ changes with time analysed in Section 4.2.1, we can now characterise the dynamics of organisational performance. The overall shape depends on where the actual hurdle rate during expansion, H_P, falls relative to the critical threshold H_{Crit}. The following theorem summarises these outcomes:

Theorem 1 (Organisational Learning Dynamics under SVA)**.** *Under the SVA policy, the shape of the expected organisational performance trajectory* $\mathbb{E}[W_t]$ *is determined by the hurdle rate during expansion,* H_P*.*

1. ***Concave (Always Decelerating):*** *If* $H_P \leq H_{Crit}$*. The trajectory is globally concave.*

2. ***S-shaped (Convex-Concave):*** *If* $H_{Crit} < H_P < A_{Tail}^{max}$*. The trajectory is initially convex during Expansion and subsequently concave during Refinement.*

3. ***Linear-Concave:*** *If* $H_P \geq A_{Tail}^{max}$*. The trajectory is initially linear during Expansion and subsequently concave during Refinement.*

The shape of the learning curve is determined by how the incentives for exploration evolve as the knowledge base expands.

- **Concave Growth (Low Value of Research):** When H_P is low (e.g., low δ or large L), the organisation prioritises exploitation. The Unserved Frontier Effect dominates: expansion eliminates unserved regions (where exploration was guaranteed), reducing overall exploration and causing immediate deceleration.

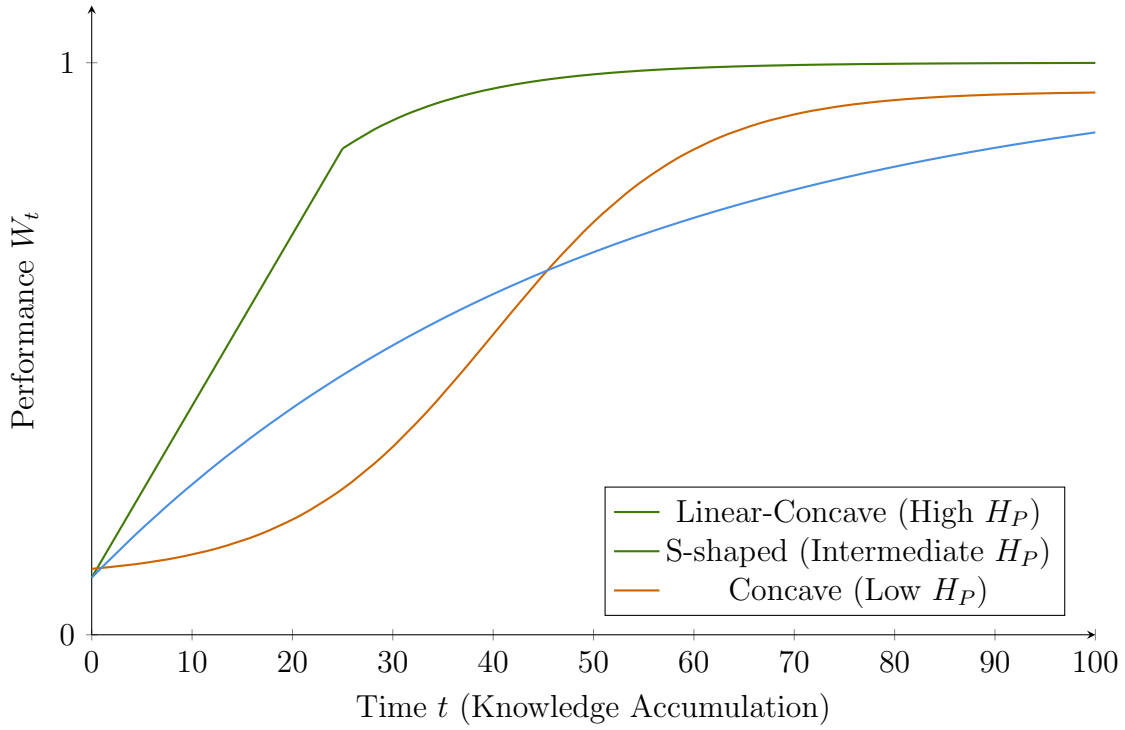


Figure 4: The three dynamic regimes of organisational performance. The shape depends on the value of research (H_P).

- **S-shaped Growth (Intermediate Value of Research):** When H_P is intermediate, the Infrastructure Effect dominates. The superior efficiency of the gap configuration encourages more exploration compared to the tail configuration. Exploration increases, leading to accelerating growth during the Expansion phase.
- **Linear-Concave Growth (High Value of Research):** When H_P is very high (e.g., high δ or small L), the value of research swamps the immediate operational advantages. Exploration is 100% throughout the Expansion phase, leading to maximum speed (linear growth).

Figure 4 illustrates these three dynamic regimes.

4.5 Interpretation of Dynamics under the SVA

The characterisation of the performance trajectory in Theorem 1 is derived under SVA. It is important to consider how the biases inherent in the SVA might affect the interpretation of these regimes relative to the true optimal policy.

As discussed in Section 3.3.2, the bias of the SVA hurdle rate H_P depends on the dynamic regime itself.

- **Concave Regime ($H_P \leq H_{Crit}$):** In this regime, exploration is decreasing during Expansion. The SVA overestimates the true hurdle rate ($H_P > H_P^*$). This implies that the SVA policy explores more than the optimal policy, but both policies predict a Concave trajectory when the value of research is sufficiently low.

- **S-shaped Regime ($H_{Crit} < H_P < A_{Tail}^{max}$):** In this regime, exploration is increasing during Expansion. The SVA underestimates the true hurdle rate ($H_P < H_P^*$) because it ignores the positive feedback where current research accelerates future learning by increasing exploration (a "learning-to-learn" effect). This implies that the true optimal policy values research more highly than the SVA suggests. Therefore, the S-shaped dynamic is likely robust: if the SVA policy predicts acceleration, the true optimal policy (with a higher hurdle rate) is even more likely to exhibit acceleration. The range of parameters supporting the S-shaped regime might be broader under the optimal policy.

- **Linear-Concave Regime ($H_P \geq A_{Tail}^{max}$):** In this regime, exploration is constant (100%). The SVA is exact, and the characterisation is accurate.

In summary, the qualitative distinction between accelerating (S-shaped/Linear) and decelerating (Concave) growth phases appears robust. The SVA approximation may slightly mischaracterise the precise thresholds between these regimes, but the underlying mechanisms—the interplay between the Infrastructure Effect and the Unserved Frontier Effect—remain central to the dynamics of organisational learning.

5 The Impact of Artificial Intelligence Adoption

The final section integrates the comparative-dynamics results from Section 4 with the discussion of artificial intelligence (AI) adoption. Recent evidence from academia and industry suggests that generative artificial intelligence is lowering the cost of expertise (Agrawal et al., 2024a). Such tools are also impacting knowledge creation (Agrawal et al., 2024b). These developments raise the possibility that the shadow cost of expertise operation may fall dramatically over time (Agrawal et al., 2025). At the same time, AI can rapidly ingest and synthesise vast amounts of data, helping routine workers bridge disparate knowledge sources (Gans, 2025a). From the perspective of this paper, AI technologies thus act on both sides of the cognitive hierarchy: *interpolation AI* enhances the productivity of routine workers (as captured by an increase in q), while *application AI* reduces the costs of experts in solving problems (as captured by a reduction in η). The motivation for these two AI classes is as follows:

1. **Interpolation AI (affecting q).** Tools such as retrieval-augmented generation, domain-specific search engines, and generative copilots allow routine workers to interpolate across codified knowledge with higher precision. They reduce the variance of routine work outputs, effectively increasing the precision tolerance q. Empirically, AI-native startups exploit such tools from the outset to operate with lean teams and to codify knowledge rapidly. Interpolation AI also enables non-experts to synthesise disparate information sources and thereby frees expert time for higher-value activities.

2. **Application AI (affecting η).** Technologies such as expert systems, advanced simulation tools, and AI copilots augment the tacit problem-solving capabilities of experts. They reduce the cost of expert application η, allowing an expert to handle more complex cases per unit time. Professional-services surveys find that early AI adopters anticipate that such tools will drive revenue growth and improve career satisfaction. However, these tools may encourage reliance on expert operations rather than investing in exploration.

We focus on how these two families of AI tools influence and reshape the organisation's performance trajectory. We then examine how their impacts differ across the firm life cycle and conclude by revisiting the broader implications for the cognitive division of labour.

5.1 Adoption of AI and the S-Curve

Theorem 1 links the shape of the performance trajectory to the hurdle rate H_P during the expansion phase, specifically whether H_P exceeds the critical threshold H_{Crit}. The adoption of AI shifts both H_P and H_{Crit} by altering the underlying parameters (η, q). Interpolation AI raises q, whereas Application AI lowers η. We analyse how these changes impact the organisational dynamics.

Proposition 3 (AI Adoption and the Performance Trajectory)**.** *Consider the SVA policy in a localised expertise regime.*

1. ***Interpolation AI (higher q).*** *Adoption increases the hurdle rate H_P by increasing the value of research $\bar{\Delta}_{\text{Tail}}$. However, it also increases the critical threshold H_{Crit}. The net effect on the likelihood of an S-shaped trajectory ($H_P > H_{Crit}$) is ambiguous.*

2. ***Application AI (lower η).*** *Adoption increases the hurdle rate H_P by increasing the value of research $\bar{\Delta}_{\text{Tail}}$. However, it also increases the critical threshold H_{Crit}. The net effect on the likelihood of an S-shaped trajectory ($H_P > H_{Crit}$) is ambiguous.*

Proof. We analyse how changes in q and η affect H_P and H_{Crit}. $H_P \propto \bar{\Delta}_{\text{Tail}}$. H_{Crit} is defined by $\Delta S(H_{Crit}) = 0$. We utilise the Envelope Theorem for H_P and the Implicit Function Theorem for H_{Crit}. Since $\partial(\Delta S)/\partial H > 0$ (Lemma **??**), the sign of dH_{Crit}/dx is opposite to the sign of $\partial(\Delta S)/\partial x$ for parameter x.

Part (1): Interpolation AI (higher q). *Effect on H_P:* $\frac{d\bar{\Delta}_{\text{Tail}}}{dq} = -\frac{\partial C_{\text{Gap}}(\bar{s})}{\partial q} > 0$. H_P increases.

Effect on H_{Crit}: We examine $\partial(\Delta S)/\partial q$. An increase in q improves V_R. $\partial V_R/\partial q = \sigma^2/q^2$. Since $\sigma^2_{Tail} > \sigma^2_{Gap}$, V_R improves more in the Tail configuration, meaning A_{Tail} decreases more than A_{Gap}. Thus, S_{Tail} increases more than S_{Gap}. $\partial(\Delta S)/\partial q < 0$. Therefore, $dH_{Crit}/dq > 0$.

Since both H_P and H_{Crit} increase, the net effect is ambiguous.

Part (2): Application AI (lower η). *Effect on H_P:* $\frac{d\bar{\Delta}_{\text{Tail}}}{d\eta} = -\frac{\partial C_{\text{Gap}}(\bar{s})}{\partial \eta} < 0$ (since $\partial C_{\text{Gap}}/\partial \eta > 0$, see Appendix A). A decrease in η increases H_P.

Effect on H_{Crit}: We examine $\partial(\Delta S)/\partial \eta$. The advantage is $A(s) = V_E(s) - V_R(s)$. An increase in η decreases $A(s)$ since $\partial A/\partial \eta = -s^2$. This increases the exploration set $S(H)$. As established in the proof of Lemma 4, S_{Gap} is more sensitive to shifts in the advantage function than S_{Tail} ($S'_{Gap}(H) > S'_{Tail}(H)$). Thus, S_{Gap} increases more than S_{Tail}. $\partial(\Delta S)/\partial \eta > 0$. Therefore, $dH_{Crit}/d\eta < 0$.

A decrease in η increases H_P and increases H_{Crit}. The net effect on the condition $H_P > H_{Crit}$ (S-shape) is ambiguous. □

While Proposition 3 shows that the impact on the *likelihood* of the S-shaped regime (the condition for acceleration) is formally ambiguous for both types of AI, the underlying mechanisms reveal distinct effects on the *rate* of organizational learning, $G(t) = \rho(t) \cdot \Delta^*(t)/L$. This rate is driven by how each technology affects the immediate incentive to explore (organisational exploration, $\rho(t) = P(A(s) \leq H_P)$).

- **Interpolation AI ($q \uparrow$)** improves routine worker performance. This reduces the expert's immediate advantage $A(s)$, making exploration more attractive today. Since both the value of research (Δ^*, proportional to H_P) and the propensity to research (ρ) increase, the growth rate $G(t)$ accelerates significantly.

- **Application AI ($\eta \downarrow$)** improves expert performance. This creates two opposing forces on exploration. First, it increases the value of research H_P, favouring exploration. Second, it increases the immediate advantage of expert operation $A(s)$, favouring exploitation. This latter effect represents a "crowding-out" of investment in codified knowledge: as experts become cheaper to deploy immediately, the organisation may

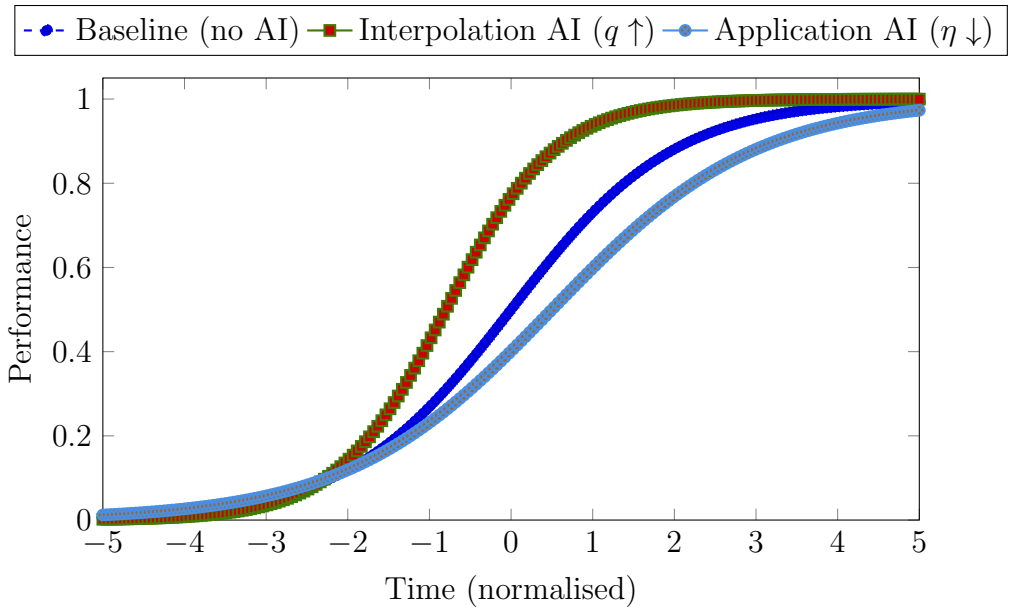


Figure 5: Illustrative effect of Interpolation AI ($q \uparrow$) and Application AI ($\eta \downarrow$) on the performance trajectory. *Schematic only.* Axes are not calibrated; curves illustrate the distinct dynamics driven by changes in organisational exploration and the value of research.

defer research. The net effect on exploration $\rho(t)$, and therefore the speed of learning $G(t)$, is ambiguous.

Figure 5 provides an illustrative sketch of these dynamics. The baseline curve (no AI) exhibits a modest S-shape. Under Interpolation AI, the increased exploration and higher value of research lead to unambiguously faster learning, illustrated by a steeper curve and an earlier inflection point. Under Application AI, the outcome is ambiguous. We illustrate a scenario where the crowding-out effect dampens the acceleration: the curve may be flatter or the acceleration phase delayed, reflecting the tension between immediate efficiency and long-term knowledge investment.

5.2 Lifecycle Dynamics: Start-ups versus Mature Firms

The effect of AI adoption is mediated by a firm's position on the organisational life cycle. The dynamic Babbage-Firm model implies that nascent firms operate near the *stagnation equilibrium*: high expert advantage $A(x)$ and low exploration probability ρ. Growth-phase firms operate on the steep part of the S-curve, investing heavily in codification. Mature firms operate near the saturation plateau, with dense infrastructure and low hurdle rates.

Nascent firms and start-ups. Start-ups have limited infrastructure and face high operational pressure. Interpolation AI can be a lifeline because it reduces the operational advantage $A(x)$ and creates organisational exploration, allowing some expert time to be redirected to research. Discovery-like effects (now captured by improvements in q) are particularly powerful in escaping the stagnation trap. Note that Application AI may be detrimental at this stage: by making expert operations cheaper, it crowds out investment in codified knowledge and entrenches reliance on experts.

Growth firms. Organisations in the growth phase already allocate significant resources to exploration. Interpolation AI accelerates learning by increasing the precision tolerance q, thereby steepening the S-curve. Application AI may improve operational efficiency, but risks diverting resources away from exploration.

Mature firms. Mature firms have dense knowledge bases and operate near the upper plateau of the S-curve. For them, Interpolation AI delivers efficiency gains: by raising q, it allows routine workers to leverage the extensive codified infrastructure more effectively. Application AI can be valuable in capturing residual tacit tasks, but it may provide diminishing returns.

The contrasting impacts of Interpolation and Application AI reflect a fundamental trade-off in knowledge-intensive organisations. Interpolation AI improves the generalists' toolkit, reduces the variance of routine work, and raises the value of codifying knowledge. Application AI makes specialists more powerful but can create an "AI-induced knowledge-debt trap": operational efficiency improves, but the organisation may defer investments in scalable infrastructure. This dichotomy aligns with empirical observations that early AI adopters in professional services report positive revenue and career outcomes, while many generative-AI pilots fail because organisations deploy tools without clear use-case specialisation. The model implies that organisations must balance efficiency gains against long-run learning. For nascent firms, the priority is to escape the stagnation equilibrium; for growth firms, to accelerate codification; for mature firms, to leverage AI for efficiency while maintaining innovation incentives.

5.3 Comparison with Ide and Talamas (2025)

The preceding discussion has emphasised how two archetypes of AI—improvements to routine interpolation ($q \uparrow$) and reductions in expert operating cost ($\eta \downarrow$)—shape the internal dynamics of the Babbage–Firm. A useful benchmark for interpreting these results is provided by Ide and Talamas (2025), who develop an equilibrium model of knowledge hierarchies

with human and machine agents. That model differs in structure and predictions from the present framework. In this subsection, we briefly compare the key mechanisms.

Ide and Talamas model firms as competitive knowledge hierarchies populated by heterogeneous human workers and a supply of machine agents produced from compute. Their distinction between *autonomous* and *non-autonomous* AI rests on whether the machines can both advise and produce output or only assist humans. The equilibrium determines who solves routine tasks and who handles exceptions, with wages adjusting to clear the labour market. By contrast, the Babbage–Firm is a single-organisation dynamic programme in which a manager allocates a fixed expert between exploitation and exploration. There is no market for labour or compute: AI enters solely by changing the parameters of the cognitive technologies (η and q), and the organisational structure evolves endogenously over the life cycle.

In Ide and Talamas's model, *basic autonomous* AI substitutes for routine workers and reassigns marginal human workers into problem-solving roles, creating a top-heavy hierarchy with many human solvers. When *advanced autonomous* AI matches the expertise of human solvers, machines take over exception handling, pushing those solvers down into routine production and flattening the human hierarchy. In the present paper, the analogue to autonomous AI is *interpolation AI*: by raising q, it expands the scope of routine work, reallocates expert time into research, and accelerates the transition to a codified organisation. Conversely, *application AI* in the Babbage–Firm lowers η and makes exploitation cheaper; this strengthens expert operations, slows codification, and can entrench specialist layers, in contrast to the flattening predicted by advanced autonomy in Ide and Talamas.

Ide and Talamas distinguish between *workers* who perform routine tasks and *solvers* who handle exceptions. Under basic autonomy, AI displaces the workers and raises the value of human solvers; under advanced autonomy, the least knowledgeable humans rely on AI solvers and marginal human solvers are pushed into routine roles. With non-autonomous AI, coproductive tools disproportionately benefit the least knowledgeable workers. With a Babbage firm, by contrast, there are only two roles: routine interpolators and tacit experts who can also conduct research. Interpolation AI raises the precision of routine work and frees experts to invest in research; application AI increases expert productivity in operations and crowds out research. The latter implies an "AI-induced knowledge-debt trap" in which short-run efficiency gains delay the build-up of codified knowledge.

The two models, therefore, highlight different channels through which AI can flatten or thicken organisational structures. Ide and Talamas obtain flatter human hierarchies when machines replace solvers at the top; the model here obtains flattening when interpolation improves routine processing and pushes experts into research. Ide and Talamas stress a dis-

tributional trade-off in autonomy regulation (advanced AI benefits extremes of the knowledge distribution), whereas the Babbage Firm model emphasises a dynamic trade-off between exploitation and exploration. The comparison underscores that predictions about AI depend critically on whether AI substitutes for human expertise or complements it by raising the quality of routine cognition (Autor and Thompson, 2025).

6 Conclusion

This paper develops a dynamic model of the cognitive division of labour and organisational learning. The framework extends the Babbage-Firm to include a research–operations trade-off, deriving a performance trajectory that can be linear-concave or S-shaped. Comparative dynamics analysis shows that the S-curve's shape depends on cognitive technology parameters: higher expertise localisation (η) and better interpolation quality (q) raise the research hurdle and lower critical thresholds, increasing the likelihood of an acceleration phase. We then incorporate two archetypes of AI—Interpolation AI (raising q) and Application AI (lowering η)—and show that they have opposite effects on the performance trajectory. Interpolation AI promotes organisational exploration and accelerates learning, while Application AI boosts operational efficiency but risks crowding out exploration.

The analysis generates several predictions. Interpolation AI should flatten hierarchies and reallocate expert time to research; Application AI should improve short-term efficiency but may slow infrastructure building; start-ups that adopt capability-broadening AI will scale faster than those that focus solely on expert augmentation; and mature firms that invest in R&D-oriented AI may rejuvenate growth. These predictions align with emerging empirical evidence: AI-native start-ups report high revenue per employee, early AI adopters in professional services expect revenue and career gains, while most generative-AI pilots fail without a specialised use case, and older firms struggle more with AI adoption.

Future research should endogenise AI adoption choices, allowing firms to select not only labour inputs but also AI investments that target generalists or specialists. Another fruitful direction is to study multi-period dynamics where codified knowledge depreciates or is superseded, and to incorporate discovery-AI channels that lower the cost of research. Understanding how AI reshapes cognitive hierarchies and organisational boundaries remains a critical challenge for economics and management science.

Appendices

A Derivation of Integrated Costs

This appendix derives the integrated costs $C_{\text{Gap}}(s)$ and $C_{\text{Tail}}(Z)$.

A.1 Tail Cost $C_{\text{Tail}}(Z)$

In a tail, under Assumption 1, the expert dominates the routine worker. The cost at distance z is $\min\{\eta z^2, 1\}$. Let $X_E = 1/\sqrt{\eta}$.

If $Z \le X_E$: $C_{\text{Tail}}(Z) = \int_0^Z \eta z^2 dz = \frac{\eta Z^3}{3}$.

If $Z > X_E$:

$$C_{\text{Tail}}(Z) = \int_0^{X_E} \eta z^2 dz + \int_{X_E}^{Z} 1 dz = \frac{\eta X_E^3}{3} + (Z - X_E) = Z - \frac{2}{3\sqrt{\eta}}.$$

A.2 Gap Cost $C_{\text{Gap}}(s)$

By symmetry, we analyse the cost on $[0, s/2]$ and double the result.

Regime 1: Expert Dominated (ED) ($s \le X_c$).

$$C_{\text{Gap}}(s) = 2\int_0^{s/2} \eta x^2 dx = \frac{\eta s^3}{12}.$$

Regime 2: Mixed Allocation (MA) ($X_c < s \le X_R$). The crossover point is $r_0 = \frac{s}{1+\eta q s}$. The routine cost is $C_R(x) = \frac{x(s-x)}{qs}$. We calculate the integrated cost as the cost if the routine worker covered the entire segment, minus the gain from using the expert on $[0, r_0]$. $C_{\text{Gap}}(s) = 2\left(\int_0^{s/2} C_R(x)dx - \int_0^{r_0}(C_R(x) - \eta x^2)dx\right)$.

The first term is $2\int_0^{s/2} C_R(x)dx = \frac{s^2}{6q}$.

The second term involves the advantage $A(x) = C_R(x) - \eta x^2 = x/q - \beta x^2$. Since $A(r_0) = 0$, $\beta = 1/(qr_0)$.

$$\int_0^{r_0} A(x)dx = \frac{r_0^2}{2q} - \frac{\beta r_0^3}{3} = \frac{r_0^2}{2q} - \frac{r_0^2}{3q} = \frac{r_0^2}{6q}.$$

The total integrated cost is:

$$C_{\text{Gap}}(s) = \frac{s^2}{6q} - 2\left(\frac{r_0^2}{6q}\right) = \frac{s^2}{6q} - \frac{r_0^2}{3q}.$$

Substituting r_0:

$$C_{\text{Gap}}(s) = \frac{s^2}{6q}\left[1 - \frac{2}{(1+\eta q s)^2}\right].$$

Partial Derivatives (MA Regime). $\frac{\partial C_{\text{Gap}}}{\partial \eta} = \frac{s^2}{6q} \cdot \frac{4(qs)}{(1+\eta q s)^3} = \frac{2s^3}{3(1+\eta q s)^3} > 0$. $\frac{\partial C_{\text{Gap}}}{\partial q} < 0$.

B Two–Period Case under Broad Expertise

For completeness, we outline the analysis of the two–period model when the cost parameter is *below* the localised expertise threshold, that is $\eta < 1/(4q^2)$ (Broad Expertise). In this case, the expert's reach $X_E = 1/\sqrt{\eta}$ exceeds $2q$. Let $S_E = 2X_E$ be the expert coverage limit, which exceeds the routinisation limit $X_R = 4q$. Because $X_c = 1/(\eta q)$ exceeds S_E, the mixed allocation (MA) regime never arises; the expert always dominates the routine worker until the expert's own coverage limit is exceeded.

The routine worker is never optimally used in interpolation. The gap cost function is:

$$C_{\text{Gap}}(s) = \begin{cases} \frac{\eta s^3}{12}, & 0 \le s \le S_E \text{ (Expert-Dominated, ED)}, \\ s - \frac{4}{3}X_E, & s > S_E \text{ (Unserved Centre)}. \end{cases}$$

The manager minimizes $f(s) = C_{\text{Gap}}(s) + C_{\text{Tail}}\big(L/2 - s\big)$. The analysis of the FOC yields:

$$s^* = \begin{cases} L/3, & L \le 3S_E, \\ \text{any } s \in [S_E,\, L/2 - X_E], & L > 3S_E. \end{cases}$$

When $L > 3S_E$, the cost function is flat, allowing for 'leap-frog' investments, which are ruled out under localised expertise.

The discounted hurdle rate is $H^* = \delta \Delta W^*/L$, where $\Delta W^* = C_{\text{Tail}}(L/2) - f(s^*)$. The calculation depends critically on whether the initial tails ($L/2$) are covered by the expert (i.e., $L/2 \le X_E$ or $L \le S_E$). This distinction leads to three regimes for H^*:

Regime 1: Fully Covered ($L \le S_E$). The initial state and the post-research state ($s^* = L/3$) are fully covered (ED).

$$H_1^* = \frac{\delta \eta L^2}{27}.$$

Regime 2: Intermediate ($S_E < L \leq 3S_E$). The initial tails are unserved ($L/2 > X_E$), but the post-research state is anchor-dominated ($s^* = L/3$) and fully covered (ED).

$$H_2^* = \delta \left(\frac{1}{2} - \frac{2X_E}{3L} - \frac{\eta L^2}{216} \right).$$

Regime 3: Coverage-Dominated ($L > 3S_E$). The initial tails are unserved, and the post-research state is coverage-dominated.

$$H_3^* = \frac{4\delta X_E}{3L} = \frac{4\delta}{3L\sqrt{\eta}}.$$

We analyse the comparative statics of organisational exploration $\rho = P(A_1(s) \leq H^*)$.

In Regime 1 (Fully Covered), H_1^* increases in η and L. An increase in η reduces the expert's advantage $A_1(s)$ while simultaneously increasing H^*; both effects raise the incentive to explore, so ρ increases in η. An increase in L increases H^*, and ρ also increases in L.

In Regime 3 (Coverage-Dominated), H_3^* decreases in η and L. As L increases, H^* falls. However, the exploitation region (where $A_1(s) > H^*$) remains bounded near the anchor, while the domain L expands. The regions far from the anchor are unserved ($A_1(s) = 0$), guaranteeing exploration. Since the exploitation region becomes a vanishingly small fraction of the total domain, ρ rises in L (approaching 1 as $L \to \infty$). As η increases, $A_1(s)$ decreases (favouring exploration) while H^* decreases (favouring exploitation). Detailed analysis confirms that the direct effect on $A_1(s)$ dominates, so ρ rises in η.

This analysis shows that when the cost parameter is small, the organisation may optimally establish widely spaced infrastructure and forgo local interpolation altogether. Such non-local expansions complicate the infinite-horizon problem and are excluded by the localised expertise assumption used in Section 3.2.

C Proofs of Lemmas 1 and 2

C.1 Proof of Lemma 1

We aim to find the split point $y \in (0, X)$ that maximises the split benefit. This is equivalent to minimising the post-split cost function:

$$K(y) = C_{\text{Gap}}(y) + C_{\text{Gap}}(X - y).$$

We take as given that C_{Gap} is strictly convex on $(0, 4q)$.

First, we find the optimal split point y. We examine the derivatives of $K(y)$:

$$K'(y) = C'_{\text{Gap}}(y) - C'_{\text{Gap}}(X - y),$$
$$K''(y) = C''_{\text{Gap}}(y) + C''_{\text{Gap}}(X - y).$$

Since C_{Gap} is strictly convex, $C''_{\text{Gap}}(s) > 0$ for $s \in (0, 4q)$. Thus, $K''(y) > 0$, which means $K(y)$ is strictly convex.

The first-order condition for an interior minimum is $K'(y) = 0$, which implies:

$$C'_{\text{Gap}}(y) = C'_{\text{Gap}}(X - y).$$

Because C_{Gap} is strictly convex, its derivative C'_{Gap} is strictly increasing. Therefore, the equality holds if and only if the arguments are equal: $y = X - y$. This yields the unique solution $y = X/2$. Since $K(y)$ is strictly convex, this unique critical point is the global minimum. Thus, splitting the gap at its midpoint maximises the split benefit.

Second, we examine the monotonicity of the maximised split benefit:

$$\Delta_{\text{Gap}}(X) = C_{\text{Gap}}(X) - 2\, C_{\text{Gap}}\big(X/2\big).$$

We differentiate with respect to X:

$$\begin{aligned}\Delta'_{\text{Gap}}(X) &= C'_{\text{Gap}}(X) - 2 \cdot C'_{\text{Gap}}\big(X/2\big) \cdot \frac{1}{2} \\ &= C'_{\text{Gap}}(X) - C'_{\text{Gap}}\big(X/2\big).\end{aligned}$$

Since C'_{Gap} is strictly increasing and $X > X/2$ (for $X > 0$), it follows that $C'_{\text{Gap}}(X) > C'_{\text{Gap}}\big(X/2\big)$. Therefore, $\Delta'_{\text{Gap}}(X) > 0$, which proves that the split benefit is strictly increasing in X on $(0, 4q)$.

C.2 Proof of Lemma 2

We aim to maximise the split benefit, which is equivalent to minimising the post-split cost function:

$$K(d) = C_{\text{Gap}}(d) + C_{\text{Tail}}(Z - d).$$

Since C_{Gap} is strictly convex and C_{Tail} is convex, $K(d)$ is strictly convex. The unique optimum d^* is characterized by the First-Order Condition (FOC): $K'(d) = 0$, or $C'_{\text{Gap}}(d) = C'_{\text{Tail}}(Z - d)$.

We first establish the derivatives of the cost functions. The tail cost derivative is:

$$C'_{\text{Tail}}(u) = \begin{cases} \eta u^2, & u \le X_E \text{ (Expert-Dominated, ED)}, \\ 1, & u \ge X_E. \end{cases}$$

Note that $C'_{\text{Tail}}(X_E) = \eta X_E^2 = 1$.

In the ED regime (when gap length $X \le X_c$), the gap cost is $C_{\text{Gap}}(X) = \eta X^3/12$, and its derivative is $C'_{\text{Gap}}(X) = \eta X^2/4$. Under localised expertise, $X_E \le X_c$.

Part (i): Optimal split $d^*(Z)$.

1. *Case 1: $Z < X_E$ (Short Tail).* If $Z < X_E$, then $d < X_E$ and $Z - d < X_E$. Since $X_E \le X_c$, both the resulting gap d and residual tail $Z - d$ are in the ED regime. We use the ED forms for both derivatives in the FOC:

   $$\frac{\eta d^2}{4} = \eta (Z - d)^2.$$

   Since $\eta > 0$, $d > 0$, and $Z - d > 0$, we take the square root: $d/2 = Z - d$. This yields $d^*(Z) = 2Z/3$.

2. *Case 2: $Z \ge X_E$ (Longer Tails).* The function $C'_{\text{Tail}}(Z - d)$ has a kink where the residual tail equals X_E, i.e., at $d_k = Z - X_E$. We evaluate the derivative of $K(d)$ at this kink:

   $$K'(Z - X_E) = C'_{\text{Gap}}(Z - X_E) - C'_{\text{Tail}}(X_E) = C'_{\text{Gap}}(Z - X_E) - 1.$$

   - *Subcase 2a: $Z \ge \bar{s} + X_E$ (Long Tail).* Here, $Z - X_E \ge \bar{s}$. Since C'_{Gap} is strictly increasing and $C'_{\text{Gap}}(\bar{s}) = 1$, we have $C'_{\text{Gap}}(Z - X_E) \ge 1$. Thus, $K'(Z - X_E) \ge 0$. By strict convexity of $K(d)$, the minimum must occur at $d^* \le Z - X_E$. In this range, the residual tail $Z - d^* \ge X_E$, so $C'_{\text{Tail}}(Z - d^*) = 1$. The FOC is $C'_{\text{Gap}}(d) = 1$. The unique solution is $d^*(Z) = \bar{s}$.
   - *Subcase 2b: $X_E \le Z < \bar{s} + X_E$ (Intermediate Tail).* Here, $Z - X_E < \bar{s}$. Thus, $C'_{\text{Gap}}(Z - X_E) < 1$, which implies $K'(Z - X_E) < 0$. By strict convexity, the minimum must occur at $d^* > Z - X_E$. In this range, the residual tail $Z - d^* < X_E$ (ED regime).

   The FOC is $C'_{\text{Gap}}(d) = C'_{\text{Tail}}(Z - d) = \eta(Z - d)^2$. A unique solution $d^*(Z)$ exists because the LHS is strictly increasing and the RHS is strictly decreasing in d. Furthermore, $Z - X_E < d^*(Z) < \bar{s}$.

Part (ii): Maximal Tail Dividend. Let $\Delta(Z)$ be the maximized split benefit: $\Delta(Z) = C_{\text{Tail}}(Z) - K(d^*(Z))$. We analyse how $\Delta(Z)$ evolves with Z. By the Envelope Theorem, the derivative of the minimised cost $K(d^*(Z))$ with respect to Z is $C'_{\text{Tail}}(Z - d^*(Z))$. Therefore:

$$\Delta'(Z) = C'_{\text{Tail}}(Z) - C'_{\text{Tail}}(Z - d^*(Z)).$$

If $Z < \bar{s} + X_E$. We analyse the regimes. If $Z < X_E$ (Short Tail), both Z and $Z - d^*$ are in the region where C'_{Tail} is strictly increasing. Since $d^* > 0$, $\Delta'(Z) > 0$. If $X_E \leq Z < \bar{s} + X_E$ (Intermediate Tail), then $C'_{\text{Tail}}(Z) = 1$. We established $Z - d^* < X_E$, so $C'_{\text{Tail}}(Z - d^*) < 1$. Thus, $\Delta'(Z) > 0$. The dividend is strictly increasing.

If $Z \geq \bar{s} + X_E$ (Long Tail). $d^* = \bar{s}$. Both Z and $Z - d^*$ are $\geq X_E$. Thus, $C'_{\text{Tail}}(Z) = 1$ and $C'_{\text{Tail}}(Z - d^*) = 1$. $\Delta'(Z) = 1 - 1 = 0$. The dividend is constant.

The maximum dividend is therefore achieved when $Z \geq \bar{s} + X_E$. In this regime:

$$\bar{\Delta}_{\text{Tail}} = C_{\text{Tail}}(Z) - [C_{\text{Gap}}(\bar{s}) + C_{\text{Tail}}(Z - \bar{s})].$$

Since $C_{\text{Tail}}(u)$ is linear with slope 1 for $u \geq X_E$, the difference $C_{\text{Tail}}(Z) - C_{\text{Tail}}(Z - \bar{s})$ equals the length of the interval, $\bar{s}$. Therefore, $\bar{\Delta}_{\text{Tail}} = \bar{s} - C_{\text{Gap}}(\bar{s})$.

D Dominance of tail dividends over gap dividends

For the localised expertise range $\eta \in [1/(4q^2),\ 1/q^2]$, this appendix proves that the maximal tail dividend exceeds the maximal interior gap dividend up to length $4q$; hence $X^\star > 4q$ in Proposition 2. We analyze the function $\Phi(\eta) := \bar{\Delta}_{\text{Tail}}(\eta) - \Delta_{\text{Gap}}(4q; \eta)$. We show $\Phi(\eta) > 0$ by evaluating it at the boundaries of the interval.

Case 1: Lower Bound $\eta_L = 1/(4q^2)$. At this boundary, $X_c = 4q$. We use the ED formulas (Appendix A).

Gap Dividend. $C_{\text{Gap}}(4q) = \eta_L(4q)^3/12 = 4q/3$. $C_{\text{Gap}}(2q) = \eta_L(2q)^3/12 = q/6$. $\Delta_{\text{Gap}}(4q) = 4q/3 - 2(q/6) = q$.

Tail Dividend. $C'_{\text{Gap}}(s) = \eta_L s^2/4$. Setting this to 1 yields $\bar{s} = 4q$. $\bar{\Delta}_{\text{Tail}} = \bar{s} - C_{\text{Gap}}(\bar{s}) = 4q - 4q/3 = 8q/3$.

Comparison. At $\eta = 1/(4q^2)$ we have $\bar{s} = 4q$. The maximal tail dividend is $\bar{\Delta}_{\text{Tail}} = 8q/3$, while the gap dividend is $\Delta_{\text{Gap}}(4q) = q$. Hence $\Phi\big(1/(4q^2)\big) = \frac{8q}{3} - q = \frac{5q}{3} > 0$.

Case 2: Upper Bound $\eta_H = 1/q^2$. At this boundary, $X_c = q$. The gaps $4q$ and $2q$ are in the MA regime.

Gap Dividend. We use the MA cost formula. $1 + \eta_H q(4q) = 5$; $1 + \eta_H q(2q) = 3$. $C_{\text{Gap}}(4q) = \frac{16q^2}{6q}[1 - \frac{2}{25}] = \frac{184q}{75} \approx 2.453q$. $C_{\text{Gap}}(2q) = \frac{4q^2}{6q}[1 - \frac{2}{9}] = \frac{14q}{27} \approx 0.519q$. $\Delta_{\text{Gap}}(4q) = C_{\text{Gap}}(4q) - 2C_{\text{Gap}}(2q) \approx 1.416q$.

Tail Dividend. We find $\bar{s}$ by solving $C'_{\text{Gap}}(\bar{s}; \eta_H) = 1$ in the MA regime. This leads to a polynomial equation whose unique positive root is $\bar{s} \approx 3.086q$. The tail dividend $\bar{\Delta}_{\text{Tail}} = \bar{s} - C_{\text{Gap}}(\bar{s}) \approx 1.689q$.

Comparison. $\Phi(\eta_H) \approx 1.689q - 1.416q > 0$.

Since $\Phi(\eta)$ is positive at both boundaries and the functions are smooth, this confirms that $\bar{\Delta}_{\text{Tail}} > \Delta_{\text{Gap}}(4q)$ throughout the localised expertise region.

E The Bellman Equation and the SVA Approximation

This appendix details the exact dynamic programming formulation and the Stationary Value Approximation (SVA) used in the main text.

E.1 The Exact Dynamic Value of Research

The continuation value $W(\mathcal{F})$ is defined by the Bellman equation, representing the NPV of expected payoffs under the optimal policy. Let $V(\mathcal{F})$ be the expected current payoff. Let $\rho(\mathcal{F})$ be the probability of exploration, and $\mathcal{F}'$ the subsequent state if research occurs.

The Bellman equation is:

$$W(\mathcal{F}) = V(\mathcal{F}) + \delta[\rho(\mathcal{F})W(\mathcal{F}') + (1 - \rho(\mathcal{F}))W(\mathcal{F})]$$

Rearranging gives the flow value equation:

$$(1 - \delta)W(\mathcal{F}) = V(\mathcal{F}) + \delta\rho(\mathcal{F})\Delta W(\mathcal{F})$$

where $\Delta W(\mathcal{F}) = W(\mathcal{F}') - W(\mathcal{F})$ is the incremental value of research. The exact Hurdle Rate is $H^*(\mathcal{F}) = \delta\Delta W(\mathcal{F})$.

To find the exact value of $\Delta W(\mathcal{F})$, we examine the difference in flow values between $\mathcal{F}'$ and $\mathcal{F}$. Let $\Delta V(\mathcal{F}) = V(\mathcal{F}') - V(\mathcal{F}) = \Delta^*(\mathcal{F})/L$ be the static gain. Subtracting the flow value equations yields:

$$(1 - \delta)\Delta W(\mathcal{F}) = \Delta V(\mathcal{F}) + \delta \underbrace{[\rho(\mathcal{F}')\Delta W(\mathcal{F}') - \rho(\mathcal{F})\Delta W(\mathcal{F})]}_{\mathcal{R}(\mathcal{F})}.$$

The exact incremental value is composed of the static gain $\Delta V(\mathcal{F})$ plus the change in the expected value of future research opportunities, $\mathcal{R}(\mathcal{F})$. This equation highlights the complexity of the exact solution, as ΔW depends on future values and the endogenous exploration ρ, while ρ itself depends on ΔW.

E.2 The Stationary Value Approximation (SVA)

The Stationary Value Approximation (SVA) assumes that the expected value of future research opportunities remains unchanged by the research conducted today. Formally, it imposes $\mathcal{R}(\mathcal{F}) = 0$.

Under this assumption, the dynamic equation simplifies to:

$$(1-\delta)\Delta W^{SVA}(\mathcal{F}) = \Delta V(\mathcal{F})$$

$$\Delta W^{SVA}(\mathcal{F}) = \frac{\Delta V(\mathcal{F})}{1-\delta}$$

This yields the SVA Hurdle Rate used in the text (Equation (2)):

$$H^{SVA}(\mathcal{F}) = \delta \Delta W^{SVA}(\mathcal{F}) = \frac{\delta}{1-\delta}\Delta V(\mathcal{F}).$$

E.3 Bias of the SVA

The bias is determined by the sign of $\mathcal{R}(\mathcal{F})$. In this model, the system is non-stationary.

Refinement Phase: During refinement, both ΔW (as Δ^* decreases) and ρ decrease over time. Thus $\mathcal{R}(\mathcal{F}) < 0$. The SVA overestimates the value of research during Refinement ($H^{SVA} > H^*$).

Expansion Phase: During expansion, ΔW is approximately constant (as Δ^* is constant), but ρ evolves.

- Concave Regime (Section 4): ρ decreases. $\mathcal{R}(\mathcal{F}) < 0$. SVA overestimates.
- S-shaped Regime: ρ increases. $\mathcal{R}(\mathcal{F}) > 0$. SVA underestimates.
- Linear Regime: ρ is constant. $\mathcal{R}(\mathcal{F}) = 0$. SVA is exact.

This confirms the interpretation provided in Sections 3.3.2 and 4.

F Dynamics of Exploration and Asymptotic Analysis

This appendix analyses the exploration probability $\rho(X;H)$ and details the asymptotic behaviour during the refinement phase. Let $\beta(X) = \eta + 1/(qX)$ and $R(X;H) = \sqrt{1 - 4q^2\beta(X)H}$. The roots solving $A_X(r_\pm) = H$ are $r_\pm = \frac{1}{2q\beta}(1 \pm R)$.

We must distinguish two cases based on whether $r_+ \leq X/2$. This condition is equivalent to $R \leq \eta q X$, or $H \geq A(X/2)$.

1. Case 1: $H \geq A(X/2)$ (High H). The exploitation region in $[0, X/2]$ is (r_-, r_+).

$$\rho(X;H) = 1 - \frac{2(r_+ - r_-)}{X} = 1 - \frac{2R}{1 + \eta q X}.$$

2. Case 2: $H < A(X/2)$ (Low H). The exploitation region in $[0, X/2]$ is $(r_-, X/2]$. The exploration region is $[0, r_-]$.

$$\rho(X;H) = \frac{2r_-}{X} = \frac{1 - R}{1 + \eta q X}.$$

The analysis of the derivative $\frac{\partial \rho}{\partial X}$ is complex because the functional form changes and depends explicitly on H. The claim in the main text that the monotonicity switch occurs exactly at X_c regardless of H is an approximation; the precise location of the minimum $X^*(H)$ varies with H.

In the long run refinement phase, $X \to 0$ and $H \to 0$. Since $A(X/2) \sim O(X)$ and $H \sim O(X^3)$, we are in Case 2 ($H < A(X/2)$). The asymptotic analysis in the text utilises the formula from Case 2.

References

Acemoglu, D., P. Aghion, C. Lelarge, J. Van Reenen, and F. Zilibotti (2007): "Technology, information, and the decentralization of the firm," *Quarterly Journal of Economics*, 122, 1759–1799.

Agrawal, A., J. S. Gans, and A. Goldfarb (2024a): "The Turing Transformation: Artificial Intelligence, Intelligence Augmentation, and Skill Premiums," *Harvard Data Science Review*.

——— (2025): "The Economics of Bicycles for the Mind," Tech. rep., National Bureau of Economic Research.

Agrawal, A., J. McHale, and A. Oettl (2024b): "Artificial intelligence and scientific discovery: A model of prioritized search," *Research Policy*, 53, 104989.

Autor, D. (2014): "Polanyi's paradox and the shape of employment growth," Tech. rep., National Bureau of Economic Research.

Autor, D. and N. Thompson (2025): "Expertise," *Journal of the European Economic Association*, jvaf023.

Bardhi, A. and S. Callander (2026): "Learning in a Correlated World," *Annual Review of Economics*, 51624, 072515.

Becker, G. S. (1964): *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, University of Chicago Press.

Callander, S. (2011): "Searching and learning by trial and error," *American Economic Review*, 101, 2277–2308.

Callander, S. and T. S. Clark (2017): "Precedent and doctrine in a complicated world," *American Political Science Review*, 111, 184–203.

Calvo, G. A. and S. Wellisz (1978): "Hierarchy, Ability, and Income Distribution," *Journal of Political Economy*, 86, 991–1010.

Carnehl, C. and J. Schneider (2025): "A Quest for Knowledge," *Econometrica*, 93, 623–659.

Gans, J. S. (2025a): "A Quest for AI Knowledge," Tech. rep., National Bureau of Economic Research.

——— (2025b): "A Simplified Quest for Knowledge," Tech. rep., National Bureau of Economic Research.

Garicano, L. (2000): "Hierarchies, markets, and the limits of knowledge-based organizations," *Journal of Economic Literature*, 38, 874–904.

Garicano, L. and T. N. Hubbard (2006): "The organization of knowledge in small firms," *Quarterly Journal of Economics*, 123, 1349–1384.

Garicano, L. and E. Rossi-Hansberg (2004): "Inequality and the Organization of Knowledge," *American Economic Review*, 94, 197–202.

Ide, E. and E. Talamas (2025): "Artificial intelligence in the knowledge economy," *Journal of Political Economy*, forthcoming.

Lucas, R. E. (1978): "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 9, 508–523.

March, J. G. (1991): "Exploration and exploitation in organizational learning," *Organization science*, 2, 71–87.

Rosen, S. (1982): "Authority, Control, and the Distribution of Earnings," *Bell Journal of Economics*, 13, 311–323.

——— (1983): "Specialization and Human Capital," *Journal of Labor Economics*, 1, 43–49.

Roy, A. D. (1951): "The distribution of earnings and of individual output," *The Economic Journal*, 60, 489–505.