The Pollinators

How specialization deepened knowledge—and why we need a better way to connect it

Éditions Maison Cœur · Editorial working draft · 1 November 2026

The following is new continuous prose, not the earlier 278-page blank manuscript template. It includes a prologue and a first chapter. Historical and conceptual claims are referenced at the end. Fictional or composite scenes are explicitly identified; no unconducted expert interview is represented as testimony. This is an editorial draft, not a final published book.

— — —

PROLOGUE — The Question Nobody Owned

Imagine a city council meeting at which everyone has done their homework. The engineer can explain why a bridge needs repairs. The financier knows what happens if the work is delayed and what else the budget might fund. A public-health specialist can describe how a diversion may affect access to a hospital. A neighborhood representative knows which routes elderly people actually use. The communications team understands which explanation will survive a thirty-second news clip. They disagree, but nobody in the room is foolish, lazy, or acting in bad faith.

This meeting is invented. Its purpose is to expose a familiar difficulty: each participant may bring excellent knowledge about a portion of the decision while the most important question belongs to no single participant. How should a city compare technical risk, time, public trust, and unequal burdens in one responsible judgment? Who notices that the options have been framed too narrowly? Who can translate the engineer's warning without turning it into a slogan, or the resident's experience without treating it as a statistically representative survey?

A conventional response is to ask for more expertise. Often that is precisely right. But the meeting may already contain a great deal of it. The problem is not necessarily that knowledge is absent. It may be that relevant knowledge is dispersed across people whose professional languages, incentives, and habits of proof do not align. They have different names for uncertainty, different meanings for urgency, different ideas about what counts as a decisive fact. The missing capacity is not another specialist to overpower the rest. It is a disciplined way of making different forms of knowledge answerable to a shared question.

This book calls that capacity *pollination*. The metaphor is attractive for its movement across fields, but it needs rules. A bee that carries pollen does not become the flower. A person who transports a concept from surgery into business does not thereby become a surgeon. A borrowed term can illuminate a neglected mechanism, or it can lend glamorous authority to an empty analogy. Pollination is not the collection of clever metaphors. It is the difficult practice of learning enough about another domain to ask better questions in one's own—and then knowing when the analogy stops working.

The history of organized knowledge has given us remarkable reasons to respect depth. Specialization is not a cultural mistake that a heroic generalist can reverse. Complex medical treatment, reliable aircraft,

scholarly translation, food safety, and modern computing depend on bodies of knowledge that few people can master in their entirety. The achievement is real. So is the cost of a world where the routes between expertise may remain weaker than the individual disciplines themselves.

In 1945, Friedrich Hayek argued that the economic challenge is not simply to compute a solution from facts assembled in one place: the relevant knowledge is scattered among people who have access to different circumstances [2]. His argument addressed economic coordination. It should not be casually promoted into a universal theory of institutions. Yet it exposes a puzzle that returns throughout this book. If important knowledge exists in many hands, the challenge becomes one of coordination and translation, not merely acquisition.

What would a decision look like if the people making it could say, plainly: *This is what our discipline sees. This is what it cannot see. This is the evidence we would need from yours. This is where our concepts fail when transferred.* Such a conversation would not end disagreement. It would make disagreement more informative. It would replace polite interdisciplinary theater with testable points of contact.

The chapters that follow will examine how skill becomes concentrated, how professional words compress experience, why institutions sometimes reward the maintenance of boundaries, and what happens when ideas cross those boundaries responsibly. There is also a darker possibility to confront: the same language of connection can be used to flatten expertise, evade accountability, and sell the promise of effortless understanding. In the age of generative AI, that temptation is especially powerful. A fluent answer may give us the sensation that knowledge has been integrated when it has merely been placed in the same paragraph.

We will therefore not treat a convincing synthesis as proof of understanding. We will ask whose judgment produced the claim, what experience supports it, what was lost in translation, and which objections would require us to revise the transfer. The ambition is not to give every reader mastery of everything. It is to make the world's separated kinds of competence more capable of meeting without surrendering their standards.

That is the pollinator's task: to cross a boundary while preserving the fact that a boundary exists.

— — —

CHAPTER ONE — The Library We Cannot Enter

1. A building made of other people's lives

Picture a library so large that no visitor could walk all its corridors. Its collections contain not only books but skills: how to recognize a deteriorating machine by its sound, how to detect a mistranslation in a diplomatic document, how to choose a safer route through an unfamiliar mountain valley, how to tell whether a clinical result is ordinary noise or a sign that demands attention. Some shelves hold formal theorems; others hold procedures, idioms, patient histories, unwritten workshop rules, and mistakes that cost someone years to understand.

The library is a metaphor, not an existing archive. Its strangest property is that a reader's access card works in only a few wings. Entering a field takes time. A competent practitioner has learned which failures matter, which distinctions are subtle, and which statements require proof. The rest of us can often borrow a definition, but a definition is not the same as the experience of deciding when it applies.

This difference is the first obstacle to horizontal wisdom. We imagine knowledge as information that can travel intact from one person's head to another. Some knowledge travels this way: a measurement, a citation, a clearly specified procedure. But much of a field's intelligence lives in judgments about *relevance*. Which signal should interrupt a routine? Which irregularity is harmless? When is a familiar solution dangerous because the surrounding conditions have changed? These are not always secrets that professionals refuse to share. They may be distinctions that a profession has learned to make before it has learned to explain them in ordinary words.

Michael Polanyi developed the idea of tacit knowing to describe dimensions of human competence that exceed explicit statements [3]. His point should not be reduced to the popular claim that experts merely have mysterious intuitions. Expertise can be partly articulated, tested, taught, and criticized. Yet the gap between performance and description matters. An expert may identify a pattern without being able to give a complete verbal reconstruction of every cue. A novice may understand the expert's words without yet possessing the discriminations that make those words useful.

We meet the same issue in language. A native speaker can use an expression appropriately in dozens of situations while finding it difficult to state its full pragmatic rule. A translator, conversely, may be alert to distinctions that ordinary speakers rarely formulate. A legal term and a similar everyday word may invite the same dictionary translation while carrying different obligations, thresholds, or consequences. The problem is not that language is inadequate. It is that language works inside communities of practice, and carrying a word outside its community can change what the word does.

A book about reconnecting knowledge must therefore begin with humility about access. The visible entrance of a field—a popular article, an executive summary, a useful analogy—is not the field itself. There is nothing wrong with starting at the entrance. Trouble begins when we mistake admission for completion.

2. The pin, the workshop, and the bargain

Adam Smith's famous account of pin manufacturing in *The Wealth of Nations* offers one of the clearest historical arguments for the productivity of division of labor. Writing in 1776, Smith described the making of a pin as a series of specialized operations rather than a task that one worker performed from beginning to end [1]. The precise quantities in his illustration belong to his historical account, not to a contemporary industrial benchmark. The important point is the mechanism: repeated attention to one operation can increase skill, reduce switching costs, and encourage improvements in tools.

The story is often told as if specialization were simply the opposite of a richer, integrated way of making things. That is too easy. Without the capacity to concentrate on a difficult subproblem, many forms of excellence would remain out of reach. A contemporary laboratory cannot simply ask every participant to become an equal master of all instruments and techniques. A software service cannot dispense with specialized understanding of security, infrastructure, design, regulation, and accessibility merely because its leaders prefer a holistic vocabulary. Specialization creates a form of collective power by making individual mastery possible at greater depth.

But Smith's example also has a detail that deserves attention. The pin is a finished object. Its operations must finally fit together. Specialization succeeds not only because people work on parts but because their work participates in a coordination system that delivers a whole. In the pin workshop, that whole is relatively easy to describe. Modern social decisions, research collaborations, and institutional reforms often lack such a simple object. What is the finished product when a university tries to improve students' judgment? What is the equivalent of a completed pin when a city tries to restore public trust? The answer depends on contested purposes, multiple measures of success, and the perspectives of people who do not always agree on what success means.

When there is no shared object, the division of intellectual labor can become a division of the problem itself. Each department completes a portion of the task according to its own standards. Reports accumulate, deadlines are met, and nobody is responsible for checking whether the reports can be combined into a coherent decision. The institution can be efficient in every room and confused in the corridor between them.

This is not a proof that specialization causes institutional failure. It is a hypothesis about a vulnerability that can emerge under certain organizational conditions. It must be tested against competing explanations: poor governance, political incentives, asymmetric power, limited resources, and simple disagreement about values. The case for reconnecting knowledge becomes stronger when we identify those conditions precisely rather than blame a single grand historical transformation.

3. The map is not the conversation

A common response to fragmentation is to draw a map. Create a directory of experts. Connect organizations through a knowledge graph. Build a dashboard in which every department has a node. These are useful tools. They can make relationships visible and lower the cost of finding a relevant perspective. But a map of expertise does not tell us how to resolve a disagreement over evidence, nor does an organizational chart reveal what each participant is actually able to contribute.

Consider an illustrative hospital design discussion. A spatial planner may propose a clearer entrance; an infection-control specialist may prioritize separations; a nurse may notice that a proposed route adds a difficult journey for patients; an administrator may worry about cost. A graph can record that these roles participate. It cannot decide, by itself, whether a shorter route has a higher or lower clinical risk in the circumstances. The valuable question is not merely *Who is connected to whom?* It is *Which observation, made from one practice, changes the decision made in another?*

Susan Leigh Star and James Griesemer studied cooperation among heterogeneous participants associated with a natural-history museum, describing how objects and practices can enable coordination across different social worlds [4]. Their idea of boundary objects offers a useful distinction. Cooperation does not always require everyone to adopt a single meaning of everything. It may require sufficiently stable shared materials that people can work together while retaining distinct purposes. A common form, a repository, or a carefully designed case record can give unlike practices a meeting point.

There is a warning here for technological enthusiasm. The availability of a shared interface does not guarantee a shared understanding. A form that asks every contributor to choose from one set of categories may coordinate their data while erasing the distinctions they would have made in their own field. A repository may preserve documents but not the tacit conditions under which someone considered them

reliable. Design decisions about categories, permissions, and evidence are therefore epistemic decisions as well as technical ones.

A better map would expose its own limits. It would distinguish an expert's actual participation from a keyword match, a verified source from a plausible summary, an observed failure from an interpreted cause, and a transferable mechanism from a merely suggestive analogy. It would allow participants to say: *There is a connection here, but we do not yet know what it means.*

4. Crossing without pretending

The term *interdisciplinary* often announces a desirable aspiration without telling us what work occurred. Two specialists speaking at the same event are not necessarily collaborating. An elegant slide with arrows between professions is not necessarily integration. Even a successful collaboration may depend on painstaking translation that cannot be repeated by copying its surface features.

The National Academies' report *Facilitating Interdisciplinary Research* identifies institutional conditions that can obstruct cross-disciplinary work, including evaluation and resource-allocation systems organized around traditional fields [5]. This matters because the walls between areas of knowledge are not only cognitive. They are also administrative and professional. What counts toward promotion? Who has authority to define a problem? Which publications are recognized? Who pays for the time spent learning the language of another field? A call for more openness will be ineffective when the institutions involved reward only narrowly legible outputs.

A second difficulty concerns authority. Harry Collins and Robert Evans have argued for taking expertise seriously while distinguishing it from democratic or political decision-making rights [6]. The distinction matters in practice. A community may have a legitimate right to decide which risks it will accept, even when it does not possess the technical knowledge needed to estimate those risks. An expert may understand a technical mechanism without possessing the sole authority to decide which social goals deserve priority. Horizontal wisdom should deepen this distinction, not dissolve it.

The pollinator is not a replacement expert, an inspirational speaker, or a universal problem solver. The role is better understood as a disciplined intermediary. A pollinator learns the structure of a problem, listens for words whose apparently familiar meanings conceal different standards, seeks mechanisms that might travel across settings, and brings the proposed transfer back to people capable of testing it. This work needs both linguistic sensitivity and intellectual restraint.

An illustrative example makes the difference clearer. Suppose a manager borrows the concept of *triage* from medicine to rank incoming business tasks. At the level of a loose metaphor, the transfer is banal: do urgent things first. A serious transfer asks what the term organizes in its home domain: time pressure, consequences of delay, available resources, decision thresholds, and moral responsibilities. It then asks which of these structures exist in the business setting and which do not. The manager may learn to distinguish urgency from importance. But if the comparison becomes a justification for neglecting people whose requests are less visible, the metaphor has concealed an ethical problem rather than solved one.

To pollinate is therefore to create an accountable sequence: identify a source-domain distinction; describe the mechanism; state the receiving problem; specify the conditions that would make the transfer useful; test the consequences; document the limits. This is slower than throwing a clever term into a meeting, but it produces something more valuable than admiration. It produces a question someone can answer.

5. Artificial fluency and the new corridor

Generative AI adds an unusual possibility to the story. It can help people locate terminology, restate technical explanations, compare plausible frameworks, and generate questions that a newcomer might otherwise never think to ask. It can also make the corridor between disciplines look shorter than it is. A coherent paragraph can bring words from medicine, ecology, operations, and political theory into close contact while leaving their evidential standards untouched.

The risk is not that every AI-generated comparison is false. Some comparisons will be useful. The risk is that fluency can hide the work still required: provenance checking, expert disagreement, cases of failed transfer, domain restrictions, and consequences of error. If a system gives five answers in polished English, the user may forget to ask whether any qualified person would defend the underlying inference. The interface has solved a communication problem while the epistemic problem remains open.

There is also a positive possibility. An intelligently designed research system could make uncertainty more visible rather than less. It could ask what is missing, what rival explanation fits the same evidence, whose domain competence would be needed to test an inference, and why a concept should not be transferred. It could preserve disagreement across disciplines as cumulative knowledge instead of compressing it into an apparently unanimous summary. But those are design objectives, not accomplishments one can assume from a language model's fluency.

A mature form of horizontal wisdom would exploit AI's reach without confusing access to text with access to competence. It would make expert attention more valuable, not less. The point would be to distribute better questions and create more productive encounters between people who know different things. That is the meaning of the library metaphor: its purpose is not to make every visitor the librarian of every room. Its purpose is to improve the passages between rooms while keeping their contents intelligible and trustworthy.

The library we cannot enter in full may still be a library in which we learn to find the right door, ask permission, discover what the inhabitants know, and bring back a carefully checked idea. That is a smaller ambition than omniscience. It is a more demanding one than collecting quotations. And it is a workable starting point for a civilization that has become very good at producing expertise and is still learning how to connect it.

— — —

Source notes and editorial status

[1] Adam Smith, *An Inquiry into the Nature and Causes of the Wealth of Nations* (1776), Book I, Chapter I, “Of the Division of Labour.” Primary text: https://en.wikisource.org/wiki/The_Wealth_of_Nations/Book_I/Chapter_1 . This is an eighteenth-century illustration, not evidence for present-day production quantities.

[2] F. A. Hayek, “The Use of Knowledge in Society,” *American Economic Review* 35, no. 4 (1945): 519–530. https://www.mercatus.org/sites/default/files/d7/the_use_of_knowledge_in_society_-_hayek.pdf . This is an economic argument about distributed information, here used only for an explicitly limited comparison.

[3] Michael Polanyi, *The Tacit Dimension* (1966), especially “Tacit Knowing.” Historical publication information and excerpt: https://tacit-knowledge-architecture.com/object/the-tacit-dimension-excerpt/ . A theory of tacit knowing is not proof that every expert judgment is inexpressible.

[4] Susan Leigh Star and James R. Griesemer, “Institutional Ecology, ‘Translations’ and Boundary Objects,” *Social Studies of Science* 19, no. 3 (1989): 387–420. https://doi.org/10.1177/030631289019003001 .

[5] National Academy of Sciences, National Academy of Engineering, and Institute of Medicine, *Facilitating Interdisciplinary Research* (Washington, DC: National Academies Press, 2005). https://doi.org/10.17226/11153 .

[6] Harry Collins and Robert Evans, “The Third Wave of Science Studies: Studies of Expertise and Experience,” *Social Studies of Science* 32, no. 2 (2002): 235–296. https://doi.org/10.1177/0306312702032002003 .

Editorial note: All city, hospital, and manager examples above are invented illustrations, not reported incidents. No interview quotations are included. The manuscript still needs the remaining chapters, further expert interviews, historical fact-checking, independent editorial review and publication clearance. It must not be advertised as a finished 278-page book.