UNIVERSITY OF WARSAW | FACULTY OF ECONOMIC SCIENCES

WORKING PAPERS
No. 30/2026 (524)

FROM INFORMATION TO COMMUNICATION: GENERATIVE INTERACTION IN THE PRODUCTION OF KNOWLEDGE

TOMASZ KOPCZEWSKI
JAN LISICKI

WARSAW 2026
ISSN 2957-0506

WORKING PAPERS 30/2026 (524)

UNIVERSITY OF WARSAW
FACULTY OF ECONOMIC SCIENCES
WORKING PAPERS

From Information to Communication: Generative Interaction in the Production of Knowledge

Tomasz Kopczewski[1], *Jan Lisicki*[1*]

[1] *University of Warsaw, Faculty of Economic Sciences*
[*] *Corresponding author: jt.lisicki@uw.edu.pl*

Abstract: Economic analysis of information explains how existing knowledge is distributed, revealed, and used, while the economics of innovation treats new ideas as outcomes of search and recombination. This paper examines a distinct role of communication in knowledge production: interaction can change the course of search itself. We define the mechanism as adaptive feedback, in which one participant responds to the current state of another participant's reasoning and thereby alters the directions subsequently explored. The resulting idea need not be contained in the feedback or previously possessed by either participant. In stronger cases, interaction can change the representation of the problem and make previously unavailable hypotheses accessible. The paper distinguishes this generative role of interaction from information transmission and access to others' knowledge. It argues that institutions shape knowledge production through both information structure and interaction structure, including the timing of interaction and the balance between independent search and mutual influence.

Keywords: communication, knowledge production, generative interaction, search, innovation, problem representation

JEL codes: D83, O31, O32

1. Introduction

1.1. Information: Transmission, Generation, and Communication

In the simplest case, communication changes how knowledge is distributed across agents. Person A knows a certain piece of information X, person B does not know it, and communication allows B to acquire it. After successful transmission, B now knows X, so the difference in knowledge between the participants has narrowed. Viewed more broadly, from the perspective of the entire environment, no new content has been created—the total stock of knowledge available in the environment has not changed. X existed prior to communication as part of A's knowledge, and the only thing that has changed is who has access to it. Such a case may be of fundamental importance for the agents' decisions, coordination, and actions, but the analysis then concerns the connections through which existing information is transmitted—nothing new is created. The pursuit of innovation or the advancement of science does not result solely from the accumulation of information. A researcher may read all the articles on a given topic and still fail to find a hypothesis that explains an observed anomaly. Similarly, an engineer may be familiar with all current technologies, but that is not enough to discover a new one. Likewise, a company may gather a large amount of market information and still not know what product to create or what business model to adopt. In such situations, the problem is not merely a matter of accumulating as much available data as possible—the data must also be critically analyzed. What is needed is a hypothesis, a way to test it, or a shift in the line of reasoning or in how the problem is framed—something that was not already available as a piece of knowledge held by any participant.

The distinction between transmitting and generating content also has a history in economics. Mirowski and Nik-Khah (2017) show that several important strands of postwar economics developed under the influence of approaches in which information could be measured, transmitted, and processed without analyzing its substantive meaning. This way of thinking is well suited to studying how existing information moves between agents and affects their decisions. It provides much less guidance on how a new hypothesis or a new way of describing a problem first becomes available to an agent.

Economics also contains a clear tradition of treating new ideas as something that must be produced. In Romer (1990), technological change results from deliberate activity in which resources allocated to research lead to the generation of new ideas. This allows us to distinguish knowledge production from its later distribution: before an idea can be communicated, applied, or disseminated, it must first be generated. Weitzman (1998) shows that this process can involve the recombination of elements that are already known. A new idea can therefore emerge from an existing body of knowledge even when no new information enters from outside. The amount of information available to an agent does not by itself explain which new hypothesis or solution will arise from it. This is why information transmission and information generation need to be treated as separate processes. In the first case, we can ask who possesses X, who receives it, and how accurately it is transmitted. In the second, we must explain how X comes into existence when it was previously absent from the participants' knowledge.

This distinction does not imply that communication is necessary for new knowledge to emerge. New content can arise for many reasons, and communication may or may not be one of them. An individual can generate a new hypothesis on their own by combining prior knowledge, recognizing a new analogy, or reformulating a problem without the involvement of others. Epistemic novelty and social interaction are therefore distinct phenomena. If every new idea were considered a product of communication simply because knowledge has a social context or draws on the prior achievements of others, it would become difficult to identify any distinct creative role for communication. The problem we are describing arises only after these two distinct processes are separated. Communication can convey content that participants already possess, and knowledge production can occur without communication. There is, however, a further possibility: interaction between participants may itself influence the process through which previously unavailable content is generated. If one person's response changes the question subsequently posed by another, directs attention to a different assumption, or leads to a reformulation of the problem, the outcome of the conversation need not be a simple sum of previously held information. The question posed by this article is therefore whether and how communication can contribute to the production of new knowledge, rather than merely conveying existing knowledge.

1.2. The Claim and the Line of Argument

The answer developed in this article is conditional: communication can contribute to the generation of new ideas and hypotheses. Communication is neither necessary nor sufficient for novelty, but critical interaction can alter the course of idea generation and thereby become part of the process through which new content is produced. The problem framed in this way goes beyond the standard framework of the economics of information, which is examined in the next section. Information economics is well suited to analyzing situations in which the information under study is already available—either to one of the participants or generally within the environment. An agent may not possess it, may not know the true state of the world, may receive signals of varying quality, may disclose private knowledge, or may draw on information dispersed among many people. In each of these cases, the fundamental question concerns what happens to the information already available within the system. A different problem arises when the set of hypotheses or categories that the agent can consider changes.

In the following sections, we will develop this argument step by step. First, we analyze what standard information economics can explain when the representation of the problem and the relevant set of possibilities are predefined. Next, we turn to the economics of innovation, where the production of new knowledge is already a central object of analysis. We then turn to the problem of *search*: if new hypotheses and solutions result from search and recombination, we must explain what guides the subsequent moves in this process. Against this backdrop, we isolate the specific role of communication between agents. Interaction may redirect search toward new answers and, in the stronger case, toward a new representation of the problem itself. Finally, we examine the economic implications: institutions can influence knowledge production by organizing interactions among participants. The claim requires an important qualification. Communication can broaden the scope of knowledge, but it can also narrow it through conformism, premature convergence among participants, or by hindering independent, individual exploration. The subject of analysis is therefore not the general benefit of "greater communication," but rather the conditions under which interaction alters the process of knowledge production. The starting point for further consideration is an analysis of a situation in which the body of available knowledge is closed and predefined.

2. Information in a Closed System

Before examining how communication can influence the emergence of new theories and inventions, we need to establish what the existing tools of information economics can already explain. The starting point is a situation in which an agent receives new information about the world while the set of relevant possibilities remains fixed and predefined. This information can significantly alter the agent's beliefs and decisions, but it does not change the overall state of knowledge within the environment. In this section, we therefore treat information as exogenously given and ask how it changes beliefs and decisions and how it is distributed or communicated among agents.

2.1. Learning Among Known Possibilities

The simplest case of acquiring new knowledge occurs when an agent is able to identify all available options from which to choose but does not know the true state of the world and cannot determine exactly which option is best for them. For example, an entrepreneur may not know whether demand for a new product will be high or low. They may assign probabilities to both scenarios, choose a production volume under this uncertainty, and then purchase a market forecast or conduct consumer research. New information may radically change the decision, but both before and after receiving it, the entrepreneur continues to consider exactly the same two states.

Savage (1954) provides a formal language for such a situation in the theory of decision-making. An agent chooses among actions whose consequences depend on an uncertain state of the world. States represent circumstances beyond the decision-maker's control; actions assign consequences to particular states; and beliefs can be represented by subjective probabilities. Decision uncertainty can be very high, but the available options must be sufficiently well-defined to allow beliefs to be assigned to them and the consequences of actions for the environment to be assessed. In a problem defined in this way, new information allows for a revision of prior beliefs. If an agent considers states H and L and then observes a signal s, the agent can update their prior beliefs by assigning different probabilities to the respective states. A strongly positive market research result may raise the probability of high demand from 0.3 to 0.8 and justify increased production of the good. The change in the agent's behavior itself, based on the newly acquired information, may therefore be significant, but it occurs only among the possibilities that were represented in the closed system at the outset.

A similar situation can be viewed from the perspective of information quality. Blackwell (1953) shows how to compare information structures in terms of how well they enable an agent to distinguish among possible states of the world and make decisions on that basis. Such a comparison, however, still assumes that the states between which the agent is choosing have already been specified. The quality of information may lead to a better choice, but the choice itself is defined in advance. The common conclusion is simple: an agent may obtain significantly better information and completely change their decision, but they are still learning about possibilities that were already part of the problem.

The examples above show how valuable new information can be even when the set of possibilities remains fixed. We now move from a single decision-maker to situations in which relevant knowledge is dispersed among different agents.

2.2. Information Dispersed Across Agents

When knowledge is dispersed among different agents—whether people or companies—its availability, the willingness to disclose it, and how it is used introduce an additional level of complexity. A seller may know something that the buyer does not; an employee may possess private knowledge relevant to the employer; market participants may have access to local information unavailable to any central decision-maker.

A prime example is Akerlof (1970), in which the seller knows the quality of a used car better than the potential buyer. Since the buyer cannot reliably distinguish between high- and low-quality cars before the purchase, the buyer's willingness to pay reflects the average expected quality. This may also lower the price offered to owners of good cars and drive them out of the market. The source of the problem is the unequal distribution of information that already exists.

A more difficult case arises when private knowledge cannot be reliably communicated directly. In Spence (1973), the employer does not directly observe a candidate's productivity but can observe their level of education. Since obtaining an education is less costly for more productive individuals than for less productive ones, the very decision to pursue it may reveal information about the candidate's productivity. In Spence's model, education can thus serve as a signal of productivity, conveying information to the employer without direct (and potentially unreliable) communication.

The problem of dispersed knowledge can be far more complex than simply transmitting a signal between the two parties to a transaction. Hayek (1945) emphasizes that information about resources, demand, production capabilities, and local conditions is dispersed among many individuals. A change affecting the supply of a good may be reflected in its final price, and market participants must adjust their decisions without full knowledge of what caused the fluctuation. In this way, price serves as a tool that enables the coordination of actions based on knowledge that remains decentralized.[1] It is hard to find a more vivid example of the aggregation of information in a price than the recent confrontation between the U.S. and Iran, during which the rapid pace of events led to frequent fluctuations in the price of oil. Simply tracking these price fluctuations was enough to determine, with a fairly high degree of certainty, the current state of affairs in the Strait of Hormuz.

It is not only price that can aggregate data; observing social behavior—which can itself be a carrier of information—can have a similar effect. Another person's actions may reveal something about their private information, even if the underlying signal remains unknown. However, information cascades can work in both directions: they can convey information, but they can also reduce the amount of independent private information revealed by subsequent actions. If the information contained in the observations and choices of other participants outweighs a given agent's private signal, the agent may make a rational decision to imitate others. When subsequent individuals act in the same way, their actions cease to contribute new private information to the aggregate state of knowledge—they act as a herd (Bikhchandani et al., 1992).

These models show how economics can analyze access to information, its reliability, indirect revelation, and aggregation when the relevant information already exists within the system. They do not exhaust the problem, however, because agents may also communicate directly and strategically. The next step is therefore to examine communication within a problem whose basic elements remain predefined.

2.3. Communication within a predefined problem

Information economics is not limited to situations in which information flows unidirectionally from one agent to another. There are also models of direct communication in which participants exchange messages, decide what information to disclose, and react to what others know or say.

Economics can analyze not only what individual agents know, but also what they know about what others know. The concept of common knowledge describes a situation in which a specific piece of information is known to all participants; everyone knows that the others know it, everyone also knows that the others know that the others know it, and so on. Aumann (1976) shows that if two people start with the same initial beliefs and their subsequent probability assessments of a specific event become common knowledge, then these assessments must be

[1] The interpretation of Hayek's later claims about markets and information is disputed. Mirowski and Nik-Khah (2017) place them within a broader development in which markets came to be treated as information-processing institutions, while Caldwell (2017) questions this interpretation. We use the narrower point that prices can coordinate actions while the knowledge underlying individual decisions remains dispersed.

consistent, even if they were previously based on different information. A situation of complete awareness regarding the level of information held by other participants is not self-evident. This is illustrated by the “Electronic Mail Game” of Rubinstein (1989), in which one of two players receives information relevant to their joint decision and forwards it to the other via email. The recipient automatically sends a confirmation; subsequent messages are also confirmed, but any of them may be lost. Each confirmation received increases the players’ certainty about what the other person knows; however, it does not provide complete certainty that both know the same thing and know that the other also knows it. The model thus shows that prolonged communication and numerous confirmations can alter the participants’ mutual knowledge without leading to common knowledge.

Information disclosure is not automatic—it is often a strategic choice by the participants. In the “cheap talk” model of Crawford and Sobel (1982), one party possesses information relevant to the other party’s decision and can convey it through a non-binding message. Since the participants’ interests may not be fully aligned, the recipient cannot simply assume that every statement is true. Nevertheless, communication can convey some information if the parties’ interests are close enough that it is in the sender’s interest to reveal something about the knowledge they possess, but the receiver should take the sender’s motivations into account when interpreting the message. A related framework is the *Bayesian persuasion* model of Kamenica and Gentzkow (2011), in which the sender can choose in advance how information about the actual state of the world will be disclosed to the receiver. The receiver interprets the information received and makes a decision based on it, while the sender chooses a method of disclosing information that increases the likelihood of a response from the receiver that is favorable to the sender. The model thus shows that the object of strategic choice may be not only the content of a single message but also the disclosure rule itself. Dewatripont and Tirole (2005) add one more element: effective communication requires the involvement of both the sender and the receiver. The sender must expend effort to prepare the message so that the information is understandable, and the receiver must devote attention to receiving and processing it. Whether the information is effectively conveyed therefore depends on whether it is worthwhile for both parties to expend the effort associated with communication.

The models discussed above differ substantially in what they treat as information and in the problems they are designed to explain. Mirowski and Nik-Khah (2017) emphasize this diversity in the development of postwar information economics. For the argument here, this makes one shared feature especially important. Across the models considered in this section, the basic elements of the problem—possible states, relevant types, available actions, or the information that can be disclosed—are specified before communication begins. New information can change beliefs and decisions, private knowledge can become observable, and communication can affect what agents learn from one another. Throughout these processes, however, the main elements from which the problem is constructed are already available. This leaves open a different question:

what happens when the set of possibilities that an agent can formulate changes during the process itself?

2.4. The Role of Unawareness

The preceding models allow information to change and beliefs to be revised while the relevant possibilities remain representable from the outset. A different problem arises when the agent cannot yet formulate one of those possibilities. An agent may initially consider possibility A to be unlikely, but after receiving information, change their mind and regard it as almost certain. Their decision may be completely reversed, but possibility A was, from the outset, something they were able to articulate and distinguish from alternatives. Another type of change occurs when a new category emerges that the agent had not previously considered possible. In such cases, not only can the information available to the agent or their beliefs change, but the set of actions the agent can formulate may change as well—a completely different system dynamic may emerge.

This distinction has been formalized in the literature on unawareness. If an agent considers event C to be extremely unlikely, C must already be available to them as an event to which a probability can be assigned. Unawareness is a situation in which a certain possibility or choice does not fall within the scope of defined events at all. Dekel et al. (1998) show that standard state-space models encounter a fundamental difficulty in representing unawareness understood in this way. Therefore, it is crucial to distinguish between two perspectives: the state space, which describes the entire world, and the representation of that world that the agent operates with. Upon discovering a new category of reality, we can retrospectively describe the environment by expanding the space of possibilities. However, this does not affect the agent's earlier line of reasoning, as the agent could not have incorporated such information into their decision-making process. In this framework, we can demonstrate the outcome of the discovery—a larger decision space—but we cannot explain how the new knowledge arose by analyzing the situation from the agent's perspective. Heifetz et al. (2006) show that this difference can be formalized without assuming a single shared state space. Their *"interactive unawareness"* model employs a structure of multiple spaces that differ in expressive power. A richer space may use a broader vocabulary, while the possibilities available to the agent are represented in a less expressive space. An event may therefore be neither probable nor improbable for the agent, but simply impossible to express using the current representation. The formal representation of an agent first in a less expressive space and then in a broader one illustrates the difference between the two states of representation, but does not yet explain exactly where this new distinction comes from.

The problem of including an unknown category in the state space under consideration does not mean that, in economics, all possibilities must be predetermined. Ellis and Masatlioglu (2022) show that available options can be grouped into different categories depending on how the problem is formulated and that this categorization can subsequently influence which options the agent chooses. The model thus allows for the categorization of already available possibilities to emerge within the decision-making problem itself. However, this is still distinct from the emergence of a completely new distinction that the agent was previously unable to analyze at all. It is precisely this second case that is relevant to the subsequent argument.

In the subsequent analysis, the agent's representation can therefore be understood operationally as a set of variables and relations that the agent can use to formulate hypotheses, interpret data, and select actions. A change in representation occurs when this framework is expanded or reorganized. The economics of information describes changes in beliefs about known possibilities, the use of knowledge distributed among agents, and communication within a defined problem. The literature on unawareness adds an important distinction between a possibility that receives a low probability and one that the agent is currently unable to formulate. These approaches allow us to describe different states of an agent's knowledge and representation. They leave open the process through which a new hypothesis or distinction first becomes available. Before turning to that process, it is useful to see how far economic coordination can go when the relevant space of possibilities is already defined.

2.5. How Far a Closed System Reaches

The closed system reaches further than the preceding examples might suggest. Gode and Sunder (1993) placed simple computerized traders in a double-auction market. The traders submitted random bids and offers subject to a budget constraint, yet the market achieved allocative efficiency close to the competitive benchmark. The result shows how much institutional structure can accomplish even when little reasoning is required from individual participants. The market can coordinate their actions because the relevant goods, constraints, bids, and possible allocations are already defined within the environment. Mirowski and Nik-Khah (2017) place this line of research within a broader development in which part of the cognitive burden is shifted from individual agents to the market as an information-processing institution. The result itself is enough to make an important distinction. A market mechanism can only select and coordinate among outcomes that are already feasible within the environment.

This tells us much about the power of institutional coordination, but it tells us much less about the process through which an agent comes to formulate a hypothesis that was absent from the set of possibilities considered before. This is the boundary that matters for the rest of the article. Institutions can allocate and coordinate within an existing space of possibilities. The mechanisms considered so far leave open the process through which the set of hypotheses available to an agent is generated or transformed. A new hypothesis must first become available before it can be evaluated, communicated, or selected. The next section therefore turns to the production of new ideas and to the search process through which new candidates enter consideration. Once this process is clear, we can ask whether interaction with other people changes its course.

3. The Production of New Knowledge

We now turn to the process through which new hypotheses and solutions are generated. Economic theories of growth and innovation already treat knowledge as the result of research activity, and inventiveness can be understood as a search through potential solutions. For the argument developed here, three sources of new ideas need to be separated. An individual can generate a hypothesis through their own reasoning and exploration. Knowledge produced by other people can expand the material available for further search. Direct interaction can also affect the direction in which that search develops. Keeping these processes separate will allow us to identify the specific contribution of communication.

3.1. New Ideas in a Closed Environment

The starting point for analyzing knowledge creation is the distinction between the exchange of information and the acquisition of data about the world, on the one hand, and hypothesis generation based on that data, on the other. These processes are interrelated and influence one another, but they are also distinct. An additional observation may alter the relative credibility of two explanations already under consideration, or it may prompt the formulation of a new one. However, a new hypothesis may also emerge without a new signal from the environment, based on knowledge and cognitive operations already available to the agent. Communication is therefore not a necessary condition for epistemic novelty. This difference can be framed as a distinction between evidence acquisition and hypothesis generation. Evidence acquisition provides material against which hypotheses can be evaluated. Hypothesis generation, on the other hand, determines which possible explanations, rules, or solutions will even become candidates for evaluation. The mere fact that an agent possesses more data does not, therefore, determine whether or which hypotheses they will be able to formulate. Furthermore, even the same set of information can lead to different conclusions, different hypotheses, and even different outcomes.

Klahr and Dunbar (1988) demonstrate precisely this distinction in their experiments. Participants performed a task designed to simulate the process of discovery: after familiarizing themselves with a specially designed robotic toy, they were to discover the operating rule of a button previously unknown to them. They had to formulate hypotheses, design experiments, and evaluate the results. The authors observed two main ways of arriving at new discoveries. Some participants relied primarily on their own experiences—rules they had stored in memory or prior knowledge. Others formulated them through experimentation and subsequent generalization of the results. In the second study, participants were asked to propose hypotheses even before beginning to experiment. Some were able to find the correct rule immediately, while others needed several attempts and empirical experiments. The same correct hypothesis could thus arise in two ways: from resources already available to the agent or as a result of further exploration of the world. On this basis, Klahr and Dunbar describe scientific discovery as a “dual search”: an interrelated search in the space of hypotheses and in the space of experiments. The first process involves hypothesis generation and selection of candidates to explain a problem; the second involves selection of actions to test them and gain experience. Both processes interact iteratively. Hypotheses guide which experiments will be conducted, and their results alter the evaluation of existing hypotheses and may influence further hypothesis generation.

For the argument that follows, we do not need to distinguish between two separate sources of new knowledge—a basic understanding of search suffices: the agent performs successive operations that influence which hypotheses, solutions, and tests become the next candidates. From this perspective, the production of new knowledge involves generating content that was not previously available and can only now be subjected to critical evaluation. The possibility of arriving at new solutions independently places an important limit on the argument developed here. Since an individual can generate new hypotheses without communicating with others, novelty cannot be identified with social interaction. However, the question remains: can contact with others alter the course of the search—what possibilities does an agent generate, which paths do they develop or abandon, and which tests do they deem worth undertaking? Before such an influence can be isolated, however, it is necessary to determine how economics already conceptualizes the production of ideas and the search leading to new solutions.

3.2. Production of New Ideas in Economics

Innovation and the production of new knowledge have also been extensively studied in economics. In growth and innovation theory, new knowledge can be produced through deliberate investment and, once produced, can become an input into subsequent innovation. Technological progress can also contribute to expanding the set of available solutions. Economics therefore has the tools to treat the development of new ideas in an endogenous manner. The remaining question is where the problem of *search* enters these models.

In many economic models, technological change is treated as an exogenous source of progress whose path is specified by the researcher or calibrated to historical values. There are, however, important exceptions. In Romer (1990), technological change results from deliberate investment decisions made by economic agents. In his model, conducting research requires specific expenditures, and decisions to undertake such research depend on economic incentives. Research activity can result in new ideas and technologies that increase productivity and thus contribute to economic growth. For this article, the most important aspect is the change in the status of innovation itself: new knowledge can be the result of the production process. Economics can therefore describe a situation in which research activity leads to the discovery of a technology that was not previously available.

Although this approach is endogenous, it remains highly general—we do not delve into the specifics of how a new idea is generated. The relationship between research expenditures and knowledge production does not precisely define the steps through which a researcher arrives at a given solution. Furthermore, research and development expenditures alone do not determine the outcomes of innovative activity. Cross-country studies show that R&D productivity also depends on the institutional environment and the way elements of the innovation system are organized and interconnected (Furman et al., 2002). If we treat research activity as a process of seeking new solutions, research expenditures alone do not yet explain how a specific idea arises. We must also consider which hypotheses and experiments researchers explore and why they pursue certain lines of inquiry while abandoning others.

Weitzman (1998) calls this phenomenon “recombination”—a process in which new ideas arise through the appropriate configuration of existing technologies, and the resulting new combinations can then enter the pool of knowledge used in subsequent explorations. Novelty, therefore, does not require an independent addition from outside the existing body of knowledge. It may consist of combining known elements into a whole that previously did not function as a separate idea. At the same time, recombination shows why the existing stock of knowledge alone is insufficient to explain which ideas will actually be discovered. If a great many different configurations can be created from the available elements, the mere fact of their existence does not lead to a discovery. The availability of ideas A and B does not explain why they will be combined specifically with one another, in a particular way, and at a given moment. In this context, Weitzman points out that the limitation may lie in the ability to process a vast number of potential new ideas into a useful form (Weitzman, 1998). Knowledge production then takes the form of selection within a vast space of possible recombinations.

Fleming (2001) develops this intuition at the level of technological search. In this view, inventiveness consists of experimenting with components and ways of combining them. The uncertainty of the outcome is greater when the inventor uses elements with which they have little prior experience, or combines them in a way that has not been tested before. New components and new combinations yield, on average, less useful results, but at the same time increase the variability of outcomes—including both failures and breakthroughs. Recombination therefore requires exploration; there is no rule stating that simply combining existing components automatically yields a valuable idea.

The *search* process thus becomes an economically significant component of knowledge production. If search paths differ in cost, probability of success, and potential value of the outcome, the result also depends on which parts of the space of possibilities are actually explored (Fleming, 2001; Romer, 1990; Weitzman, 1998). The problem shifts from the question of whether ideas can be produced to the question of what guides the search among many potential solutions and why some achieve better results than others. Regardless of the field, prior knowledge is always one of the most important factors in innovation research. It can alter the starting point, allow for additional elements in new combinations, and open up paths that a single agent would not be able to find or replicate on their own at a reasonable cost.

3.3. Knowledge produced by others expands individual search

The search process does not begin in a vacuum. An inventor is usually familiar with previous solutions, experimental results, theories, and other elements produced by past generations and earlier research. Thanks to this, they do not have to independently retrace the entire path leading to the current state of knowledge. Like Newton, every (successful) inventor and researcher stands on the shoulders of giants. Someone else's result can become the starting point for the next inquiry, the basis for a new combination, or the rationale for choosing a specific experiment. The social nature of knowledge production thus emerges through accumulation: the result of one search changes the conditions under which the next one is conducted.

A good illustration of this mechanism is the phenomenon of collective invention. In the case analyzed by Allen (1983), companies disclosed information to competitors about new installation projects and the results achieved, allowing subsequent firms to build on earlier solutions and develop them further in subsequent implementations. Progress arose through a series of modifications in which one entity's results served as the foundation for another's. From a search perspective, this means that some trials had already been conducted, certain configurations had become known, and the next agent could begin the search further along the path than it would have reached on its own without incurring additional costs. When such a flow of knowledge is systematic and repeated, a positive feedback loop can emerge between the public disclosure of earlier results and subsequent innovation. Cowan and Jonard (2003) formalize collective invention as a process in which agents share knowledge, and the resources received from others raise their own level of knowledge and enable them to produce further improvements. The pace of accumulation depends on the structure of the network through which knowledge circulates. The outcome of someone else's experiment, a technology design, or even a partial solution to a problem can enter the set of potential combinations, suggest new tests, or create a potentially viable path that was previously unfeasible.

Moreover, the significance and impact of socially generated knowledge grow as it accumulates. Jones (2009) describes the "burden of knowledge": successive generations of innovators must master an ever-increasing body of prior achievements before they can push the frontiers of knowledge further. One result of this situation is increasingly narrow specialization, which, however, limits the scope of knowledge of any single individual and increases the importance of teamwork when innovation requires combining competencies distributed among specialists. Accumulation thus has a dual effect. On the one hand, it expands the total pool of possibilities available in the innovation process; on the other hand, it limits the breadth of knowledge that can be mastered by a single researcher—increasing the importance of specialization and collaboration across specialists.

Individual search is thus embedded in a socially constructed structure of knowledge. An individual may independently perform the operations leading to a new hypothesis, but relies on concepts, solutions, tools, and competencies developed by others (Allen, 1983; Cowan & Jonard, 2003; Jones, 2009). The production of ideas can thus be a social process, even though a specific hypothesis originates with a single agent. However, a mechanism defined in this way still does not require direct interaction between the agents. The result of another person's work may be recorded in an article, a project, a database, or a technology and enter another person's search long after the author's reasoning process has ended. Knowledge can circulate, accumulate, and increase the productivity of research without the creative process of questioning ideas and iterative scientific work. Similarly, the mere bringing together of various specialists on a single project does not yet determine whether and how their knowledge is combined into a new result. To attribute a distinct mechanism to interpersonal interaction, one must separate access to the outputs of earlier search from interaction during an ongoing search process.

3.4. Access Is Not Interaction

This difference is illustrated by a simple example: a researcher with access to an ideal library may have access to a vast number of articles, experimental results, models, and arguments. Each of these resources can alter the researcher's search: suggest a hypothesis, reveal a method previously unknown to them, or enable the recombination of elements from distant fields. The library, however, remains insensitive to the current course of the researcher's reasoning. A text written earlier will not change the next sentence simply because the reader has identified a specific error, made an incorrect assumption, or formulated a question the author did not anticipate. The product of an earlier search is available, but that product does not respond to the state of the reader's search.

The value of a different perspective can be very significant even without dialogue. Jeppesen and Lakhani (2010) analyzed 166 scientific problems announced via a "broadcast search" format, to which external solvers submitted solutions; the study sample included over 12,000 scientists. The likelihood of a winning solution increased with the distance between the solver's field of technical specialization and the field of the problem being solved, which the authors interpret as a significant influence of a different perspective on innovation. Thus, the very act of submitting a problem to the judgment of a broad audience with diverse perspectives increases the chances of resolving it. This demonstrates the power of a social approach to search, though it does not yet allow us to isolate the value of mutual interaction among participants. This example allows us to sharpen the distinction between access to a different perspective and reciprocal interaction. Another person can influence a researcher's search either because they possess information the researcher lacks or because they represent a perspective the researcher would not have reached independently. Both effects can be conveyed through a finished article, a review, or a contest entry. Mutual interaction, however, introduces an additional possibility: another person's subsequent contribution may depend on what the researcher has just said and on the specific point at which the search has encountered difficulty. It therefore also matters whether others' responses are adapted to the current course of the researcher's reasoning. Such a contribution may relate directly to an issue that has only come to light as a result of the interaction.

The literature on innovation has long recognized interactions as an important element in the creation of new solutions. In many approaches, feedback loops and the relationships among participants in the innovation process play a special role. Lundvall (1988) describes innovation in user–producer relationships as a process based on the exchange of qualitative information, collaboration, and mutual learning. In the chain-linked model of Kline and Rosenberg (1986), innovation proceeds through numerous back-and-forth cycles between successive stages: problems revealed during design, production, or use may redirect attention to earlier solutions, existing knowledge, or further research. Both approaches thus depict a process in which subsequent actions are shaped by the results of earlier steps, and feedback constitutes a permanent element of innovation. Similarly, Nonaka (1994) conceives of knowledge formation as a dynamic and social process, based on interaction and the conversion between tacit and explicit knowledge. New knowledge develops at the individual level, and the organization can reinforce it through processes of interaction. This literature shows that interaction has long been treated as an important part of innovation. These approaches demonstrate that feedback, collaboration, and dialogue matter for knowledge creation, but they do not yet isolate the precise mechanism by which a specific partner's response alters the next move in a single participant's search.

By breaking this mechanism down into its components, we obtain a distinct object that we will analyze in the following section. While access to others increases the pool of knowledge and perspectives an agent can draw upon, mutual interaction adds a further possibility: another person can respond to the current state of the agent's reasoning and redirect the next step. We refer to such responsive interaction as *adaptive feedback*. Adaptive feedback is a response tailored to the current state of another participant's reasoning that can influence that participant's next step. This process distinguishes the interaction from mere access to another person's knowledge, as the interlocutor can react to information revealed only in the course of reasoning. The key question, therefore, is how such a reaction can change the direction of further inquiry and contribute to the formation of a new hypothesis.

4. The Impact of Communication

Incorporating a different perspective into the creative process can broaden the scope of the search even without direct contact with its author. However, in direct conversation, the other person can respond to their partner's line of reasoning and adapt their response to the situation. We refer to this type of interaction as adaptive feedback. In the following section, we analyze how such a reaction can alter the course of the search for a solution and what impact it may have on the discovery of new hypotheses.

4.1. Interaction Introduces Feedback

A useful point of reference is the model of information transmission developed by Shannon (1948). The sender selects a message, the system transmits it through a channel, and the technical problem is to reproduce the message accurately at another point. The semantic meaning of the message lies outside the model; the analysis concerns channel capacity, interference, and transmission reliability. This provides a clear benchmark for the communication problem considered here. Transmission can be evaluated by comparing the message sent with the message received. A conversation introduces a further element because the receiver's interpretation and response can affect what the sender says or considers next.

Real conversations between people have a much richer structure—they cannot be reduced to the technical transmission of information alone. Grice (1957) shifts attention from the words transmitted to the speaker's communicative intention and the listener's interpretation of that intention. The interlocutor must therefore determine what the utterance refers to, what the scope of a given statement is, what conclusions are intended, and why a specific element of the message is relevant in the current context. The same utterance can be understood differently depending on the context, and thus lead to different reactions from the listener. The mere fact that information is exchanged correctly does not, therefore, guarantee a shared understanding of it. The speaker may flawlessly convey the words they intended to say, yet the listener may interpret them differently. Text-based communication illustrates the problem: contextual cues such as irony may be harder to interpret, especially when interlocutors know little about one another. In direct conversation, misunderstanding can itself generate an observable response, allowing the speaker to clarify or reformulate the message.

In Clark (1996), language use is treated as "joint action," in which one participant's utterance and the other's reaction are treated as a single coordinated process. A central element of this framework is "grounding": an ongoing process through which interlocutors establish whether their contributions have been understood well enough for the conversation to proceed. Each response from the recipient simultaneously reveals something about their interpretation. It may confirm that the interlocutors share the same understanding, but it may also indicate that the message was unclear. In such a situation, the communication process can take a step back, return to earlier information, and clarify what was meant. Through clarification and repair, one can adjust the wording and reframe the presentation of the issue.

The structure of the communication process can be represented as a sequence:

$$A_1 \to B_1 \to A_2 \to B_2 \to \cdots$$

Communication thus introduces an additional element of conditionality to subsequent messages. B_1 arises as a response to A_1: it depends on what B heard, how they interpreted it, and how they reacted. A_2 is then formulated under conditions altered by B_1. The content of subsequent messages is therefore endogenous to the interaction and depends on its previous course. This conditionality is significant from the perspective of the process of producing new knowledge. The other person's response may address ambiguities that arose in the preceding utterance and therefore could not have been predicted at the initial stage. Adaptive feedback thus makes the path of a conversation endogenous: subsequent contributions depend on earlier responses rather than merely adding fixed pieces of content to the recipient's information set.

4.2. Feedback Can Redirect Search

Not every communicative interaction matters for the production of new knowledge; the relevant cases are those that can change the course of search. If response B to A_1 merely conveys an existing fact, the interaction remains within the familiar domain of information transmission. A different structure may emerge when a response challenges an assumption or introduces an alternative interpretation that redirects the argument along a path the agent had not previously considered. The other person's contribution can then alter the subsequent steps of the search, regardless of whether it provides new knowledge about the world. This matters because an agent cannot enumerate and evaluate every logically possible hypothesis or answer. The agent must therefore explore only a subset of the available paths—selecting specific analogies and deciding which issues require further clarification and which are likely dead ends. These decisions simultaneously influence the possible choices in subsequent steps, and the interlocutor's responses can intervene precisely in this process of iterative selection. The interlocutor need not provide a ready-made solution. It is enough to change the agent's search trajectory, question a premise, or demand a more detailed justification for a claim.

The difference from standard learning is particularly well captured by the distinction between "signal" and "reason." A signal alters the agent's beliefs regarding possibilities that it is already capable of representing. If an observation o is more likely in state X than in state Y, its occurrence may strengthen the agent's belief that X is indeed the case. These relationships between o, X, and Y remain defined within the closed model of the problem. The interlocutor's response, however, can challenge the structure that links observations to the possibilities represented in the model; for example, the division into X and Y may itself turn out to be inadequate. The difference can be summarized as follows: a signal changes beliefs within a model; a reason can change the model itself. Of course, not every argument leads to a change in the representation of the problem. What is important is that communication can subject inferential relations and assumptions—which remain constant in a pure update model—to criticism. The argumentative theory of reasoning developed by Mercier and Sperber (2011) helps explain why interaction can function in precisely this way. In their view, reasoning is strongly linked to the production and evaluation of arguments in a social context. An argument provides the interlocutor with a rationale intended to support a specific conclusion, thereby subjecting the connection between the premise and the conclusion to evaluation. The recipient may accept an observation while questioning its interpretation, but may also request an explanation of a step that the interlocutor had taken for granted without much reflection. In this process, the very structure of reasoning can become the subject of the interlocutor's feedback. A related mechanism specific to interaction is articulation. When reasoning individually, an agent can move between different thoughts without explicitly revealing all intermediate assumptions. Presenting a position to another person requires intermediate assumptions to be made explicit and the relevance of the proposed solution to be explained. Once articulated, those assumptions become available both for self-scrutiny and for the interlocutor's criticism. A problem may therefore become visible only when an initially private line of reasoning is expressed and subjected to a response.

Nemeth (1986) shows that the direction of exploration can be influenced even when an agent disagrees with another person's position. Her experiments indicate that exposure to a consistent opposing position can foster more divergent processing of the problem: participants consider more of its aspects and more frequently re-examine their accepted premises, whereas confirmation promotes convergence. This occurs because a differing response disrupts the established search path and can trigger alternative modes of analysis, even if the recipient ultimately does not accept it. Further insights into the creative process are provided by direct observations of scientific work. Dunbar (1997) tracked projects conducted in four molecular biology laboratories, observing, among other things, meetings during which researchers presented results, asked questions, and proposed experiments, hypotheses, and interpretations. These discussions served as a forum where new ideas and concepts emerged and were modified, and the reactions of other participants led to changes in the approach to specific problems. These observations do not allow us to conclude that a particular hypothesis would not have arisen without a given interaction. However, to identify the mechanism under analysis, it suffices to observe that another person's reaction influenced the subsequent course of reasoning, and a new solution emerged only in subsequent steps of this process. Communication can therefore affect both the search for new answers and the critical examination of existing hypotheses. This does not, however, automatically mean that through interaction we can discover entirely new categories that were not previously within the researcher's conceptual apparatus.

4.3. Search Can Produce New Representations

Thus, we can distinguish four effects of communication according to how it changes the agent's knowledge and representation of the problem. In the simplest case, we have "new information," which changes the agent's assessment of a previously known possibility. "New inference" arises when a previously unformulated conclusion is derived from previously available elements (similar to Weitzman's recombination). The next level is a "new hypothesis," which emerges when an agent begins to consider C, even though it was originally outside the scope of their analysis. The final level is a "new representation," in which the way the problem is formulated changes and elements become available that were not expressible within the researcher's previous representation. A new category may distinguish cases previously treated as a single type; a new variable may reveal an overlooked dimension of a phenomenon; and a different formulation of the problem may alter the very question the agent is seeking to answer.

A change in representation therefore has a different structure than a simple expansion of the list of hypotheses. It is not merely a transition from

$$\{A, B\} \to \{A, B, C\},$$

but rather the possibility of a transition from

$$R_0 : \; H(R_0) = \{A, B\} \to R_1 : \; H(R_1) = \{A', B', C\},$$

where R_0 and R_1 denote operational representations of the problem, and $H(R)$ is the set of hypotheses or answers available within a given framework. The symbols A' and B' indicate that the emergence of C may alter the way earlier possibilities are understood. A new distinction may require reclassifying the material, altering the relationships among previous alternatives, and changing the type of evidence considered relevant. The representation itself may therefore be an outcome of search. Simon (1995) describes discovery as a search taking place across multiple spaces, encompassing observations and hypotheses, but also possible problems, representations of the problem, and new concepts. Discovery, therefore, does not always consist of finding the missing answer to a previously defined question. Sometimes it is necessary to construct a way of describing a phenomenon that only then makes a new set of questions and answers accessible.

Here we arrive at the essence of the mechanism through which interaction can influence the creative process. Another person's response may reveal to the agent that the distinctions made so far or the categories under consideration are insufficient for solving the problem. A counterexample may undermine a line of reasoning, a question may prompt the agent to distinguish between cases previously treated as a single unit, and an analogy may suggest a relationship that is difficult to discern in the existing description. The interlocutor does not need to have a ready-made new representation—their contribution can alter the constraints on further reasoning so that the existing framework ceases to be sufficient. The next stage of the search will then consist of seeking an alternative structure for the representation of the problem.

Similar cognitive processes have been studied in the literature on model-based reasoning. Nersessian (1999) shows that scientists can create and modify their understanding of a phenomenon under study by building models, using analogies, and mentally testing how a given model would behave under various conditions. These models are constructed based on knowledge of the phenomenon under study and are then evaluated and modified in the course of further reasoning. A change in the way a problem is understood may thus result from the restructuring of the model using information the researcher already possesses. This does not require the emergence of new data or observations. The boundary between a new hypothesis and a new representation corresponds to the distinction made in Section 2.4 between an unlikely possibility and a possibility that the agent is not yet able to formulate (Dekel et al., 1998; Heifetz et al., 2006). In the latter case, C becomes operationally accessible only as a result of a cognitive process.

The way in which the agent frames the problem at a given moment should not be equated with the complete model that a researcher might later construct. Once the process is complete, it is possible to describe all the possibilities considered earlier and later within a single model. However, such a description does not explain how the agent began to consider a new possibility or why their understanding of earlier possibilities changed—and these are precisely the most important issues analyzed in this article. Economics allows for situations in which the way available possibilities are categorized depends on the context of the decision (Ellis & Masatlioglu, 2022). However, the mechanism examined here concerns a different process: responsive interaction alters the course of the search, and in the course of this search, the agent constructs a distinction or representation of which they were previously unaware. The subject of this analysis is therefore the way in which dialogue can contribute to the emergence of such a change in representation.

If we represent the structure of the interaction as:

$$A_1 \rightarrow B_1 \rightarrow A_2.$$

Then B_1 need not contain C as a pre-existing idea. If it changes the direction of A's search in such a way that a new hypothesis appears in A_2, C was not transferred from B to A. It arose in a process whose subsequent steps depended on the participants' earlier reactions. The strongest case occurs when a new element also reorganizes earlier possibilities, and the result of the process is a new representation. However, a mechanism formulated in this way remains abstract. Its significance for knowledge production depends on whether, in actual processes of discovery, one can identify a sequence comprising an initial representation of the problem, a specific intervention by another person, a change in the subsequent search, and an epistemic result that did not previously exist as the ready-made content of a message.

4.4. The Mechanism in Scientific Discovery

Dunbar (1997) suggests that this sequence also occurs in scientific work: laboratory discussions can alter the interpretation of results, lead to new analogies, and change the way a problem is understood. In this section, we will analyze two historical case studies of biological discovery that allow us to understand this process more thoroughly and break it down into its components. In the first case, interaction helped connect two problems that had previously been represented separately, and in the second, thanks to a correction from an interlocutor who had a different perspective, the representation of the elements used to construct the model changed.

The first case concerns the development of the operon model, for which François Jacob and Jacques Monod received the Nobel Prize in Physiology or Medicine in 1965. Initially, they worked on problems belonging to different research programs—Jacob focused on the regulation of viral activity, while Monod focused on the induction of enzymes involved in lactose metabolism. According to Jacob's later account, the two problems initially seemed completely unrelated, even though they were being studied at the same institute and on the same organism (Jacob, 2011). In one case, the researchers investigated under what conditions a virus remains inactive in a bacterial cell and under what conditions it begins to replicate. In the other, researchers analyzed when a bacterium initiates the production of an enzyme needed to utilize lactose. Both research programs employed separate sets of tools: variables, techniques, and explanations. The shift began when a similarity between the two phenomena under study was recognized. Jacob noted that both might be governed by a similar principle of gene expression regulation. This analogy then became the starting point for further collaboration with Monod—it was not yet a ready-made solution, but a hypothesis that needed to be further developed and tested. A common explanation had to fit both research programs simultaneously, so the model was refined gradually through the verification of successive hypotheses. Their search thus shifted from explaining the two phenomena separately to seeking a single principle that could explain both. This new approach to both problems also changed the way earlier experimental results were interpreted. Findings from studies on lactose utilization by bacteria began to be interpreted as indicating that the activity of specific genes could be inhibited by a separate regulatory element. The similarity to studies on viral activity suggested that the same general regulatory mechanism might explain both phenomena. The operon model organized these relationships by distinguishing between elements responsible for performing specific functions and elements controlling their activity (Jacob, 2011). The result was thus a shift in how both problems were understood. Phenomena previously analyzed separately began to be viewed as instances of a similar mechanism of gene activity control. Jacob initially presented the first intuition regarding the similarity between the two phenomena as his own insight, and the significance of his collaboration with Monod became apparent only at a later stage, when this initial analogy became the subject of joint development and testing. Monod brought experience from a different research program, allowing the proposed explanation to be tested against findings related to both phenomena. In the course of this collaboration, the direction of the ongoing search evolved, and the joint model was gradually refined by combining elements previously used in the two separate studies. This interaction allowed them to connect the two lines of inquiry and integrate the two research programs, resulting in a new solution that explained both problems.

The case of Watson and Donohue illustrates another variation of the same mechanism. In early 1953, James Watson was constructing a model of DNA in which, initially, the bases were supposed to pair with bases of the same type. He based this on the standard representation of DNA bases accepted at the time. This representation determined which bonds between adenine, thymine, guanine, and cytosine appeared chemically possible. The accepted depiction of these elements thus limited the set of configurations that Watson considered when building his DNA model. Jerry Donohue, an American chemist who was in Cambridge at the time, shared an office with Watson and Crick. Thanks to his knowledge of molecular structure, he recognized a problem in the way Watson represented the DNA bases in his model. He pointed out that their actual chemical structure differed from the one Watson had assumed based on the textbooks available at the time. This correction led to a change in Watson's approach to base pairing, which allowed for further work on the model. Although Donohue did not provide a ready-made solution in the form of A-T and G-C pairs, he corrected one of the key initial elements of the problem, and on this basis, Watson was able to consider other representations. In this way, Donohue's correction changed the direction of Watson's further work—once the correct chemical representation of the bases was adopted, the earlier idea of pairing bases of the same type no longer fit the model being constructed. Watson therefore began to experiment further with various ways of pairing them and thus ultimately arrived at the adenine-thymine and guanine-cytosine pairs. This pairing fit both the structure of the DNA model he was building and the already known data on the proportions of individual bases. As Small (2023) reconstructs the episode, the fundamental data available to Watson remained unchanged at that time. What changed was the way they were organized within the model, which made a new solution accessible. The case of Watson and Donohue thus illustrates how another person's response can alter the subsequent search for a solution. The way Watson initially presented the DNA bases limited the configurations he considered—Donohue corrected this aspect but did not provide Watson with a ready-made DNA structure. This correction later led Watson to explore other ways of pairing the bases, and it was only in the course of this search that the A–T and G–C pairs emerged. For this discovery, Watson and Crick were awarded the Nobel Prize in Physiology or Medicine in 1962.

Jacob–Monod and Watson–Donohue illustrate two ways in which responsive interaction can influence the discovery process. In the first case, bringing two previously separate problems into contact helped produce a shared explanation that could account for both. In the second case, an interlocutor's different perspective and criticism led to a modification of an element of the earlier representation, opening up new avenues for experimentation. In both cases, the interlocutor's contribution changed the subsequent course of search rather than supplying the final solution. The relevant point is not that no information was transmitted during the interaction, but that the eventual result was not contained in what was transmitted. The contribution altered what the researchers explored next, and the new solution emerged only in the subsequent search process. Of course, we cannot conclude from these examples that the operon model or correct base pairing would never have been discovered without the interactions described. However, to support the thesis regarding the mechanism of discovery, it suffices to demonstrate a real, documented sequence of events in which another person's contribution influenced subsequent elements of the discovery process, and the resulting solution was not initially available to any participant as ready-made knowledge. At the same time, the ability to influence the search process does not mean that interaction will always have a positive effect on the stock of available knowledge. The same mechanism that can guide an agent toward the correct solution can just as easily redirect them toward unproductive paths or prematurely terminate promising lines of inquiry.

4.5. Why Interaction Sometimes Helps and Sometimes Harms

The impact of communication on the creative process is not necessarily positive. Another agent's response can either redirect the development of an idea toward a productive line of inquiry, or interrupt the exploration of a promising path that might otherwise have been developed further. Whether such redirection is beneficial or harmful can therefore be assessed only ex post. It is worth distinguishing between two effects here: the impact of contact with others on a single researcher's inquiry and its impact on the diversity of directions pursued by the entire group. Exposure to others' solutions can improve a researcher's individual reasoning, allowing them to avoid mistakes or capitalize on a promising idea. However, when many researchers react to the same visible solutions, they begin to direct their searches along similar paths. Individuals may thus learn from others, while the group as a whole explores fewer and fewer distinct possibilities. Ideas that did not attract the interest of others at an early stage may be abandoned before there is sufficient basis to assess their value.

Bernstein et al. (2018) demonstrate this trade-off in an experiment in which participants in groups of three worked on the problem of finding the shortest possible route between a set of cities. The researchers compared three types of groups: those working completely independently, groups with constant access to the solutions of other participants, and groups in which such access occurred only occasionally. Independent work led to greater diversity in solutions, while constant interaction improved the average result but more quickly steered participants toward similar solutions. The most beneficial combination of both effects was achieved through intermittent interaction: participants could periodically draw on others' ideas and, in between those moments, independently develop their own solutions. This finding shows that the significance of interaction also depends on when it occurs in the search process—a period of independent work allows different solutions to develop separately, before they are subjected to comparison and criticism.

Interaction can also hinder the very generation of new ideas. Research on brainstorming shows that a group working together may generate fewer ideas than the same number of people working independently. Diehl and Stroebe (1987) attribute a significant portion of this difference to "production blocking": when one person is speaking, the others must wait for their turn, which can interrupt the development of their own idea or cause it to be forgotten before it can be voiced. The source of the problem, therefore, lies in the very structure of group discussion, which can limit the individual flow of idea generation.

Communication can thus limit the generation of new ideas through at least two different mechanisms. First, exposure to others' solutions may cause many people to begin developing similar ideas, resulting in fewer alternative possibilities being explored in parallel. Second, continuous group discussion may hinder the development of individual thought processes, as the participant's attention is partially occupied with following and sustaining the interaction. Greater access to others' ideas may therefore improve individual reasoning while reducing the diversity of independently developed avenues from which alternative solutions could emerge.

The productivity of interaction therefore also depends on whether participants retain sufficiently different ways of thinking. Encountering a different perspective can reveal the limitations of one's own understanding of a problem and steer further exploration in a different direction. Nemeth (1986) shows that the presence of a genuinely dissenting viewpoint can prompt participants to consider a greater number of possibilities. However, this effect requires that there be a genuine difference among participants in how they approach the problem. If prior communication has caused everyone to start thinking alike, the subsequent discussion may only further develop a shared direction rather than introduce new perspectives. Both factors—independent search and interaction—therefore play important but distinct roles in a single process of knowledge production. The period of independent work allows participants to develop different lines of thought and arrive at various solutions, while subsequent interaction allows these solutions to be compared. Another person's reaction may then reveal a weakness in an idea and redirect further reasoning; for this to occur, however, the idea must be developed enough to be articulated and examined. A productive process should therefore allow room both for the independent development of different ideas and for their subsequent interaction.

These findings change the way we should evaluate communication in knowledge-producing organizations. The sheer number of messages, the frequency of meetings, or the intensity of contact does not in itself indicate whether communication fosters the emergence of new ideas. More frequent exchanges may accelerate the dissemination of existing solutions, but at the same time leave less time for the independent development of other ideas. Complete isolation provides more space for independent work but limits the opportunity to contrast one's own way of thinking with a different perspective that could alter the future direction of the inquiry. It therefore becomes important to combine periods of independent work with moments of interaction, during which previously developed ideas can be compared and subjected to criticism. The creative role of communication leads us directly to an institutional problem. Interaction can foster the emergence of new knowledge, as another person's response can alter the course of the inquiry and help develop a new solution. However, a communication process that is too intense or poorly organized can limit the independent development of various ideas and steer participants in a similar direction too early. The economic problem, therefore, lies in organizing work and communication in such a way that different ideas can first develop independently and then influence one another through interaction.

5. Economic Implications

Since interaction can change the direction of inquiry, the way contact between participants is organized also influences the process of knowledge creation. Institutions should shape conditions in which different approaches can first develop independently and later come into contact with one another. An economic analysis of knowledge production should therefore take into account both the way knowledge is transmitted and the way interactions among the people who create it are organized.

5.1. Institutions as Interaction Architectures

A natural point of reference for organizing these concepts is the "information structure," which describes who possesses specific information and how it can be observed or disclosed to other participants—this is a standard component of economic theory discussed in Section 2. The mechanism analyzed here requires supplementing this approach with "interaction structure," which specifies who can engage in direct interaction with whom, when such contact occurs, and how participants can respond to one another's reasoning. The way interactions are organized also influences whether different lines of inquiry can develop independently and when they begin to influence one another.

This perspective on institutional design gives the participants' own reasoning a central role. Mirowski and Nik-Khah (2017) discuss approaches in which institutional structure allows a market to produce outcomes with limited dependence on what individual participants understand or how they reason. Interaction structure concerns a different part of the process. Participants develop the hypotheses, arguments, and representations themselves, while the organization shapes the conditions under which their separate lines of reasoning develop and meet. It can influence how long different approaches remain independent, when they come into contact, and whether one participant has an opportunity to respond to the reasoning of another.

Organizational economics provides a natural starting point for thinking about interaction architecture, as it analyzes how knowledge is distributed among members of an organization and how access to it is organized. In Garicano (2000), different individuals specialize in solving different types of problems, and the hierarchy allows more difficult cases to be directed to those with the appropriate expertise. Employees master the knowledge needed to solve frequently occurring problems, while rarer or more difficult problems are referred to specialized experts. The organization thus determines who should possess a specific type of knowledge and to whom problems requiring a more specialized solution are directed. The structure of communication and hierarchy therefore influences when and how employees can draw on the knowledge of others.

However, in the process of creating new knowledge, a problem may arise even earlier: a researcher may not know exactly what the difficulty is or how to properly formulate it. In Garicano (2000), the problem is already known at the outset, and it is clear what kind of knowledge is needed to solve it. In the actual process of discovery, where the problem may not be clearly defined, what matters is not only access to an expert but also the opportunity for a conversation in which the expert or colleague responds to the current framing of the problem and helps reshape how the problem is subsequently understood. Structure can thus influence not only who the problem is directed to but also whether and how the problem is developed.

A similar distinction can be made when analyzing a network of contacts within a company or at a university—members of the same group often have access to similar information and share a similar way of thinking shaped by shared experiences. In such networks, key figures are those who bridge different groups or research areas that, under normal circumstances, would not intersect. Burt (2004) describes gaps in the network between otherwise disconnected groups as "structural holes." In the organization he studied, individuals in positions that bridged such groups were more likely to propose new ideas, and their ideas were more likely to receive high ratings. The structure of a network therefore also influences whether different ways of thinking can come into contact with one another; those designing such a network should strive to build bridges between different fields.

Connecting different parts of an organization is beneficial only when participants are able to understand the differences in how problems are defined across various areas. Crémer et al. (2007) demonstrate that organizational language is subject to an economic trade-off in this regard. A specialized unit can use its own precise language tailored to the problems it encounters most frequently. Collaboration between different units, however, requires finding a common language that facilitates communication, even though it may be less precise or more complex for each of them. One solution is a specialized translator who facilitates coordination between units that retain their own languages; this role can just as easily be filled by an expert who serves as a liaison between areas. In knowledge production, a similar problem arises when different groups develop their own ways of describing and understanding problems. Having their own language can help each group develop its approach with greater precision, but excessive linguistic differences make it difficult to compare these approaches later on. Therefore, a common language is needed—one that is sufficient for mutual understanding while still allowing room for differences in how the problem is framed.

Even with a strong network of contacts and a common language, interaction may remain of little use if participants are afraid to openly question others' ideas or voice their own doubts, or if they are closed off to differing points of view and prefer to remain within familiar conceptual frameworks. Asking questions, admitting mistakes, or challenging someone else's interpretation can come at an interpersonal cost. The need to involve others in one's problems may also be seen as an unnecessary burden. Edmondson (1999) defines team psychological safety as the shared belief among team members that taking such risks is safe. In a study of 51 teams, higher psychological safety was associated with more frequent engagement in behaviors that promote learning, such as asking questions, seeking feedback, and discussing mistakes. Psychological safety thus creates conditions in which participants can bring problems to light and address them collectively.

Interaction architecture thus encompasses several organizational elements, each of which influences a different aspect of collaboration in knowledge creation. The network structure determines which differing perspectives have the opportunity to come together. Hierarchy influences where problems are directed and who can contribute knowledge to resolving them. Within-team norms, in turn, influence whether participants openly express doubts, subject them to collective discussion, and are open to collaboration. Timing can also be significant—interaction that is too early or too frequent may cause participants to begin developing similar ideas, while a complete lack of contact prevents the later comparison and critique of differing solutions. An organization that generates knowledge must therefore create conditions both for the independent development of various ideas and for their subsequent encounter and mutual influence.

What matters, therefore, is not simply increasing cooperation, but how it is organized. Two organizations may have similar knowledge resources and similar formal channels for information exchange, yet still create different conditions for the emergence of new ideas. In one, different individuals or teams may independently develop their own ideas for a period of time, and then compare them with alternative approaches and subject them to criticism. In the other, frequent contact may standardize lines of thought too quickly, and barriers between groups may hinder mutual understanding and the open questioning of ideas. An assessment of a knowledge-producing organization should therefore encompass both the efficiency of sharing existing information and the conditions under which new solutions emerge. This requires analyzing how the organization shapes the process of searching for solutions that participants do not yet possess. Institutions thus differ in both their information structure and their interaction structure. The former determines who possesses specific knowledge and under what conditions it may be disclosed to others. The latter determines who interacts with whom, whether differences in ways of thinking can be recognized, and whether one person's reaction can influence another's subsequent reasoning. An economic analysis of knowledge-producing organizations should therefore take both of these dimensions into account.

5.2. Conclusion: From Information to Generative Communication

The economics of information provides a well-developed framework for studying situations in which relevant information already exists, even though it may be unevenly distributed among participants, difficult to obtain, or strategically concealed. This information can change participants' beliefs, decisions, and patterns of coordination, and the way it is distributed can have a fundamental impact on economic dynamics. We see this mechanism particularly clearly in the simplest learning models: an additional signal can change beliefs and improve a decision when an agent already knows which options it is considering (Blackwell, 1953; Savage, 1954). More social models also take into account the manner of exchange and the decision to disclose information, but they still assume that the fundamental elements of the problem are already defined before communication begins.

Accounting for the production of new knowledge requires broadening the scope of the analysis. The economics of innovation shows that new ideas can arise as a result of deliberate search, including by combining pre-existing elements into new configurations (Fleming, 2001; Romer, 1990; Weitzman, 1998). Others' knowledge can enhance the productivity of such search even without direct interaction between actors—access to others' findings and solutions can spark new ideas and open up avenues for further search. Such access, however, differs from direct interaction, in which another person can respond to the current course of an agent's reasoning and influence its further development.

During communication, subsequent steps in one agent's reasoning can depend on the other's earlier responses. Each response can refer to what the interlocutor said earlier, while the next step can in turn incorporate the feedback already received. Adaptive feedback defined in this way can reveal a problem in the original reasoning or steer the line of reasoning in a direction that had not been previously considered. Interaction therefore need not provide new information; it is sufficient for it to influence the next step in reasoning and thereby alter the further course of the search process. Observations of scientific work show that questions and discussions can change the way researchers approach problem-solving (Dunbar, 1997). If the interlocutor's response changes the direction of the search, and in the further course of this process a hypothesis emerges that none of the participants had previously possessed, this result cannot be described as simple information transmission. In the stronger case, interaction can alter the very way the problem is represented and can cause a participant to recognize possibilities that they had not previously considered at all. After this process is complete, when analyzing it retrospectively, one can of course describe all possibilities within a broader model, assigning some non-zero probability to a previously unknown solution; however, such an approach fails to explain where the new idea actually came from.

Therefore, an analysis focused exclusively on the final state overlooks part of the economically significant process. The same final belief may arise from evidence gathered independently, from accepting someone else's opinion, or as a result of critical discussion—the result may be the same, but the processes leading to it can vary significantly. In the course of interaction, a participant may change their understanding of a problem and begin to consider a possibility that had not previously occurred to them. However, the impact of interaction on the generation of new knowledge depends not only on the intensity of communication but also on how it is organized. Contact with others can change the direction of the search, but contact that is too early or continuous can also limit the independent development of different ideas. Excessive exposure to others' solutions can lead to the convergence of search directions, and synchronous group work may be subject to production blocking. Research on group work and intermittent interaction shows that the benefits of learning from others are greater when participants also have time to independently develop their own solutions (Bernstein et al., 2018; Diehl & Stroebe, 1987). The organization of the process should therefore first allow for the development of different lines of thought and then create conditions in which they can interact with one another. This makes the institutional organization of interaction central to the analysis. The way work is organized influences whether different ideas can first develop independently and then interact with one another. Institutions thus influence knowledge production not only through the way they distribute existing knowledge among participants but also by designing conditions in which different ways of reasoning can be constructively confronted with one another. Interaction structure thus complements information structure by adding an important dimension of knowledge organization that is particularly significant during the stage of new idea generation.

We therefore distinguish three perspectives on how the stock of knowledge in a community can change. The economics of information explains how existing information is acquired, transmitted, and utilized. The economics of innovation shows that new ideas can emerge through the processes of search and recombination. The main contribution of this article is to demonstrate that the analysis of interactions introduces a mechanism through which another person's response can alter the course of this search and contribute to the emergence of a new idea or an entirely new representation of the problem. An economic analysis of knowledge should therefore encompass both the flow of existing information and the conditions under which new ideas arise.

Communication affects knowledge production in two ways: it transmits existing knowledge between agents, and through interaction it can also contribute to generating new ideas or ways of understanding a problem that were not available to the participants beforehand.

References

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, *84*(3), 488–500. https://doi.org/10.2307/1879431

Allen, R. C. (1983). Collective invention. *Journal of Economic Behavior & Organization*, *4*(1), 1–24. https://doi.org/10.1016/0167-2681(83)90023-9

Aumann, R. J. (1976). Agreeing to Disagree. *The Annals of Statistics*, *4*(6), 1236–1239. https://doi.org/10.1214/aos/1176343654

Bernstein, E., Shore, J., & Lazer, D. (2018). How intermittent breaks in interaction improve collective intelligence. *Proceedings of the National Academy of Sciences*, *115*(35), 8734–8739. https://doi.org/10.1073/pnas.1802407115

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy*, *100*(5), 992–1026. https://doi.org/10.1086/261849

Blackwell, D. (1953). Equivalent Comparisons of Experiments. *The Annals of Mathematical Statistics*, *24*(2), 265–272. https://doi.org/10.1214/aoms/1177729032

Burt, R. S. (2004). Structural Holes and Good Ideas. *American Journal of Sociology*, *110*(2), 349–399. https://doi.org/10.1086/421787

Caldwell, B. (2017, December). Review of The Knowledge We Have Lost in Information: The History of Information in Modern Economics, by philip mirowski and edward nik-khah.

Clark, H. H. (1996). *Using language*. Cambridge University Press.

Cowan, R., & Jonard, N. (2003). The dynamics of collective invention. *Journal of Economic Behavior & Organization*, *52*(4), 513–532. https://doi.org/10.1016/S0167-2681(03)00091-X

Crawford, V. P., & Sobel, J. (1982). Strategic Information Transmission. *Econometrica*, *50*(6), 1431–1451. https://doi.org/10.2307/1913390

Crémer, J., Garicano, L., & Prat, A. (2007). Language and the Theory of the Firm. *The Quarterly Journal of Economics*, *122*(1), 373–407. https://doi.org/10.1162/qjec.122.1.373

Dekel, E., Lipman, B. L., & Rustichini, A. (1998). Standard State-Space Models Preclude Unawareness. *Econometrica*, *66*(1), 159–173. https://doi.org/10.2307/2998545

Dewatripont, M., & Tirole, J. (2005). Modes of Communication. *Journal of Political Economy*, *113*(6), 1217–1238. https://doi.org/10.1086/497999

Diehl, M., & Stroebe, W. (1987). Productivity loss in brainstorming groups: Toward the solution of a riddle. *Journal of Personality and Social Psychology*, *53*(3), 497–509. https://doi.org/10.1037/0022-3514.53.3.497

Dunbar, K. (1997). How scientists think: On-line creativity and conceptual change in science. In T. B. Ward, S. M. Smith, & J. Vaid (Eds.), *Creative thought: An investigation of conceptual structures and processes.* (pp. 461–493). American Psychological Association. https://doi.org/10.1037/10227-017

Edmondson, A. (1999). Psychological Safety and Learning Behavior in Work Teams. *Administrative Science Quarterly*, *44*(2), 350–383. https://doi.org/10.2307/2666999

Ellis, A., & Masatlioglu, Y. (2022). Choice with Endogenous Categorization. *The Review of Economic Studies*, *89*(1), 240–278. https://doi.org/10.1093/restud/rdab011

Fleming, L. (2001). Recombinant Uncertainty in Technological Search. *Management Science*, *47*(1), 117–132. https://doi.org/10.1287/mnsc.47.1.117.10671

Furman, J. L., Porter, M. E., & Stern, S. (2002). The determinants of national innovative capacity. *Research Policy*, *31*(6), 899–933. https://doi.org/10.1016/S0048-7333(01)00152-4

Garicano, L. (2000). Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy*, *108*(5), 874–904. https://doi.org/10.1086/317671

Gode, D. K., & Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, *101*(1), 119–137. https://doi.org/10.1086/261868

Grice, H. P. (1957). Meaning. *The Philosophical Review*, *66*(3), 377–388. https://doi.org/10.2307/2182440

Hayek, F. A. (1945). The Use of Knowledge in Society. *The American Economic Review*, *35*(4), 519–530.

Heifetz, A., Meier, M., & Schipper, B. C. (2006). Interactive unawareness. *Journal of Economic Theory*, *130*(1), 78–94. https://doi.org/10.1016/j.jet.2005.02.007

Jacob, F. (2011). The Birth of the Operon. *Science*, *332*(6031), 767–767. https://doi.org/10.1126/science.1207943

Jeppesen, L. B., & Lakhani, K. R. (2010). Marginality and Problem-Solving Effectiveness in Broadcast Search. *Organization Science*, *21*(5), 1016–1033. https://doi.org/10.1287/orsc.1090.0491

Jones, B. F. (2009). The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder? *Review of Economic Studies*, *76*(1), 283–317. https://doi.org/10.1111/j.1467-937X.2008.00531.x

Kamenica, E., & Gentzkow, M. (2011). Bayesian Persuasion. *American Economic Review*, *101*(6), 2590–2615. https://doi.org/10.1257/aer.101.6.2590

Klahr, D., & Dunbar, K. (1988). Dual Space Search During Scientific Reasoning. *Cognitive Science*, *12*(1), 1–48. https://doi.org/10.1207/s15516709cog1201_1

Kline, S. J., & Rosenberg, N. (1986). An overview of innovation. In R. Landau & N. Rosenberg (Eds.), *The positive-sum strategy: Harnessing technology for economic growth* (pp. 275–306). National Academy Press.

Lundvall, B.-Å. (1988). Innovation as an interactive process: From user-producer interaction to the national system of innovation. In G. Dosi, C. Freeman, R. R. Nelson, G. Silverberg, & L. Soete (Eds.), *Technical change and economic theory* (pp. 349–369). Pinter.

Mercier, H., & Sperber, D. (2011). Why do humans reason? Arguments for an argumentative theory. *Behavioral and Brain Sciences*, *34*(2), 57–74. https://doi.org/10.1017/S0140525X10000968

Mirowski, P., & Nik-Khah, E. (2017, June). *The Knowledge We Have Lost in Information* (Vol. 1). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780190270056.001.0001

Nemeth, C. J. (1986). Differential contributions of majority and minority influence. *Psychological Review*, *93*(1), 23–32. https://doi.org/10.1037/0033-295X.93.1.23

Nersessian, N. J. (1999). Model-based reasoning in conceptual change. In L. Magnani, N. J. Nersessian, & P. Thagard (Eds.), *Model-based reasoning in scientific discovery* (pp. 5–22). Kluwer Academic/Plenum Publishers. https://doi.org/10.1007/978-1-4615-4813-3_1

Nonaka, I. (1994). A Dynamic Theory of Organizational Knowledge Creation. *Organization Science*, *5*(1), 14–37. https://doi.org/10.1287/orsc.5.1.14

Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy*, *98*(5, Part 2), S71–S102. https://doi.org/10.1086/261725

Rubinstein, A. (1989). The electronic mail game: Strategic behavior under "almost common knowledge". *The American Economic Review*, *79*(3), 385–391.

Savage, L. J. (1954). *The foundations of statistics*. John Wiley & Sons.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*, 379–423, 623–656.

Simon, H. (1995). Machine discovery. *Foundations of Science*, *1*(2), 171–200. https://doi.org/10.1007/BF00124609

Small, H. (2023). Bayesian history of science: The case of Watson and Crick and the structure of DNA. *Quantitative Science Studies*, *4*(1), 209–228. https://doi.org/10.1162/qss_a_00233

Spence, M. (1973). Job Market Signaling. *The Quarterly Journal of Economics*, *87*(3), 355–374. https://doi.org/10.2307/1882010

Weitzman, M. L. (1998). Recombinant Growth. *The Quarterly Journal of Economics*, *113*(2), 331–360. https://doi.org/10.1162/003355398555595

University of Warsaw
Faculty of Economic Sciences
44/50 Długa St.
00-241 Warsaw
www.wne.uw.edu.pl
ISSN 2957-0506