Perspective—Organizing for Verification

Daniel Obermeier, Michael Kummer, Tobias Kretschmer

Working Paper · November 2026

ABSTRACT

Organizations process information to cope with uncertainty, and organization theory has long held that processing more information, and more reliable information, leads to better decisions. Both tenets presuppose that the reliability of information is known. We argue that this presupposition rested on three conditions, which are simultaneously in question under human–genAI co-production: (1) producing a convincing form required expertise or care, (2) truth served as a common point of convergence, and (3) independent sources erred independently. Under these conditions, the form of information carried a free reliability signal, and verification beyond a quick read rarely paid. Digitization weakens these conditions, and generative AI removes all three. We develop a framework of two blind spots, at the sender who cannot tell whether an AI-generated output is correct and at the receiver who can no longer infer reliability from form, and trace how they turn into costs for receivers and harms for organizations once unverified information is aggregated. Applying the framework to hierarchies, communities of practice, markets, and open source communities, we find that no form acquired its verification capacity by design and that each is prone to failure resulting from ceremonial verification, the erosion of apprenticeship, or both. Verification, we conclude, replaces provision as the scarce resource in organizational information processing and becomes a problem of organizing in its own right.

INTRODUCTION

Organizations process information to cope with complexity and uncertainty (Simon 1955, 1956, Cyert and March 1963, Galbraith 1977, Tushman and Nadler 1978) by acquiring, interpreting, and synthesizing information (Burton and Obel 1984, p. 614; see also Galbraith 1974, Puranam et al. 2012, Van Knippenberg et al. 2015). The role of organization design then is to design a structure that increases the information-processing capability to match the level of uncertainty the organization faces (Burton and Obel 2004, Galbraith 1977, Gulati et al. 2005, Tushman and Nadler 1978). Consequently, much of the information-processing literature has focused on questions of designing effective decision and aggregation rules and their role in organizational choices (for a review, see Joseph and Gaba 2020).

Two central tenets of the literature are the beliefs that "more is better," i.e., that processing information up to the level required by environmental complexity will lead to better decisions (Galbraith 1974, Tushman and Nadler 1978), and "better is better," i.e., that more accurate or reliable information should be assigned more weight in the aggregation process and the extent to which it actually enters organizational choices (DeGroot 1970). Hence, each piece of information carries two components: First, an estimate, that is, what the information asserts, be it a value, a description of a state of affairs, or a recommended course of action; The second component is a reliability signal, that is, how much confidence that estimate deserves.

In the past, assessing the reliability of a new piece of information was easy, because the reliability signal was a cheap byproduct of sharing it. To share information, the sender needed to understand and synthesize the underlying material into a format *others* could understand. For example, consider a regional sales manager's monthly pipeline report, whose totals reconcile with account-level entries and whose notes refer to specific client conversations. This report is easy to produce for a manager who reviewed each account and hard to fake for one who did not.

The format thereby became a signal in the sense of Spence (1973): costly to produce but cheap to interpret, so that recipients could quickly infer the reliability of new information and weigh it accordingly. Because this reliability signal typically sufficed for successful information aggregation, full verification would have added little value to the reliability signal provided by the format and was therefore often not pursued explicitly—a rational and efficient choice.

In this perspective paper, we argue that, in light of technological advances that make generating, replicating, and transmitting information much cheaper, the close link between information format and reliability will be further loosened, thereby turning information verification into the new binding constraint on organizations' information-processing capabilities. Recognizing verification, which we understand as the process of establishing the epistemic origin that warrants a piece of information circulating in an organization, as a new binding constraint for organizational decision-making and a problem of organizing is important for two reasons. First, it is important because it helps us understand that measures that improved information processing and decision-making in the past may jeopardize them today by exacerbating the verification problem. For instance, implementing vertical information systems (Galbraith 1973, Tushman and Nadler 1978) and creating lateral relations (Galbraith 1973, Tushman and Nadler 1978, Hansen 1999) both increase the efficiency of information exchange by eliminating the need to route the information through an additional managerial layer. These measures were effective because human synthesis was required to translate information into the format needed for sharing. Therefore, the format signaled the reliability and enabled receiving parties to aggregate more information, as correctly weighing information based on the format was easy. Today, because digitization has drastically reduced the cost of replicating and transmitting information (Goldfarb and Tucker 2019), and generative AI (genAI) enables the

generation of completely new—and potentially hallucinated—information in a plausible and well-formatted form without any human oversight, the format of a piece of information does not signal its reliability anymore. Absent verification, the reliability of a piece of information is not merely low but unknown. Vertical information systems and lateral relations become liabilities, because they bypass the actors best positioned to verify and restore the correct weights for aggregation, and hence adding more information might deteriorate decision-making instead of helping it.

As a second benefit the newly proposed perspective enables us to understand that verification is not just a problem of misaligned incentives. Incentives can only make senders reveal what they know or engage more in verification. However, if a sender uses genAI to produce information she cannot verify, she cannot supply a reliable signal of reliability even if she is fully aligned. On the receiver side, a flawless AI-generated format makes it difficult for the receiver to gauge the reliability of the information. The way in which sender and receiver are embedded in the organization, how information is routed, and what decision structures are in place mitigate the limited capacity of individual members to aggregate pieces of information into decisions; however, it matters how unverified information propagates through organizations and affects their decision-making. Under which conditions processing more information becomes a liability is therefore a question of how information processing is organized.

Shifting our attention to information verification as a problem of organizing has important implications not only for organizations and managers but also for scholarly research on organizational design, and it requires a new perspective. To develop this perspective, we start by specifying three conditions under which information about reliability is a cheap byproduct of information processing, making it rational to forgo verification. The first condition (*costly form*)

is that synthesizing and sharing information produces a signal that is costly to produce but cheap to interpret. The second condition (*truth as a shared convergence point*) is that faking is costlier than producing accurate information, so that truth is the natural convergence point of individual efforts. The third condition (*independent errors*) is that individual pieces of information are independent. These conditions explain why it was reasonable to neglect verification as a standalone problem in the past. Because digitization weakens some conditions and genAI all three, they also explain why verification should receive scholarly attention now.

We then develop a framework that conceptualizes information verification as a sender-receiver problem with two blind spots. The first blind spot lies with the sender. A sender who uses genAI cannot tell whether its output rests on valid underlying information or on hallucination. The second blind spot lies with the receiver. Because format no longer signals quality, a receiver can no longer cheaply infer a report's reliability from its format alone — a quick read no longer provides sufficient information about the quality of a report, independent of whether she knows that genAI was used or not. Our framework explains how unverified information may enter an organization through individuals and be aggregated into a final organizational decision. Our paper, therefore, aligns with the ecological approach suggested by Joseph and Gaba (2020) and bridges the constraint and aggregation views of information processing. To demonstrate the usefulness of this framework, we use it to analyze four stylized forms of organizing (i.e., hierarchies, communities of practice, markets, and open source communities). This analysis reveals that these four forms can be categorized into two general approaches to solving the information provision problem that offer verification capacity as a byproduct, shared-cognition-based forms and interface-based forms. Although both approaches exhibit some degree of verification capacity, they suffer from distinct pathologies that follow

from our framework. Shared-cognition-based forms infer reliability from members' tacit familiarity with how their colleagues produce information. This familiarity is built through apprenticeship, which erodes when junior members delegate synthesis to genAI (*pathology 1: erosion of apprenticeship*). Interface-based forms check whether information conforms to a standardized interface rather than where it originates. Once the format and process are decoupled, such checks become *ceremonial verification* (*pathology 2*) (Meyer and Rowan 1977). This demonstrates that, while both forms show some verification capacity, neither was designed with information verification as a design objective.

Accordingly, we propose treating information verification as a new problem in organizational design and encourage scholars to rethink organizational information processing from a verification perspective that asks who should verify, how, and by what means. This perspective also challenges scholars to shed new light on organizational sensemaking, which Weick famously described as "driven by plausibility rather than accuracy" (Weick 1995, p. 55). Under the cheap reliability conditions, plausibility was a valid proxy for accuracy, and relying on it was a rational adaptation. Once format and process are decoupled, the same reliance becomes a liability. Further, this new perspective puts information at the center and pays more attention to its characteristics, such as its form, epistemic origin, and verifiability. This will help scholars refine existing and develop new theories that follow the call by Joseph and Gaba (2020) to shift their focus from a uniform perspective on information to one that acknowledges information as a potential source of conflicts that may compromise organizational decision-making. Additionally, this reconceptualization opens interesting avenues for empirical research.

FROM INFORMATION PROCESSING TO INFORMATION VERIFICATION

Two Tenets in the Information-Processing Literature

The view of organizations as information-processing systems dates back to the Carnegie school of organizational theory, which conceptualizes organizations as systems that acquire, communicate, and transform information into decisions (Simon 1947, March and Simon 1958, Cyert and March 1963). Galbraith (1974, 1977) and Tushman and Nadler (1978) built on this work to develop a contingency perspective, according to which organizations need to adjust their structure to increase their information-processing capacity up to the level required by internal complexity and environmental uncertainty (see also Gulati et al. 2005). Two tenets run through this literature and the research that builds on it. They are rarely stated explicitly, but they become most apparent in models of information aggregation.

The first tenet holds that aggregating more information is better ("more is better"). It underlies the contingency perspective, according to which greater uncertainty requires the organization to process more information (Galbraith 1974, 1977, Tushman and Nadler 1978). It also underlies the decision-rule models, which generalize the logic of the Condorcet jury theorem to organizations (Sah and Stiglitz 1986, 1988). If evaluators judge independently and each is more likely to be right than wrong, every additional judgment brings the collective decision closer to the correct one, and structures can be designed so that reliability increases as evaluators are added (Christensen and Knudsen 2010). The second tenet holds that aggregating better information is better ("better is better"). More reliable information should receive more weight in aggregation (DeGroot 1970), more capable evaluators improve decision quality, especially in centralized structures (Sah and Stiglitz 1991), and evaluator precision is the central parameter in models of organizational search and evaluation (Knudsen and Levinthal 2007).

Both tenets presuppose that the reliability of information is known. Each piece of information carries two components, an estimate, that is, what the information asserts, and a reliability, that is, how much confidence that estimate deserves. The literature has studied extensively how estimates are aggregated, but it has largely taken their reliability as given. Accordingly, low reliability is seen as a design problem that structure can solve by adding evaluators, shifting toward hierarchy, or weighting inputs accordingly (Knudsen and Levinthal 2007, Christensen and Knudsen 2010). Unknown reliability poses a different problem, because more information of the same kind cannot resolve it. Only verification can. Yet the literature has not treated verification as a question of organization design in its own right, and its definition of information processing reflects this. Although no canonical definition exists, most scholars define information processing as the gathering, interpretation, and synthesis of information (Burton and Obel 1984, p. 614, Tushman and Nadler 1978; see also Galbraith 1974, Puranam et al. 2012, Van Knippenberg et al. 2015). The absence of verification from this definition is not a negligent omission. As we show below, verification rarely posed a significant problem in the past, because reliability was revealed as a byproduct of producing and sharing information. The omission, however, leaves the literature without a perspective on how to organize for verification once digitization and genAI turn reliability into an unknown property of information.

We define information verification as the process of establishing the epistemic origin that warrants a piece of information circulating in an organization. This process entails building and maintaining a mapping between that information and the grounds on which it is held to be true. It is this mapping that tells the organization why an estimate deserves the weight it receives. In the remainder of this section, we show where the presupposition of known reliability enters the literature and how our perspective relates to existing explanations of why aggregating more

information can fail. In the next section, we explain why the presupposition was reasonable in the past and why it no longer is.

Three Assumptions Support "More Is Better" and "Better Is Better"

Joseph and Gaba (2020) identify four streams within the literature on organizational structure and decision-making, which treat structure as a means of matching problems to skills (e.g., Garicano 2000), screening proposals (e.g., Sah and Stiglitz 1986), integrating interdependent choices (e.g., Rivkin and Siggelkow 2003), and shaping what information is noticed and how it is interpreted (e.g., Gavetti 2005). Because the verification problem concerns most directly how unverified information enters an organization, propagates within it, and affects its decisions, the screening stream, which explains how individual errors in judging information combine into organizational errors, is the most natural starting point. Within this stream, the presupposition of known reliability rests on three assumptions.

The first two assumptions stem from aggregation studies, which model structure as a decision rule and examine how it shapes errors of omission and commission (e.g., Sah and Stiglitz 1986, Christensen and Knudsen 2010). They understand information processing as evaluation followed by aggregation. During evaluation, each evaluator observes a project of unknown quality through a noisy signal and decides whether to accept or reject it. During aggregation, the structure combines these judgments. The first assumption is that signal precision is exogenous, meaning that evaluators draw an informative signal of given and known quality at no cost. Sah and Stiglitz (1986) assign each evaluator a fixed screening function, Knudsen and Levinthal (2007) treat evaluator ability as a fixed noise parameter, and Csaszar (2013) models individual discrimination with signal detection theory, in which ability is likewise a given parameter. The second assumption is that evaluators' errors are independent (Sah and Stiglitz

1986, 1988). Under these assumptions, the stream shows how structures between hierarchy and polyarchy trade off errors of omission and commission (Christensen and Knudsen 2010, Csaszar 2013).

The third assumption becomes visible in constraint studies, which abstract from how structure combines judgments and focus instead on how structure shapes the judgments themselves (Reitzig and Maciejovsky 2015, Joseph and Gaba 2020). They show that structure determines what individuals are willing to pass on, how they interpret information from others, and how honestly they report (McNamara et al. 2002, Fang et al. 2014, Keum and See 2017). The constraint view thereby makes explicit an assumption that the aggregation view leaves implicit, namely that reports are sincere. Where senders distort what they pass on, signal precision is no longer exogenous and errors are no longer independent.

The two tenets rest on these three assumptions. "More is better" requires that errors are independent and that each judgment is more likely to be right than wrong. "Better is better" requires that precision is known, so that better information can be recognized and weighted accordingly. Both require that reports are sincere. We follow the call by Joseph and Gaba (2020) to integrate the aggregation and constraint views, but we hold the motivational concerns of the constraint view in the background and assume for most of the article that senders and receivers are sufficiently motivated. This isolates the most significant change that digitization and genAI introduce, namely how signals about the reliability of a piece of information are generated and transmitted.[1] It also allows us to treat the information itself as a determinant of information

[1] Motivational and social factors may provide important remedies to the problems this change creates, but they are more universal. We also deviate from the aggregation view in two respects. Instead of assuming that structure aggregates binary decisions on independent proposals, we assume that it aggregates pieces of information to inform organizational decisions. Further, we allow pieces of information to be interdependent, meaning that they provide different estimates of the same state of the world. This aligns with the adaptation stream (Joseph and Gaba 2020)

processing, a factor the literature has largely neglected by assuming that information is interpreted, shared, and stored uniformly (Joseph and Gaba 2020, p. 284).

Why Aggregating More Information Can Fail

We are not the first to question the first tenet. Existing explanations for why aggregating more information can fail point to scarce attention (O'Reilly 1980, Simon 1971, Piezunka and Dahlander 2015), ambiguity (Daft and Lengel 1986, Weick 1979), correlated observations and errors (Banerjee 1992, Bikhchandani et al. 1992, Lorenz et al. 2011), or symbolic motives (Feldman and March 1981). Each might explain parts of the verification problem, but its core lies orthogonal to them, in whether the link between a piece of information and the process that produced it can be established and transmitted. The problem therefore persists with unlimited attention, for unambiguous claims, and among sincere senders whose interests are aligned. More attention does not help a receiver who lacks the skills to verify. Ambiguity concerns information that admits multiple interpretations, whereas unverified information can consist of objective but hallucinated facts. In the herding literature, errors become correlated because agents observe each other (Banerjee 1992, Bikhchandani et al. 1992), whereas with genAI they become correlated because senders draw on the same models without observing one another. Finally, symbolic action can be addressed through incentives, whereas the verification problem extends beyond misaligned incentives.

Research on artificial intelligence in organizations has begun to examine how algorithms change organizational decision-making. It has studied the tension between automating and

---

and with the more fundamental idea that organizations need to process more information to cope with environmental uncertainty.

augmenting managerial tasks (Raisch and Krakowski 2021), algorithmic control of work (Kellogg et al. 2020), the allocation of decision rights between humans and algorithms (Shrestha et al. 2019, Puranam 2021), and the use of artificial intelligence as an analogy for organizations (Csaszar and Steinberger 2022). This work typically treats algorithmic output as input of greater or lesser accuracy. It has paid less attention to the possibility that genAI disrupts the signals from which receivers infer how accurate a piece of information is.

None of these explanations considers that the information itself, and what technology does to it, may be a cause of the problem. The assumptions of the screening stream were not naive, however. As we show next, they were reasonable because three conditions held that made verification cheap.

WHY VERIFICATION WAS CHEAP AND WHY IT NO LONGER IS

We develop our argument in four steps. We first specify three conditions under which the assumptions of the screening stream, and with them the two tenets, were reasonable. We then show how digitization and genAI challenge these conditions, derive the consequences for the two tenets, and conclude that verification becomes a problem of organizing in its own right.

Three Conditions for Cheap Reliability

Focusing on the information itself reveals that, in the past, the way information was produced made explicit verification either cheap or unnecessary. Three conditions ensured this, and each made one of the three assumptions of the screening stream reasonable (Table 1). This does not imply that all information circulating in organizations was reliable or correctly judged. It explains, however, why unreliable information was an occasional rather than a systematic

problem. The last two columns of Table 1 anticipate how digitization and genAI challenge each condition, which we discuss below.

Table 1. How Three Conditions Made the Assumptions of the Screening Stream Reasonable, and How Digitization and GenAI Challenge Them

Assumption	Condition that made it reasonable	What receivers obtained for free	Digitization	GenAI
Known signal precision	Costly form	Reliability could be read from the form of a report	Eroded in some settings, as templates, autofill, and remote status signals make form cheap	Removed, as flawless form is produced without epistemic work
Sincere reports	Truth as a shared convergence point	Deviating from the truth required collusion and was eventually exposed	Largely intact, as persistent records make deviations easier to detect	Weakened, as a single sender can produce a consistent false account without collusion
Independent errors	Independent errors	Agreement among sources counted as evidence	Eroded, as copying and social influence correlate nominally independent sources	Eroded, as overlapping training data, recirculated outputs, and self-corroborating checks correlate errors

Costly form. In the past, the form in which information was shared was tightly coupled to the epistemic work required to create it, that is, to acquiring the underlying material, understanding it, and synthesizing it into a form others can understand. By form, we do not merely refer to the format of a document, but its coherence, internal consistency, and specificity. Producing such a form typically required internalizing explicit knowledge and externalizing it again (Nonaka 1994), which is comparatively cheap for those who have done the work and expensive for those who have not. The form therefore served as a signal in the sense of Spence (1973), costly to produce for senders who had not done the epistemic work and cheap to read for receivers. Deficits in the underlying work usually showed as gaps, inconsistencies, or vagueness. For most information, a quick read of the form therefore told receivers how to weight it, and explicit verification was reserved for critical information, if conducted at all.

Truth as a shared convergence point. Whereas costly form made it expensive for a single sender to fake a report, the second condition made it expensive to deviate from what others reported. The reports of different senders about the same matter were anchored in the same reality, which all of them shared and had to keep acting on. The truth therefore served as a focal point (Schelling 1960). Truthful reporting was an equilibrium from which no single sender had an incentive to deviate, because a false report would contradict the reports of others and, eventually, the facts themselves. Coordinated misreporting required collusion, which is deliberate, costly, and difficult to sustain. Receivers therefore had little reason to guard against deliberate misreporting, and a misreport that occurred anyway would be exposed once the facts asserted themselves.

Independent errors. Truth as a focal point deterred deliberate misreporting, but it did not prevent honest errors. The third condition ensured that such errors did not accumulate. Because producing an account required one's own epistemic work, each source drew on its own observations, data, and reasoning, and its errors arose from its own idiosyncrasies. Where sources did share errors, for instance because they relied on the same database, the dependence usually had an identifiable origin. Agreement among sources was therefore itself evidence, because independent accounts that are each more likely to be right than wrong aggregate toward the truth (Condorcet 1785). Receivers could economize on verification by checking whether independent sources agreed and reserve closer scrutiny for cases in which they diverged. Where the truth was not yet observable, agreement served as verification in its place.

Consider a regional sales manager (the sender) who submits a monthly pipeline report to the sales director (the receiver). For a manager who conducted the client conversations, producing a report whose figures reconcile with account-level entries and with notes on specific

conversations is easy. For a manager who had not, producing the same consistency and detail would require reconstructing the conversations, so the director can treat reconciled figures and specific notes as a cheap proxy for verification (costly form). She has little reason to suspect a deliberate misreport, because colleagues' account entries and, eventually, the bookings at the end of the quarter are anchored in the same reality (truth as a shared convergence point). Finally, she can compare the pipeline with estimates from marketing, based on a customer survey, and from finance, based on industry reports. Because these estimates were produced from different material (independent errors), she can treat their agreement as strong evidence without redoing any of the analyses.

Together, these three conditions made verification a byproduct of producing, sharing, and aggregating information. Verification becomes a problem in its own right if one of them fails. Only the failure of the second condition requires deliberate deception. The first and third can arise among fully motivated actors, because a sender may share a polished report without knowing that it is unsupported, and sources may draw on the same flawed material without knowing that they do. The failure of the third condition also undermines redundancy, the mechanism on which existing verification measures most often economize (Landau 1969).

Today's Reality: How Digitization and GenAI Increase the Need for Verification

Digitization, which has been going on for several decades, and genAI, which has recently accelerated it, have become common because of the compelling benefits they offer for organizational information processing, particularly increased speed and volume. Both can also lower the cost of checking a single claim by enhancing transparency, improving the richness of information flows, and easing search and access to new information sources (Goldfarb and

Tucker 2019; see also Dellarocas 2003, Tadelis 2016). However, each also challenges at least one of the three conditions, as the last two columns of Table 1 summarize.

Digitization. Digitization attacks independent errors most directly. The properties that make information cheap to provide, namely open contribution and near-costless replication, also make it cheap to copy, but copies are not independent sources. Exposure to the judgments of others likewise reduces the diversity of estimates without improving their accuracy (Lorenz et al. 2011). Agreement among sources that are copies of a single origin reflects nothing beyond that origin. In some settings, digitization also erodes costly form. Autofill, linked cells, and last period's file carried forward make conformance to a reporting template nearly costless, whether or not the underlying work was done, so that a mistaken figure takes exactly the form a competent one takes.[2] By contrast, digitization largely leaves truth as a shared convergence point intact and may even strengthen it, because persistent digital records lower the cost of tracking and verifying claims (Goldfarb and Tucker 2019).

Generative artificial intelligence. GenAI challenges all three conditions at once. Consider first costly form. A large language model (LLM) produces information that carries every mark of competent work, including fluency, correct formatting, and convincing argumentation, without the epistemic work these marks once certified (Ji et al. 2023). Earlier automation had already cheapened polish, but carelessness still revealed itself through copy-and-paste mistakes or totals that did not reconcile. GenAI removes even these visible errors that once flagged invisible ones, while the consequential errors remain. The polished pipeline report on the

[2] The spreadsheet error in Reinhart and Rogoff's influential analysis of public debt survived in such competent-looking form until the underlying file was obtained and re-run (Herndon et al. 2014). Remote work has a similar effect, because it replaces the ambient signals of colocation with availability indicators and status updates that are nearly free to produce (Cramton 2001, Yang et al. 2022).

director's desk looks exactly as it always did, but its polish no longer testifies to anything. Consider next truth as a shared convergence point. Deviating from the truth once required bringing every colleague who knew better into line. With genAI, a single sender can produce a consistent false account, complete with supporting notes and explanations, at almost no cost and without colluding with anyone. The facts may still assert themselves, but such an account can survive long enough to shape decisions. Consider finally independent errors. Generated outputs differ on the surface, but this variation can mask a common origin. Models are trained on largely overlapping corpora, their outputs recirculate into the data of successor models (Shumailov et al. 2024), and they tend to align their answers with what users appear to believe (Sharma et al. 2024), so asking a model to confirm its own output tests little. Errors may therefore be correlated across sources that appear independent. Whereas digitization copies content that is true somewhere, genAI can produce plausible content that is true nowhere.

The case of *Mata v. Avianca* illustrates how these failures combine. An attorney who could not find supporting precedent in the firm's database asked ChatGPT instead, which returned judicial opinions that did not exist, with fabricated quotations and citations in convincing legal form (*Mata v. Avianca* 2023). By the standards of information processing, the work was organized sensibly. Research and filing were separated, and the brief passed to a second attorney, who signed and filed it. Neither attorney stood to gain from fictitious authority. The failure lay in the information, and it occurred despite an attempt to check it. Asked whether the cases were genuine, the model affirmed that they were. Because the check was addressed to the very system that had produced the claims, the system corroborated itself. The second attorney then performed exactly the quick read that form had always licensed. Structural redundancy therefore added a signature but no verification.

Consequences for "More Is Better" and "Better Is Better"

We can now return to the two tenets. The difficulty that digitization and genAI create is not that verification has become impossible or that checking a single claim has become expensive. It is that reliability is no longer revealed as a byproduct of provision. In *Mata*, the information was processed well by every conventional measure. It reached those who needed it quickly, at low cost, and in the expected form. Its reliability, however, was not low in a way the receiver could see. It was unknown, and the receiver weighted the information as if it were reliable.

If entering information is weighted by how reliable it appears instead of how reliable it is, organizational structures are matched to the wrong signal and increase their errors of omission or commission. The two tenets are affected differently. The second tenet may remain true, because better information still leads to better decisions. However, it becomes useless, because the organization can no longer distinguish information that is better from information that merely looks better. The first tenet no longer holds. When the reliability of additional inputs is unknown, aggregating them adds error that is weighted as evidence. If these inputs share a source, their errors are correlated, and each additional input adds little (Clemen and Winkler 1985). If evaluators who judge by form are more likely to be wrong than right about polished but unsubstantiated information, the Condorcet logic reverses, and each additional judgment moves the collective decision further from the correct one. Like ambiguity, unknown reliability cannot be resolved by aggregating more information of the same kind (Joseph and Gaba 2020, p. 285). It can be resolved only by verification.

Our perspective therefore proposes a scope condition for the information-processing view and the decades of research built upon it. The scope condition is that the three conditions hold, that is, form remains costly to produce, truth remains the shared convergence point, and errors

remain independent. Where they do, reliability arrives with the information, the quick read suffices to infer it correctly, and both tenets retain their validity undiminished. Where one or more of the conditions fails, verification must be designed rather than assumed. This need already arises with digitization. It is particularly acute for genAI, because all three conditions fail together and across domains, and they fail without the visible signs that once alerted organizations to the need for dedicated verification.

Verification as a Problem of Organizing

Evaluating the reliability of information has always been part of information processing, and hence of organizing. The requirement is constitutive in the sense of the universal problems of organizing (Puranam et al. 2014). Any organization whose members act on information that others produced must maintain some mapping between that information and the grounds that warrant it, or its members could not divide labor over information at all. What varies is not whether the requirement exists but whether it binds. While the three conditions held, the quick read met it at almost no cost. Where a condition visibly failed, as in financial auditing, intelligence analysis, or court proceedings, organizations built dedicated verification (Power 1997). Elsewhere, they had no need to. Organizing for verification requires decisions that the quick read made unnecessary, namely who establishes the reliability of a piece of information, by which procedures, and at which point in the workflow. Redundancy, provenance requirements, and the allocation of verification effort have each been studied (Landau 1969, Sagan 2004, Christensen and Knudsen 2010), but not as responses to verification as a binding constraint.

The fabricated citations in *Mata* illustrate the problem rather than define it. AI developers have made considerable progress in reducing hallucinated citations, but the underlying pattern

remains. When senders use genAI outside their area of expertise, they cannot detect its conceptual errors, and the form of its output no longer signals to receivers that such errors exist. This failure becomes an organizational problem through the behavior of senders and receivers, which we conceptualize in the next section as a problem of two blind spots.

A PROBLEM OF TWO BLIND SPOTS

To make information verification a design objective with distinct solutions, we first need to conceptualize the underlying problem. So far, we have argued that digitization and, more acutely, genAI challenge the three conditions that made verification cheap in the past, thereby turning verification into a binding constraint. We showed that the challenge of these conditions is a property of the technological regime and the information itself rather than just a motivational problem. Still, it only becomes an organizational problem through the behavior of the senders who produce information and the receivers who act on it. Therefore, we conceptualize the fundamental verification problem as a sender-receiver problem. When a sender uses genAI to produce a report (i.e., the piece of information that gets aggregated later on) and passes it on to a receiver, there are two verification issues: once between the sender and the genAI model, and once between the receiver and the sender. We refer to these as two blind spots (Figure 1). The first is informational: the sender may not be able to tell whether the model got it right (Dell'Acqua et al. 2026, Randazzo et al. 2025). The second is inferential: the receiver can no longer infer the report's reliability, or the diligence behind it, from its form. The two blind spots are distinct, but they reinforce each other. A receiver who can no longer read reliability from a report's form might simply ask the sender about the accuracy of the report. However, because of the first blind spot, the sender might not know whether the report contains a mistake, and so the

sender's assurance cannot replace the lost signal. Probing whether the sender understands the report can, but such probing is itself a form of costly verification, as we will discuss below.

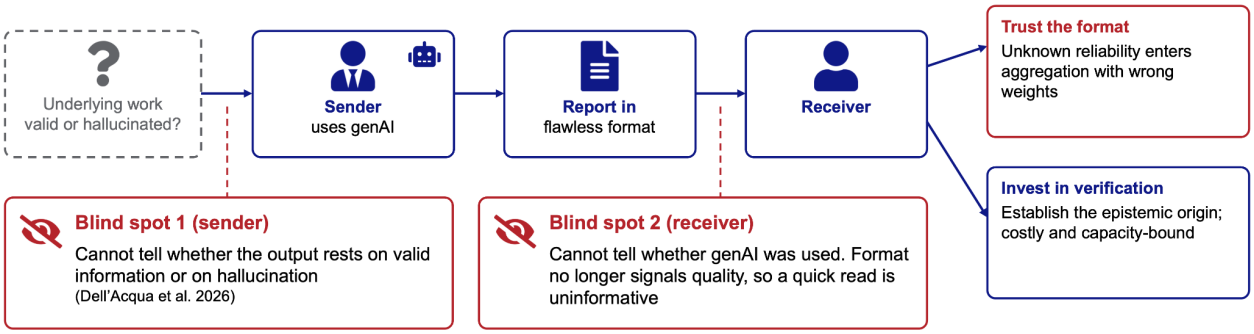


Figure 1. A Problem of Two Blind Spots

The Sender's Blind Spot

The first blind spot arises when a sender extrapolates trust in genAI beyond the domain in which the sender can check its output. It results from three factors. First, the sender calibrates trust in the model on easy tasks within their own domain, where they know the answer or can check it cheaply, and concludes that the technology is trustworthy. Second, the sender then uses it for tasks outside that domain, where they cannot tell which problems are hard for the model and which are easy. Third, the model's capability frontier is jagged: models solve some tasks that experts find difficult and fail on some that a novice would catch, so that even sophisticated users cannot reliably predict where a model will fail (Dell'Acqua et al. 2026, Gans 2026). Moreover, a model's expressed confidence is largely independent of its correctness, so the sender cannot use it as a guide (Randazzo et al. 2025). As a result, the sender cannot verify the model's output precisely where they rely on it most, outside their own expertise. This holds even for senders who are fully motivated to get it right. The blind spot results from limited knowledge and extrapolated trust, not from a lack of effort or attention. This distinction is central to our argument.

For a task outside the sender's domain, four cases can arise. Figure 2 depicts how the jagged frontier interacts with the sender's domain to create the four cases. In the best case, the model gets it right and the sender passes the information on. In the worst case, the model gets it wrong and the sender passes it on without noticing, so that the mistake travels onward carrying the sender's implicit endorsement. Between these sit two intermediate cases: the sender attempts to validate a correct answer but does not trust it, so that correct information is discarded, or the sender attempts to validate an incorrect answer and catches the mistake. Only the best case and the caught mistake leave the organization correctly informed. The first blind spot is therefore not merely a failure to catch errors but creates both errors of omission and errors of commission.

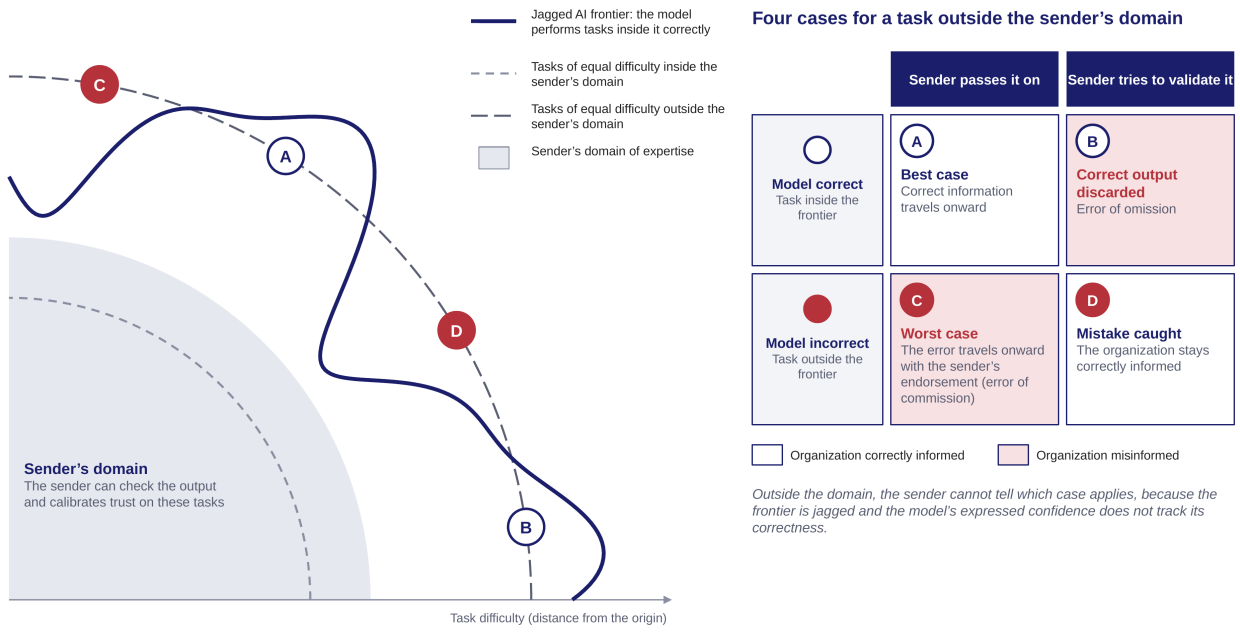


Figure 2. The Sender's Blind Spot: Four Cases

The Receiver's Blind Spot

The report then reaches the receiver. Facing more information than can be evaluated on its merits, the receiver judges a report by the quality of its cover and relies on the Spencian signal we described above. Before genAI, this worked reasonably well, because form and substance

were produced by the same underlying effort and ability: a carefully written report was, more often than not, a carefully researched one. The proxy was imperfect, but it performed well in expectation.[3] GenAI breaks this link. It raises the polish of every report, and disproportionately the polish of reports built on thin work. A uniform rise in polish would preserve the ranking of reports; it is the disproportionate rise for thin work that destroys the information that form used to carry. All reports now "look good," and form no longer reveals which of them deserve weight.

To summarize, the second blind spot is that the quick read no longer reveals a report's reliability. If the receiver fails to realize that genAI was used, unverified information is overweighted and laundered into the aggregation process. However, even knowing that genAI was used does not remove the receiver's blind spot, because such knowledge does not reveal whether a given output can be trusted.

Two Costs

The discussion above yields two costs that should be analyzed separately, because they differ in severity and in duration. The first is biased over-reliance on insufficiently verified reports. It can lead to severe outcomes, as *Mata v. Avianca* illustrates. However, it is transitional: it exists only while receivers still read polish as a signal of quality and act on unverified information. It does not require long chains of transmission, because a single sender–receiver handoff suffices for unverified error to enter the organization; propagation through further layers can amplify the harm, as we explain below. Still, this cost does not persist in equilibrium: it

[3] The receiver also holds other signals: prior knowledge of the sender, and whatever the sender communicates alongside the report. We hold these fixed. Neither substitutes for the lost signal. The prior certifies the person, not the report, and the first blind spot breaks that link on AI-assisted tasks. The sender's assurance is cheap talk that, even when sincere, runs into the first blind spot, as discussed above.

recedes as receivers become familiar with the technology and recalibrate, although errors that entered in the meantime remain in circulation.

The second cost is permanent, and it arises precisely because receivers recalibrate. Recalibration avoids the mistakes that come from overweighting, but trades them for a different problem. Careful senders can no longer use polished form to set themselves apart from careless ones: their work looks alike and can be told apart only through verification. The recalibrated receiver therefore faces two options, both costly. The first is to verify substance directly: reading carefully, probing the sender, checking sources, and redoing parts of the analysis. This is too expensive to do for every report, so it will be reserved for high-stakes decisions. For everything else, the receiver must fall back on the second option: to pool, that is, to give all reports the same average weight in aggregation rather than weighting each by its reliability.

Pooling is cheap, but it creates three problems. First, careful and careless senders are heard equally, so careful work goes unrewarded. Second, careful senders anticipate this: if scrutiny of the LLM's output is not rewarded, the incentive to scrutinize may erode, and the second blind spot feeds back into the first. Third, wrong information still enters the organization's aggregation logic, where it is discounted, but not filtered out. Whether they verify or pool, receivers are thus worse off than they were when the quick read was informative. Because verification is costly (Christensen and Knudsen 2010), this is an unhappy choice for most organizations, and especially for those that lack the capacity to verify: a standing level of unverified output, or a permanently elevated cost of verification. This is the choice as it presents itself to a receiver acting alone within existing structures. Organizations, however, need not leave this choice to individual receivers; they can design for verification, as we discuss below.

From One Bad Report to Organizational Harm

We argue that the two costs discussed above are borne first by receivers and then translate into harm for the organization, in ways that depend on how receivers respond. The receiver's blind spot leaves three responses: to remain unaware, to verify, or to pool. First, where receivers remain unaware, unverified reports receive too much weight in aggregation. Decisions worsen, and the errors settle into the organization's decision-making. This managerial response is short-lived but the most damaging, and the consequences of mistaken decisions outlast the point at which receivers recalibrate. Second, where receivers are aware and verify, verification becomes the bottleneck. Managerial attention shifts from deciding to checking, decisions slow down, and part of genAI's productivity gain is consumed by establishing whether its output can be trusted. In Galbraith's (1974) terms, the hierarchy is overloaded with exceptions, except that every report is now a potential exception. Finally, where receivers pool, decisions become noisier, because reliable reports no longer outweigh unreliable ones. Over time, the inputs themselves deteriorate. Because pooling weights careful and careless reports equally, the extra effort that careful work requires no longer earns its author more influence or recognition. As a result, fewer senders invest in costly extra effort. Juniors, in particular, respond by delegating their tasks to genAI. In doing so, however, they no longer build the expertise and skill that come from performing those tasks themselves, including, eventually, the judgment needed to verify the work of others.

Aggregation amplifies these harms rather than canceling them. As discussed above, aggregation delivers on the promise that more information is better only under two conditions. First, each input must be weighted by its reliability; second, the errors of different inputs must be independent, so that they cancel each other out rather than accumulate. GenAI undermines both. Weights become wrong where receivers remain unaware and flat where receivers pool. And

because senders across the organization draw on the same few genAI models, their errors are correlated even when the senders never observe one another. The more departments and levels use genAI, the more likely it becomes that a given decision rests on at least one flawed input. Moreover, adding further inputs does little to improve the final outcome. Worse, copies of the same error can masquerade as independent confirmation, so that the error appears more reliable with each repetition.

The harm from unverified reports also accumulates over time, because organizational knowledge is cumulative. Reports are rarely consumed once. They become premises for further work, and each typically feeds several downstream uses. Downstream actors then optimize conditional on these premises, so that a flawed starting point shapes every decision built on it, and correcting it later means unwinding all of these decisions. The longer an error survives, the larger this correction bill becomes. Some uses also embed the error in a stock, such as a database, a codebase, an organizational routine, or eventually the training data of the models themselves (Shumailov et al. 2024). From there, the error taxes every future use that draws on the stock.

Before genAI, two checkpoints kept errors from accumulating in the organization's stock of knowledge, and they did so at no extra cost. First, errors tended to betray themselves on receipt, because the quick read screened out careless work. Second, they tended to surface in use, because building on a flawed premise usually made someone stumble. GenAI removes the first checkpoint, as discussed above. The second checkpoint is weakened as well: a large language model that produced a flaw is poorly placed to detect the flaw later, as *Mata v. Avianca* showed, and fluent machine processing removes the friction over which human users used to stumble. In addition, where further downstream steps also deploy genAI, each step can add errors of its own, so that more errors accumulate along the chain while fewer of them are caught.

How far errors propagate therefore depends on whether the use of information confronts ground truth. Some uses do so by themselves: code either runs or fails, and accounts either reconcile or do not. In such settings, reality limits how long an error can go unnoticed, and mistakes due to genAI will be caught sooner, if potentially through a costly failure. Many organizational uses of information, however, never meet such a test, or meet it only after long delays, such as strategy memos, market assessments, policy analyses, and the internal reports on which they build. There, errors that once died out within a step or two now travel further, because neither checkpoint stops them anymore, and settle into the organization's stock of knowledge.

In sum, because genAI disables the quick read, errors that would once have been screened out at the first handoff now either require costly verification or enter the organization's stock of knowledge. There, the number of decisions resting on an error grows, the cost of correcting the error grows with time, and each repetition lends the falsehood a misleading air of confirmed reliability. All three effects follow from the same source: the reliability of a report can no longer be read off its form. This leaves organizations with the three harms we have developed in this section. The first is worse decisions, wherever unverified information is overweighted or pooled. The second is a permanently higher bill for verification, the resources organizations must now spend to keep unreliable reports out of circulation. The third unfolds more slowly: an erosion of the organization's knowledge base, as errors accumulate in its stock and the effort and expertise that once produced reliable inputs decline. How well organizations can contain these harms depends on how they are organized, the question to which we now turn.

Can AI Be the Remedy?

A natural response for management seeking to cut the verification bill is to use AI for verifying the information that senders provide, that is, to use AI to catch AI. We are skeptical that this approach works in its naive form and without reorganizing for verification. For the reasons discussed above, a large language model is poorly placed to recognize the very mistake it has just made, and the errors of systems trained on similar data are correlated, so that even a second opinion from a second LLM is less independent than it appears.

Systematic strategies that leverage LLMs for verification may fare better. Candidates include independent verification by alternative LLMs, adversarial review, and verification targeted at task types where the frontier is known to be jagged. Whether such strategies succeed depends on two parameters: how often the verifying LLMs err, and how strongly their errors correlate with those of the LLMs they check. Implementing adversarial review and comparable arrangements, however, is itself an act of organizing; it is precisely what we mean by organizing for verification. We return to such designs in the Discussion.

EXISTING FORMS OF ORGANIZING AND THE VERIFICATION PROBLEM

To show its analytical usefulness, we now apply our framework of two blind spots to existing forms of organizing to show what consequences it reveals. The two blind spots arise at every handoff between a sender and a receiver. What follows from them, however, depends on the form of organizing in which sender and receiver are embedded. We therefore ask how different forms respond once information verification rather than provision becomes the binding constraint. To answer this question, we analyze four stylized forms of organizing, namely hierarchies, communities of practice, markets, and open source software communities. The selection is not exhaustive, since our objective is to demonstrate the framework's diagnostic

value rather than to validate it empirically. For each form, we compare its verification capacity in the old regime, in which reliability was a cheap byproduct of provision, with its capacity in the new one, in which verification is becoming a new binding constraint. We again take genAI as the driving force, since it constitutes the acute case.

In the past, when information provision, that is, supplying members with the information they need to act, coordinate, and contribute to organizational decision-making, was the bottleneck of organizations' information-processing capacity, organizations could rely on two basic solutions (March and Simon 1958, Puranam et al. 2012, 2014). They could cultivate common ground, so that members share enough knowledge to anticipate one another's actions (Clark 1996, Srikanth and Puranam 2011), or they could reduce the need for such knowledge through modularity and standardized interfaces (Sanchez and Mahoney 1996, Baldwin and Clark 2000). We refer to forms that rely primarily on the first solution as shared-cognition-based forms and to forms that rely primarily on the second as interface-based forms. As we argued above, the act of providing information also served as a signal to gauge its reliability. Accordingly, each solution shapes how receivers infer the reliability of the information they aggregate, and hence the weight this information receives. In shared-cognition-based forms, receivers judge reliability behind the signal. They assess whether an output could have issued from competent work, drawing on their own knowledge of how such work is produced. In interface-based forms, receivers judge reliability at the signal. They check whether an output conforms to a predefined standard, such as a price, a specification, a credential, or a passing test. Neither mode was designed for verification.

While provision was the bottleneck, this arrangement was reasonable, because the three conditions made both modes informative at almost no cost. Once verification becomes the

bottleneck, our framework reveals that the way a form solves provision determines which of the receiver's responses remain feasible when the quick read fails, namely remaining unaware, verifying, or pooling, and whether the form's verification capacity reproduces itself. Hierarchies and markets are natural anchors for this analysis, since the screening stream itself began by comparing them, with polyarchy as a stylized representation of a market economy (Sah and Stiglitz 1986). We treat communities of practice as the pure shared-cognition-based form and markets as the pure interface-based form. Hierarchies add authority to shared cognition, and open source communities combine an interface layer with a review layer that operates behind the scenes.

We analyze each form in five steps. First, we describe how the form solves the problem of information provision. Second, we identify how its receivers infer reliability and which of the three conditions this inference rests on. Third, we locate where genAI enters production and where the sender's blind spot arises. Fourth, we examine what the quick read loses and which responses the form leaves to its receivers. Fifth, we ask whether the form's verification capacity is replenished or diminished under genAI. After each family of forms, we derive the pathology that follows from the fourth and fifth steps. Figure 3 shows how the two blind spots enable unverified information to enter each form of organizing.

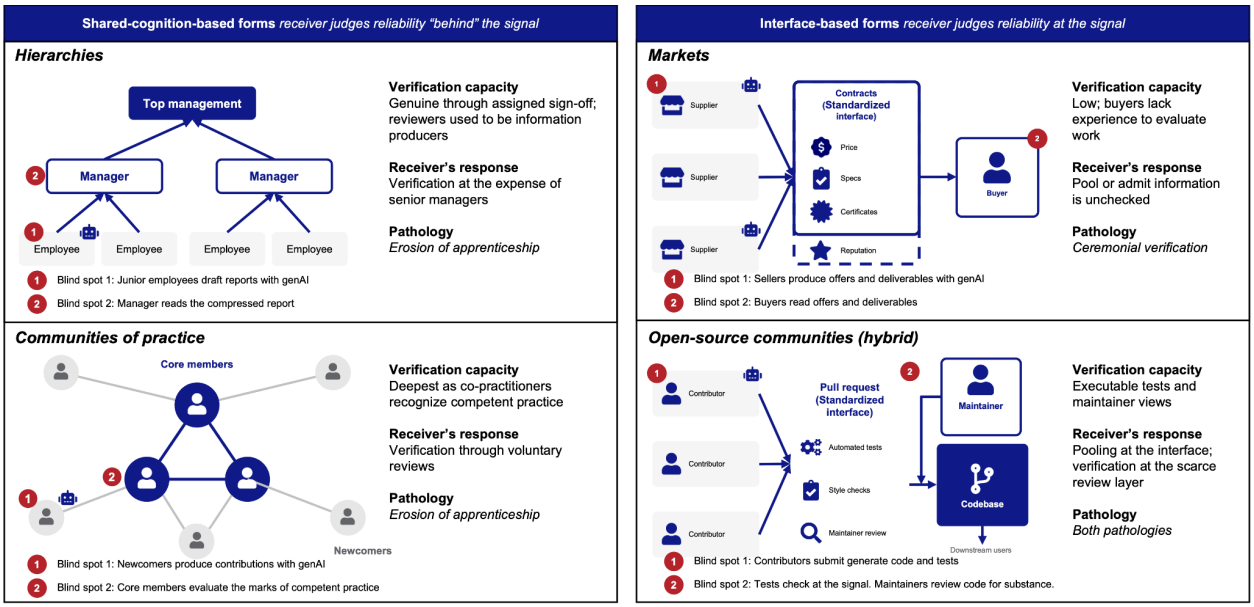


Figure 3. Forms of Organizing and the Verification Problem

Shared-Cognition-Based Forms

Hierarchies

Hierarchies solve the provision problem through authority, which designs the distribution of information, and through plans and standards that reduce the need for information, while colocation and shared tenure build common ground (Galbraith 1974, Puranam et al. 2014). Because managerial attention is bounded (Simon 1947), hierarchies compress work into legible standard formats and route the compressed version upward. Receivers infer the reliability of these reports behind the signal, because reviewers were once producers and can judge whether a report could have issued from competent work. Authority adds three features that no other form offers. It can assign verification, its sign-off establishes an addressable warranting agent, and its layers provide redundancy (Landau 1969). Expertise is also concentrated where exceptions are resolved (Garicano 2000). This capacity rests on costly form, because the compressed report carried the reliability signal, and on independent errors, because each layer constituted an independent draw.

GenAI enters at the bottom of the hierarchy, where junior members draft reports, often on tasks outside their own domain (Dell'Acqua et al. 2026). The compressed report becomes uninformative, because it now looks the same whether or not the work behind it was done. Layered redundancy offers less protection than it appears to, because layers that consult the same model cease to be independent draws. Unlike receivers in the other forms, managers can still verify behind the signal. The cost, however, falls on senior members, and every report becomes a potential exception that overloads the hierarchy (Galbraith 1974). *Mata v. Avianca* shows what remains when a sign-off lacks independent access to the underlying facts. The redundancy operated as designed, but it added a signature rather than verification.

In the long run, a hierarchy's verification capacity proves endogenous to its own promotion ladder. Senior managers can verify behind the signal because they produced the relevant information themselves under the previous regime. The junior tier that would supply tomorrow's reviewers, however, has offloaded production. Recent evidence that organizations deprioritize hiring for precisely the skills genAI can substitute for, including information provision and monitoring, suggests that this erosion is already under way (Iyer et al. 2025). Authority can mandate verification effort, but might compromise future verification capacity for short-term efficiency gains.

Communities of Practice

Communities of practice solve the provision problem without relying on authority or interfaces. Members coordinate through mutual engagement, joint enterprise, and a shared repertoire (Brown and Duguid 1991, Lave and Wenger 1991, Wenger 1998). Their receivers possess the deepest capacity to judge behind the signal of the four forms. Contributions are assessed by co-practitioners who have performed the same work, and the assessment proceeds

not by checking output against a rule but by recognizing the marks of competent practice, a tacit judgment available only to someone who has practiced (Polanyi 1966, Nonaka 1994). This capacity rests mainly on costly form, since the marks of competent practice were hard to produce without the practice itself.

GenAI enters wherever members, and newcomers in particular, generate contributions they could not produce unaided. Polished contributions initially pass as competent practice. Once members recalibrate, reviewing every contribution with care becomes costly, and because review in a community is voluntary, nobody can assign it. GenAI also affects how the community reproduces itself. Under legitimate peripheral participation, newcomers perform the simple, well-specified, and low-stakes portions of the work and thereby acquire the judgment that eventually makes them competent verifiers (Lave and Wenger 1991). These are the tasks on which genAI helps novices most (Noy and Zhang 2023, Brynjolfsson et al. 2025). They are economically least valuable but developmentally most valuable, so the community's efficiency-rational decision to offload them is simultaneously the decision to stop producing its own verifiers (see also Beane 2019). Because nobody in the community holds authority, nobody monitors the remaining stock of judgment.

Pathology 1: The Erosion of Apprenticeship

Both shared-cognition-based forms respond to the verification problem genAI creates in the same way, because their receivers can still verify behind the signal. In the short run, receivers pass through the two phases described above. Polished contributions initially receive too much weight, and once receivers recalibrate, they must either verify at a cost or pool. In the medium run, both forms retain their verification capacity, because senior members trained under the previous regime can still verify behind the signal. Doing so, however, imposes a higher

verification load on precisely the members whose time is scarcest, so that verification competes with the work it is meant to protect. In the long run, the capacity itself erodes. Senior members acquired their judgment by producing the relevant class of information themselves. GenAI substitutes for that production at the junior level, thereby shortening the path by which members grow into senior roles (Lave and Wenger 1991, Beane 2019). We refer to this pathology as the erosion of apprenticeship.

The erosion is delayed, cohort-paced, and difficult to reverse. It also conceals itself, because the incumbent stock of judgment masks the severed flow for as long as its holders remain. By the standards of information processing, it is even invisible. Offloading production to genAI is an unambiguous improvement in provision, since it is faster, cheaper, and richer. The cost registers only once verification is treated as a design objective in its own right, and otherwise it is classified as a training or human capital problem lying outside organization design. Shared-cognition-based forms thus trade better provision now at the price of worse verification later. The erosion of apprenticeship also turns the sender's blind spot into a structural condition. Recognizing where the jagged frontier lies is itself a practice-based judgment (Dell'Acqua et al. 2026), because when nobody produces tasks unaided, every task effectively lies outside the sender's domain.

Interface-Based Forms

Markets

Markets solve the provision problem by economizing on knowledge so thoroughly that buyers need not see inside the work they purchase (Hayek 1945). We focus on markets for knowledge work, in which the traded object is itself information, as in freelance writing, translation, software development, consulting, or research services. In such markets, sellers are

senders and buyers are receivers. The contract specifies what is to be delivered. Knowledge work is an experience or even a credence good, whose quality becomes known only after use, if at all (Nelson 1970, Darby and Karni 1973). To reduce the resulting information asymmetry (Akerlof 1970), markets have built an apparatus of devices around the price system, including tailored offers, credentials, gatekeepers, reputation systems, warranties, and professional liability (Dranove and Jin 2010, Tadelis 2016).

Receivers in markets infer reliability at the signal at two stages. Before contracting, tailored offers, credentials, gatekeepers' screening, and reputation signal the seller's ability. At delivery, buyers perform a quick read of the delivered work against the specification. Both stages rest on costly form. Reputation also rests on independent errors, because a rating system aggregates the judgments of buyers who assessed the work separately (Dellarocas 2003). Buyers rarely look behind the signal. At least before a first purchase, they share no history with the seller, and they often buy the work precisely because they lack the expertise to produce, and hence to judge, it themselves.

GenAI enters where sellers produce offers and deliveries. At the offer stage, a tailored offer once worked as a signal in the sense of Spence (1973), because producing one required understanding the task, which was cheaper for able sellers than for others. Once genAI makes such offers nearly free, they stop separating able from less able sellers. Buyers first overweight polished offers and, once they recalibrate, pool, treating sellers as if they were of average quality. Pooling invites the adverse selection that the apparatus was meant to prevent (Akerlof 1970). Galdin and Silbert (2025) show that employers on Freelancer.com paid more for workers whose applications were tailored to the job before LLMs became available, but that the return to tailoring collapsed once LLMs were widely used. In their counterfactual, in which written

applications no longer signal ability, workers in the top quintile of the ability distribution are hired 19% less often and those in the bottom quintile 14% more often. At delivery, the quick read becomes uninformative, because a careless delivery no longer looks careless. Reputation systems carry this failure back to the offer stage. A rating records the buyer's quick read at delivery rather than the work's eventual quality, and because polish is available to every seller, ratings are misled in the same direction. Their errors are correlated, so their agreement adds less evidence than their number suggests (Clemen and Winkler 1985). GenAI also makes plausible fake reviews nearly free (Luca and Zervas 2016). Ex post remedies, finally, apply only once a defect has been discovered and attributed to the seller, which becomes harder when the defect surfaces long after delivery.

Buyers cannot respond by verifying behind the signal. Their incentive to verify is as strong as it gets, since they bear the cost of not checking, but they lack the capacity to judge. The fallbacks that markets offer are weakened as well. Purchased verification, such as an audit, an inspection, or due diligence, is itself knowledge work, increasingly produced with genAI and bought in markets that face the same problem. Staged contracting, in which buyers learn about sellers on small tasks before raising the stakes (Pallais 2014), builds trust on exactly the small, well-specified tasks that genAI handles best, so that trust may be extended to larger tasks beyond the jagged frontier (Dell'Acqua et al. 2026). Moving the transaction into the firm, the classical response when measuring what is traded becomes too costly (Barzel 1982, Williamson 1985), meets the erosion of apprenticeship described above and therefore buys time rather than solving the problem.

Pathology 2: Ceremonial Verification

Interface-based forms respond to genAI differently from shared-cognition-based forms. Their checks are, in effect, a quick read written into a rule. An interface admits a contribution if its form conforms to a standard, much as a receiver accepts a report whose form looks right. Like the quick read, the interface worked only as long as producing the right form required doing the work. When the quick read fails, receivers in these forms cannot fall back on verification behind the signal. Once they recalibrate, they can only pool or admit contributions unchecked, and the interface keeps running. It consumes resources, still catches crude errors, and still counts as verification, but it confirms conformance rather than substance. We refer to this pathology as ceremonial verification. Verification turns into ritual (Power 1997), a formal structure that conforms to legitimated expectations while remaining decoupled from the activity it governs (Meyer and Rowan 1977). It resembles means-ends decoupling, in which practices are implemented but disconnected from the outcomes they are meant to achieve (Bromley and Powell 2012). Here, however, the decoupling is epistemic rather than institutional, and it requires no intent.

Ceremonial verification sets in immediately, because the interface ceases to separate competent work from its imitation as soon as genAI makes conformance free. It conceals itself through the check's own verdict, which creates the illusion that verification has taken place. Receivers therefore recalibrate slowly, especially where failures surface late and far from the check. In markets, pooling additionally looks like ordinary price formation. Ceremonial verification is the receiver's blind spot hardened into routine.

A Hybrid Form: Open Source Software Communities

Open source software communities solve the provision problem through modular architecture, self-selection, and the fact that the code is its own representation, while documentation and version control substitute for shared tenure (Lerner and Tirole 2002, von Hippel and von Krogh 2003, Puranam et al. 2014). Their receivers infer reliability in two layers. Contributions first pass through automated tests and continuous integration, which check at the signal. Maintainers who know the codebase in detail then review them behind the signal. Of the four forms, open source communities come closest to treating verification as an explicit design objective. They institutionalize redundancy by exposing each change to inspection by many contributors (Raymond 1999). Their commit history records the author, content, and timing of every change and thereby approximates the mapping to epistemic origin that our definition of verification requires. Their tests render checking executable. Finally, the code provides a partial non-social arbiter, since whether it executes does not depend on anyone's judgment, even if whether it executes correctly still does. In this respect, open source communities preserve part of the second condition, truth as a shared convergence point.

GenAI enters at the contribution stage. Where GitHub Copilot became available, open source contributions rose by 28% to 40%, and incremental changes to existing code rose far more than contributions adding new functionality (Yeverechyahu et al. 2025; see also Peng et al. 2023). At the interface layer, a contribution with passing tests and clear documentation was once costly to produce without understanding the code. This no longer holds. When the same model writes both the code and its tests, both inherit the same misconception, so that the third condition, independent errors, fails inside the test suite. Plausible documentation becomes nearly free and describes what the code should do rather than what it does. The checks still run and still pass, but they confirm conformance rather than substance. At the review layer, the ratio between

contributions and review collapses, because contributions have become cheap, whereas review remains voluntary and scarce (Eghbal 2020). In the long run, the capacity to review erodes as well. Maintainers typically grew into their role by contributing themselves, beginning with small, well-specified fixes to existing code (von Krogh et al. 2003, Dahlander and O'Mahony 2011). These are the contributions genAI produces most readily, so the path to maintainership is cut. Self-selection, which used to keep contributors away from tasks beyond their competence, fails too, because contributors who submit generated code may not know where the jagged frontier lies.

Open source communities therefore suffer from both pathologies, and the two reinforce each other. The interface still rejects code that does not compile, or that breaks obvious functionality, but it no longer teaches new contributors the codebase. When genAI produces conforming contributions, passing the interface teaches nothing, so the ceremony at the interface feeds the erosion of apprenticeship at the review layer. Hybrid forms thus do not necessarily average the two pathologies. Instead, they may suffer from both.

Predictions

Beyond classifying the four forms, our framework allows predictions about how forms of organizing respond to genAI, when their verification fails, and how the failure can be observed.

First, the way receivers infer reliability predicts how a form responds. Forms that judge behind the signal respond to genAI by verifying, so we expect the verification load on senior members to rise and decisions to slow (Galbraith 1974). Forms that check at the signal respond by pooling, so we expect the returns to signals of quality to flatten and adverse selection to increase (Akerlof 1970, Galdin and Silbert 2025).

Second, the two pathologies differ in timing. Ceremonial verification sets in as soon as genAI makes conformance free. The erosion of apprenticeship lags by cohorts, so the quality of review should decline with the share of reviewers whose formative work was AI-assisted, with a delay of years rather than months.

Third, the pathologies leave different signatures. Ceremonial verification predicts that checking activity stays constant or rises while downstream errors rise. The erosion of apprenticeship predicts a rising verification load per senior member, followed by declining verification quality once incumbents leave.

Fourth, hybrid forms suffer from both pathologies. We expect the erosion of apprenticeship to proceed faster in hybrids than in pure shared-cognition forms, because the ceremonial interface no longer trains the contributors who would become reviewers.

Finally, both pathologies should be less severe where ground truth asserts itself quickly, as when code runs or accounts reconcile, that is, where the second condition still holds. They should be most severe where errors surface late or never, as in strategy memos, market assessments, and policy analyses. However, none of the four organizational forms produces the independent bits of information on which verification by corroboration rests, because each inherited reliability signals from a regime in which they came free. Organizing for verification requires manufacturing independence deliberately, a point to which we return in the Discussion.

DISCUSSION

Our essay proposes a shift in perspective on organizational information processing, from the provision of information to its verification. The traditional information-processing view treats the reliability of information as given and asks how organizations can process more of it. We have shown why this was reasonable for most of organizational history, as long as three

conditions were satisfied: (i) A convincing form required expertise or care (costly form), (ii) truth served as a common point of convergence, and (iii) independent sources erred independently. While these three conditions held, the form in which information arrived carried a free reliability signal, which was good enough for the effective aggregation of information. Verification beyond a quick read rarely paid, and provision was costly and hence the scarce resource. Where these conditions fail, verification becomes the scarce resource and a design objective in its own right.

A recent episode illustrates the shift. In September 2026, the Washington Post reported that a single business-school professor had authored or co-authored more than 200 working papers in nine months with the help of generative AI (Wallack 2026). In the same report, a senior editor of *Organization Science* estimated that submissions to the journal had doubled since 2022 and warned "that the system is not equipped right now to handle" the increased volume of research. Producing papers has become cheap; reading, refereeing, and judging them has not.

In the following, we outline the opportunities this shift opens for theory development and empirical research, its implications for practice, and its limitations.

New Opportunities for Theory Development

Treating reliability as endogenous rather than given invites revision of several established theories. We illustrate this with four examples.

From Span of Control to Span of Verification. Research on organization design explains the number of hierarchical levels and the span of control by managers' capacity to process information (Simon 1947, Galbraith 1974). Knowledge hierarchies trade off the cost of communicating against the cost of acquiring knowledge (Garicano 2000), and technologies that lower either cost change hierarchies in predictable ways (Bloom et al. 2014). A common

expectation is therefore that genAI will flatten hierarchies by automating the aggregation and reporting that occupy middle managers. A verification perspective qualifies this expectation. When the quick read fails, a manager's span is limited by the number of reports she can verify rather than the number she can process, and the span of verification falls below the span of control. We therefore expect hierarchies to become flatter where errors surface quickly in use or where checks can be made executable and independent, because there the gains in provision dominate. Where errors are consequential and surface late, we expect additional levels or dedicated verification roles. Such levels add value only if they are independent. Accordingly, we expect verification layers that rely on different sources and tools rather than replicated review. Finally, where genAI absorbs junior work, hierarchies may hollow out, trading short-term savings for a future shortage of the members who verify.

From Coordination Costs to Verification Costs. Transaction cost economics explains the boundary of the firm by asset specificity and opportunism (Williamson 1985), and the measurement-cost view adds that firms integrate when the quality of what is traded is costly to measure (Barzel 1982). Because information technology lowers coordination costs, it was expected to shift transactions toward markets (Malone et al. 1987). GenAI lowers the costs of coordination and provision even further, but it raises the cost of verifying quality at the interface, even between parties who are not opportunistic. This is measurement-cost logic without opportunism. We therefore expect integration to increase in exchanges whose quality is hard to verify at the interface and surfaces late, such as consulting, research, and analytics, and outsourcing to continue where outputs can be tested in use. Firms would integrate not to control opportunism but to secure an independent reference and the capacity to verify behind the signal, which extends the view of the firm as a repository of knowledge (Kogut and Zander 1992). As

our analysis of markets suggests, however, integration buys time only as long as the firm's own apprenticeship remains intact. Alternatively, firms may buy verification separately from the work, through certification and third-party review (Dranove and Jin 2010). When firms make, buy, or verify, and how markets for verification can guarantee their own independence, are open questions.

From Inherited to Designed Verification. Existing research has studied verification mechanisms such as redundancy (Landau 1969, Sagan 2004), sequential screening (Sah and Stiglitz 1986), and audit (Power 1997) largely in isolation. Our perspective suggests classifying them by the condition they restore. Mechanisms that restore costly form recouple signals to the underlying work, for instance by probing the sender, requiring evidence of the process, or attaching a named and accountable warrant to a sign-off. Mechanisms that restore truth as a shared convergence point create provenance, such as records that trace claims to their sources or logs of human–AI co-production. Mechanisms that restore independent errors design diversity, for instance by assigning verifiers who rely on different models, data, or training lineages and rewarding them for finding errors rather than for agreeing. In this sense, using AI to verify AI is a question of organization design rather than of technology, because its value depends on how strongly the verifier's errors correlate with those it checks. Each mechanism raises questions about its organizational consequences, including who bears the cost of verification, what status and career value verification work carries, how unobservable verification effort can be rewarded, and how the mechanisms themselves can be kept from turning ceremonial.

From Four to Five Problems of Organizing. Puranam et al. (2014) argue that every organization must solve four universal problems, namely dividing tasks, allocating them, providing rewards, and providing information, and that new forms of organizing are new

solutions to these problems. Our analysis suggests that, before genAI, each form solved verification alongside a different one of them. In hierarchies and communities of practice, verification was solved alongside information provision, in open source communities alongside task division and allocation, and in markets alongside reward provision, since counterparties check because they pay. Once the free signal is gone, verification may become a fifth universal problem, a dimension of information provision, or a constraint on all four. If it is a problem in its own right, forms of organizing can be characterized by how they solve it, and new forms may emerge around it. Candidates include organizations specialized in verification, such as AI audit firms, crowd-based verification with stakes, such as bug bounties, and communities that restrict AI-generated contributions. The perspective also extends to knowledge creation. In Nonaka's (1994) account, making tacit knowledge explicit was costly, which tied explicit knowledge to a tacit origin. When externalization becomes free, organizations accumulate explicit knowledge that nobody holds tacitly. Organizational sensemaking, which is "driven by plausibility rather than accuracy" (Weick 1995, p. 55), then loses the conditions under which plausibility was a valid proxy for accuracy.

Challenges and Opportunities for Empirical Research

Verification is harder to study than provision, because it leaves traces mainly when it fails. We see four challenges, each of which offers opportunities for new empirical studies.

Measuring reliability when the truth arrives late. Because the quick read no longer reveals reliability, studies need outcome-linked measures that are observed with a delay, such as rework, reverted code, retractions, or sanctions for fabricated citations. Errors of omission, that is, correct outputs that were discarded, rarely leave traces at all. They require designs in which

the ground truth is known, such as experiments that vary where a task lies relative to the jagged frontier and the sender's domain and measure calibration rather than accuracy.

Measuring and signaling provenance. Self-reports of AI use are unreliable, detectors err frequently and disproportionately flag non-native writers (Liang et al. 2023), and watermarks are fragile. Process data such as version histories and prompt logs offer new traces of epistemic origin. Whether such traces can restore a reliability signal that separates careful from careless work is itself an empirical question, which signaling experiments and natural experiments around the diffusion of genAI can address (Galdin and Silbert 2025).

Measuring independence. Corroboration is informative only if sources err independently. Research can measure how far apparently independent sources share models or data and how strongly their errors correlate, tracking claims rather than organizations (Hsiao and Schneider 2021, Shumailov et al. 2024).

Observing slow pathologies against a moving target. Settings in which verification is recorded, such as code review, audit, and editorial review, allow tests of the signatures of the two pathologies. Ceremonial verification predicts checking activity that stays constant or rises while downstream errors rise. The erosion of apprenticeship predicts declining review quality with the share of reviewers whose formative work was AI-assisted, lagged by years, which calls for longitudinal data. Because model capabilities change, hypotheses should be stated in terms of the three conditions rather than specific models. At the same time, the window for collecting baselines from before the adoption of genAI is closing.

Implications for Practice

Organizations can either bear the costs of unverified information or organize for verification. Organizing for verification means deciding who establishes the reliability of a piece

of information, by which procedure, at which point in the workflow, and against which independent reference, decisions that the quick read used to make for free. Our conceptualization of the verification problem reveals five immediate implications for organizations. First, verification needs its own resources, metrics, and responsibilities, because it otherwise competes with production for time and budget. Second, each class of consequential information needs a designated warrant that concerns substance rather than form. Third, independence must be designed, since two checks drawn from the same model are one check. Fourth, organizations should protect the junior work through which future verifiers acquire their judgment. Fifth, checking activity that stays constant while downstream errors rise is a warning rather than an achievement. How much verification an organization should buy depends on how quickly and at what cost its errors surface. Organizations whose errors reveal themselves in use may rationally verify little, whereas those whose errors surface late or in front of customers and regulators require substantially more.

Limitations and Boundary Conditions

Our core argument is that verification becomes a new binding constraint that requires an organizational response. This argument has several important boundaries. First, it applies only where at least one of the three conditions fails. Where a convincing form still requires the underlying work, the quick read remains informative and more information remains better. Second, we assumed motivated senders and receivers to isolate the informational mechanism. Strategic uses of genAI, such as disguising its use or simulating effort, are likely to aggravate the problem, and how verification interacts with incentives remains open. Third, our four forms are stylized and illustrative rather than exhaustive, and real organizations combine them. Fourth, we treated the capabilities of genAI as given. Better calibration would narrow the sender's blind

spot, and more diverse model lineages would restore part of the lost independence. Only an AI that never erred and was perfectly aligned with the organization would render verification unnecessary. Short of that limit, an organization that delegates verification to a model has still decided who warrants what it believes. Fifth, our argument is conceptual. A formal model could specify, for instance, the threshold at which the Condorcet logic reverses. Finally, we focused on knowledge work and honest error and set aside deliberate manipulation, such as the injection of adversarial content.

CONCLUSION

Organizations have always depended on information they could not fully check. For most of their history this did not matter, because the form in which information arrived carried a free reliability signal: a convincing form required expertise or care. Moreover, truth served as a common point of convergence, and independent sources erred independently. As long as these three conditions were satisfied, the free reliability signal was good enough to ensure the successful aggregation of information. Verification beyond a quick read that determined reliability rarely paid. Digitization, and generative AI in particular, remove these conditions.

The central consequence, and the core of this essay, is that verification replaces provision as the scarce resource in organizational information processing. We traced this consequence to two blind spots, at the sender who produces information and at the receiver who judges it, and we showed how these blind spots turn into costs for receivers and harms for organizations once unverified information is aggregated and accumulated. We then applied the framework to four existing forms of organizing knowledge work, hierarchies, communities of practice, markets, and open source communities, and found that none of them acquired its verification capacity on purpose, so that each is prone to failure. The pathologies we identified are ceremonial

verification in markets, the erosion of apprenticeship in hierarchies and communities of practice, and both in open source communities. We conclude that the information-processing view needs to be extended by a scope condition. Where the free signal is gone, every organization must treat verification as a design objective in its own right and organize for verification. Doing so requires deciding who warrants what the organization believes, on what evidence, and against which independent reference.

REFERENCES

Akerlof GA (1970) The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84(3):488–500.

Baldwin CY, Clark KB (2000) *Design Rules, Vol. 1: The Power of Modularity* (MIT Press, Cambridge, MA).

Banerjee AV (1992) A simple model of herd behavior. *Quarterly Journal of Economics* 107(3):797–817.

Barzel Y (1982) Measurement cost and the organization of markets. *Journal of Law and Economics* 25(1):27–48.

Beane M (2019) Shadow learning: Building robotic surgical skill when approved means fail. *Administrative Science Quarterly* 64(1):87–123.

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.

Bloom N, Garicano L, Sadun R, Van Reenen J (2014) The distinct effects of information technology and communication technology on firm organization. *Management Science* 60(12):2859–2885.

Bromley P, Powell WW (2012) From smoke and mirrors to walking the talk: Decoupling in the contemporary world. *Academy of Management Annals* 6(1):483–530.

Brown JS, Duguid P (1991) Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organization Science* 2(1):40–57.

Brynjolfsson E, Li D, Raymond L (2025) Generative AI at work. *Quarterly Journal of Economics* 140(2):889–942.

Burton RM, Obel B (1984) *Designing Efficient Organizations: Modelling and Experimentation* (North-Holland, Amsterdam).

Burton RM, Obel B (2004) *Strategic Organizational Diagnosis and Design: The Dynamics of Fit*, 3rd ed. (Kluwer, Boston).

Christensen M, Knudsen T (2010) Design of decision-making organizations. *Management Science* 56(1):71–89.

Clark HH (1996) *Using Language* (Cambridge University Press, Cambridge, UK).

Clemen RT, Winkler RL (1985) Limits for the precision and value of information from dependent sources. *Operations Research* 33(2):427–442.

Condorcet M de (1785) *Essai sur l’application de l’analyse à la probabilité des décisions rendues à la pluralité des voix* (Imprimerie Royale, Paris).

Cramton CD (2001) The mutual knowledge problem and its consequences for dispersed collaboration. *Organization Science* 12(3):346–371.

Csaszar FA (2013) An efficient frontier in organization design: Organizational structure as a determinant of exploration and exploitation. *Organization Science* 24(4):1083–1101.

Csaszar FA, Steinberger T (2022) Organizations as artificial intelligences: The use of artificial intelligence analogies in organization theory. *Academy of Management Annals* 16(1):1–37.

Cyert RM, March JG (1963) *A Behavioral Theory of the Firm* (Prentice-Hall, Englewood Cliffs, NJ).

Daft RL, Lengel RH (1986) Organizational information requirements, media richness and structural design. *Management Science* 32(5):554–571.

Dahlander L, O'Mahony S (2011) Progressing to the center: Coordinating project work. *Organization Science* 22(4):961–979.

Darby MR, Karni E (1973) Free competition and the optimal amount of fraud. *Journal of Law and Economics* 16(1):67–88.

DeGroot MH (1970) *Optimal Statistical Decisions* (McGraw-Hill, New York).

Dell'Acqua F, McFowland E III, Mollick ER, Lifshitz-Assaf H, Kellogg KC, Rajendran S, Krayer L, Candelon F, Lakhani KR (2026) Navigating the jagged technological frontier: Field experimental evidence of the effects of artificial intelligence on knowledge worker productivity and quality. *Organization Science* 37(2). https://doi.org/10.1287/orsc.2025.21838.

Dellarocas C (2003) The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science* 49(10):1407–1424.

Dranove D, Jin GZ (2010) Quality disclosure and certification: Theory and practice. *Journal of Economic Literature* 48(4):935–963.

Eghbal N (2020) *Working in Public: The Making and Maintenance of Open Source Software* (Stripe Press, San Francisco).

Fang C, Kim JH, Milliken FJ (2014) When bad news is sugarcoated: Information distortion, organizational search and the behavioral theory of the firm. *Strategic Management Journal* 35(8):1186–1201.

Feldman MS, March JG (1981) Information in organizations as signal and symbol. *Administrative Science Quarterly* 26(2):171–186.

Galbraith JR (1973) *Designing Complex Organizations* (Addison-Wesley, Reading, MA).

Galbraith JR (1974) Organization design: An information processing view. *Interfaces* 4(3):28–36.

Galbraith JR (1977) *Organization Design* (Addison-Wesley, Reading, MA).

Galdin A, Silbert J (2025) Making talk cheap: Generative AI and labor market signaling. Working paper, arXiv:2511.08785.

Gans JS (2026) A model of artificial jagged intelligence. Working paper, arXiv:2601.07573.

Garicano L (2000) Hierarchies and the organization of knowledge in production. *Journal of Political Economy* 108(5):874–904.

Gavetti G (2005) Cognition and hierarchy: Rethinking the microfoundations of capabilities' development. *Organization Science* 16(6):599–617.

Goldfarb A, Tucker C (2019) Digital economics. *Journal of Economic Literature* 57(1):3–43.

Gulati R, Lawrence PR, Puranam P (2005) Adaptation in vertical relationships: Beyond incentive conflict. *Strategic Management Journal* 26(5):415–440.

Hansen MT (1999) The search-transfer problem: The role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly* 44(1):82–111.

Hayek FA (1945) The use of knowledge in society. *American Economic Review* 35(4):519–530.

Herndon T, Ash M, Pollin R (2014) Does high public debt consistently stifle economic growth? A critique of Reinhart and Rogoff. *Cambridge Journal of Economics* 38(2):257–279.

Hsiao TK, Schneider J (2021) Continued use of retracted papers: Temporal trends in citations and (lack of) awareness of retractions shown in citation contexts in biomedicine. *Quantitative Science Studies* 2(4):1144–1169.

[Iyer et al. 2025: full reference missing; please add]

Ji Z, Lee N, Frieske R, Yu T, Su D, Xu Y, Ishii E, Bang YJ, Madotto A, Fung P (2023) Survey of hallucination in natural language generation. *ACM Computing Surveys* 55(12):Article 248, 1–38.

Joseph J, Gaba V (2020) Organizational structure, information processing, and decision-making: A retrospective and road map for research. *Academy of Management Annals* 14(1):267–302.

Kellogg KC, Valentine MA, Christin A (2020) Algorithms at work: The new contested terrain of control. *Academy of Management Annals* 14(1):366–410.

Keum DD, See KE (2017) The influence of hierarchy on idea generation and selection in the innovation process. *Organization Science* 28(4):653–669.

Knudsen T, Levinthal DA (2007) Two faces of search: Alternative generation and alternative evaluation. *Organization Science* 18(1):39–54.

Kogut B, Zander U (1992) Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science* 3(3):383–397.

Landau M (1969) Redundancy, rationality, and the problem of duplication and overlap. *Public Administration Review* 29(4):346–358.

Lave J, Wenger E (1991) *Situated Learning: Legitimate Peripheral Participation* (Cambridge University Press, Cambridge, UK).

Lerner J, Tirole J (2002) Some simple economics of open source. *Journal of Industrial Economics* 50(2):197–234.

Liang W, Yuksekgonul M, Mao Y, Wu E, Zou J (2023) GPT detectors are biased against non-native English writers. *Patterns* 4(7):100779.

Lorenz J, Rauhut H, Schweitzer F, Helbing D (2011) How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences* 108(22):9020–9025.

Luca M, Zervas G (2016) Fake it till you make it: Reputation, competition, and Yelp review fraud. *Management Science* 62(12):3412–3427.

Malone TW, Yates J, Benjamin RI (1987) Electronic markets and electronic hierarchies. *Communications of the ACM* 30(6):484–497.

March JG, Simon HA (1958) *Organizations* (Wiley, New York).

Mata v. Avianca, Inc., 678 F. Supp. 3d 443 (S.D.N.Y. 2023).

McNamara G, Moon H, Bromiley P (2002) Banking on commitment: Intended and unintended consequences of an organization's attempt to attenuate escalation of commitment. *Academy of Management Journal* 45(2):443–452.

Meyer JW, Rowan B (1977) Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology* 83(2):340–363.

Nelson P (1970) Information and consumer behavior. *Journal of Political Economy* 78(2):311–329.

Nonaka I (1994) A dynamic theory of organizational knowledge creation. *Organization Science* 5(1):14–37.

Noy S, Zhang W (2023) Experimental evidence on the productivity effects of generative artificial intelligence. *Science* 381(6654):187–192.

O'Reilly CA III (1980) Individuals and information overload in organizations: Is more necessarily better? *Academy of Management Journal* 23(4):684–696.

Pallais A (2014) Inefficient hiring in entry-level labor markets. *American Economic Review* 104(11):3565–3599.

Peng S, Kalliamvakou E, Cihon P, Demirer M (2023) The impact of AI on developer productivity: Evidence from GitHub Copilot. Working paper, arXiv:2302.06590.

Piezunka H, Dahlander L (2015) Distant search, narrow attention: How crowding alters organizations' filtering of suggestions in crowdsourcing. *Academy of Management Journal* 58(3):856–880.

Polanyi M (1966) *The Tacit Dimension* (Doubleday, Garden City, NY).

Power M (1997) *The Audit Society: Rituals of Verification* (Oxford University Press, Oxford, UK).

Puranam P (2021) Human–AI collaborative decision-making as an organization design problem. *Journal of Organization Design* 10(2):75–80.

Puranam P, Alexy O, Reitzig M (2014) What's "new" about new forms of organizing? *Academy of Management Review* 39(2):162–180.

Puranam P, Raveendran M, Knudsen T (2012) Organization design: The epistemic interdependence perspective. *Academy of Management Review* 37(3):419–440.

Raisch S, Krakowski S (2021) Artificial intelligence and management: The automation–augmentation paradox. *Academy of Management Review* 46(1):192–210.

Randazzo S, Joshi A, Kellogg KC, Lifshitz H, Dell'Acqua F, Lakhani KR (2025) GenAI as a power persuader: How professionals get persuasion bombed when they attempt to validate LLMs. Working Paper 26-021, Harvard Business School, Boston.

Raymond ES (1999) *The Cathedral and the Bazaar: Musings on Linux and Open Source by an Accidental Revolutionary* (O'Reilly, Sebastopol, CA).

Reitzig M, Maciejovsky B (2015) Corporate hierarchy and vertical information flow inside the firm: A behavioral view. *Strategic Management Journal* 36(13):1979–1999.

Rivkin JW, Siggelkow N (2003) Balancing search and stability: Interdependencies among elements of organizational design. *Management Science* 49(3):290–311.

Sagan SD (2004) The problem of redundancy problem: Why more nuclear security forces may produce less nuclear security. *Risk Analysis* 24(4):935–946.

Sah RK, Stiglitz JE (1986) The architecture of economic systems: Hierarchies and polyarchies. *American Economic Review* 76(4):716–727.

Sah RK, Stiglitz JE (1988) Committees, hierarchies and polyarchies. *Economic Journal* 98(391):451–470.

Sah RK, Stiglitz JE (1991) The quality of managers in centralized versus decentralized organizations. *Quarterly Journal of Economics* 106(1):289–295.

Sanchez R, Mahoney JT (1996) Modularity, flexibility, and knowledge management in product and organization design. *Strategic Management Journal* 17(S2):63–76.

Schelling TC (1960) *The Strategy of Conflict* (Harvard University Press, Cambridge, MA).

Sharma M, Tong M, Korbak T, Duvenaud D, Askell A, Bowman SR, Cheng N, et al. (2024) Towards understanding sycophancy in language models. *Proceedings of the Twelfth International Conference on Learning Representations (ICLR)*.

Shrestha YR, Ben-Menahem SM, von Krogh G (2019) Organizational decision-making structures in the age of artificial intelligence. *California Management Review* 61(4):66–83.

Shumailov I, Shumaylov Z, Zhao Y, Papernot N, Anderson R, Gal Y (2024) AI models collapse when trained on recursively generated data. *Nature* 631(8022):755–759.

Simon HA (1947) *Administrative Behavior* (Macmillan, New York).

Simon HA (1955) A behavioral model of rational choice. *Quarterly Journal of Economics* 69(1):99–118.

Simon HA (1956) Rational choice and the structure of the environment. *Psychological Review* 63(2):129–138.

Simon HA (1971) Designing organizations for an information-rich world. Greenberger M, ed. *Computers, Communications, and the Public Interest* (Johns Hopkins Press, Baltimore), 37–72.

Spence M (1973) Job market signaling. *Quarterly Journal of Economics* 87(3):355–374.

Srikanth K, Puranam P (2011) Integrating distributed work: Comparing the effectiveness of coordination mechanisms. *Strategic Management Journal* 32(8):849–875.

Tadelis S (2016) Reputation and feedback systems in online platform markets. *Annual Review of Economics* 8:321–340.

Tushman ML, Nadler DA (1978) Information processing as an integrating concept in organizational design. *Academy of Management Review* 3(3):613–624.

Van Knippenberg D, Dahlander L, Haas MR, George G (2015) Information, attention, and decision making. *Academy of Management Journal* 58(3):649–657.

von Hippel E, von Krogh G (2003) Open source software and the "private-collective" innovation model: Issues for organization science. *Organization Science* 14(2):209–223.

von Krogh G, Spaeth S, Lakhani KR (2003) Community, joining, and specialization in open source software innovation: A case study. *Research Policy* 32(7):1217–1241.

Wallack T (2026) Nine months, 200 academic papers. One professor's AI-assisted output. *Washington Post* (September 23), https://www.washingtonpost.com/education/2026/09/23/nine-months-200-academic-papers-one-professors-ai-assisted-output/.Weick KE (1979) *The Social Psychology of Organizing*, 2nd ed. (Addison-Wesley, Reading, MA).

Weick KE (1995) *Sensemaking in Organizations* (Sage, Thousand Oaks, CA).

Wenger E (1998) *Communities of Practice: Learning, Meaning, and Identity* (Cambridge University Press, Cambridge, UK).

Williamson OE (1985) *The Economic Institutions of Capitalism* (Free Press, New York).

Yang L, Holtz D, Jaffe S, Suri S, Sinha S, Weston J, Joyce C, et al. (2022) The effects of remote work on collaboration among information workers. *Nature Human Behaviour* 6(1):43–54.

Yeverechyahu D, Mayya R, Oestreicher-Singer G (2025) The impact of large language models on open-source innovation: Evidence from GitHub Copilot. Working paper, arXiv:2409.08379.