A COMPANION PUBLICATION

An Addendum to the Five Readings

by Dr. Barbara van Veen

Dr. Barbara van Veen

INFORMATION GOVERNANCE · BEHAVIOURAL SCIENCE · EPISTEMOLOGY

CONTENTS

Questions at the boundary of governance

A companion reflection on responsibility, attention, knowledge, technological change and citizenship.

Why invite this perspective?

Guido van Nispen Amsterdam, August 2026

FOREWORD · GUIDO VAN NISPEN

The Five Readings began with a practical and normative concern: how can an information environment preserve the conditions in which citizens form independent judgements? The Information Wheel responds by identifying functions that turn claims into information on which people can responsibly act.

Dr. Barbara van Veen approached the Wheel from a different intellectual direction. Her work concerns emerging technologies, strategic decision-making, weak signals and the assumptions through which people interpret an uncertain future. Her feedback did not reject the need for Information Governance. It asked where governance should place its centre of gravity: in the landscape, in technology, or in the people who produce, select and interpret information.

Why invite this perspective? · CONTINUED

That question matters because Barbara does not treat the citizen as the passive destination of a well-designed information supply. She treats the citizen as an active information processor whose attention, experience, knowledge and unconscious filters shape what becomes visible in the first place. Her challenge is therefore complementary but not merely supportive. It asks the original model to clarify its boundaries.

The essay that follows retains the openness of that challenge. It does not pretend that every question has already been resolved. It gives Barbara's observations, research and disciplinary language room to stand on their own, including where they create productive tension with the Five Readings.

A different starting point

INTRODUCTION

I read the Five Readings with pleasure because they ask a necessary question about the governance of information. I agree with the normative choice at the centre of the model: government, journalism and artificial intelligence are not the final purpose of the information landscape. The purpose is to sustain people's ability to form an independent judgement.

At the same time, I found myself returning to a different starting point. What, or whom, should governance govern? The landscape? The technology? Or the human beings who develop, open, use, select and interpret information? These are not interchangeable units of analysis. The answer affects what we understand the problem to be and what remedies we consider plausible.

My questions arise from strategy and foresight. In those fields, information does not enter decision-making as neutral material. It is noticed through existing knowledge, professional roles, expectations and mental models. Before evidence can be evaluated, it must first survive attention. Before an emerging development can influence strategy, someone must categorise it as relevant. Those prior acts of selection are easy to overlook precisely because many of them are not conscious.

A different starting point · CONTINUED

The Information Wheel describes important functions within an information ecosystem. My concern is what happens at its human boundary: where information is selected, where representations acquire meaning, where provisional knowledge becomes a basis for action, and where technologies that are still developing are too quickly treated as stable.

I offer five questions rather than five answers. Some lead towards a refinement of the original framework. Others may remain alongside it. A useful governance model should be capable of living with such open questions.

Dr. Barbara van Veen

Who—or what—should governance govern?

The unit of analysis determines where responsibility is placed

QUESTION 01

In the preparatory material, generative AI sometimes appears to have an independent status as an accelerator of journalistic problems. I hesitate at that formulation. AI does not decide to develop, release or use itself. People and organisations make those choices. The problems attributed to AI are therefore also problems we create—and problems whose management ultimately requires changes in human and institutional behaviour.

This distinction is more than linguistic. When we say that an algorithm decides, a platform prefers or a model invents, responsibility can disappear into the system. Governance should travel in the opposite direction. Who commissioned the system? Who selected its objective? Who accepted its limitations? Who benefits from its operation, and who can intervene when its effects become harmful?

Generative AI also seems to me one stage in a longer development that began before today's models. News gathering and consumption increasingly run through social media, video platforms, aggregators and other intermediaries. The Reuters Institute's Digital News Reports document this changing source environment and the growing importance of third-party platforms.[1] AI may intensify the pattern, but it did not create the underlying shift.

Who—or what—should governance govern? · CONTINUED

This is why I begin with human agency. If citizens are to move from passive to active information processing, governance must consider the behaviour of developers, publishers, institutions and users—not only the formal properties of a technological artefact.

Yet a question remains. Even if technology is not a moral actor, can a consequential infrastructure still be a legitimate object of governance? Financial and communications infrastructures are regulated without being treated as persons. Perhaps AI belongs in that category. If so, the framework should say clearly that the technology is governed through duties imposed on the people and organisations that design, own and deploy it.

Where does information become meaning?

What people do not read may be as revealing as what they consume

QUESTION 02

I once asked approximately 200 executive-education students to keep a record for one week of what they did not read and then to reflect on their omissions. They were executives from many industries, generally over 35, studying strategy. Only through the exercise did many realise how much information they excluded before consciously considering its quality.

They passed over material because they lacked time and other tasks competed for attention. They did not see its relevance, often because relevance was defined by their function, industry or personal preferences. Some information was cognitively expensive: too difficult, too negative or too demanding. Other material felt too familiar to deserve attention. Their selective filters were therefore partly unconscious and acted before the news had truly entered awareness.

This suggests that the problem described by Information Governance is not only the quality of available information. It is also the early categorisation through which people decide what counts as information for them. A reliable report that never survives the relevance filter cannot contribute to judgement. A plural information landscape does not automatically produce a plural information diet.

Where does information become meaning? · CONTINUED

The same students tended to consult a limited number of sources. Their source mix was rarely the product of explicit reflection or triangulation. It was shaped by what was usual in their industry, company or professional role. They did not necessarily ask whether those sources were capable of saying something meaningful about future developments—the question that matters especially in strategy.

This is not simply a deficit of individual diligence. Attention is finite, and selection is unavoidable. But governance that focuses exclusively on information quality misses the mechanisms determining which information reaches a person at all.

Information becomes meaning only through a relation between representation, interpreter and context. The same material can appear relevant, threatening, trivial or implausible depending on the knowledge and expectations brought to it. That behavioural layer belongs next to the Wheel even if it does not belong inside the same diagram.

QUESTION 02 · CONTINUED

My doctoral research into the perception and interpretation of weak signals showed that diversity matters when people attempt to recognise emerging developments such as new technologies.[2] Here diversity meant variety in kinds of knowledge and experience. What people have previously learned and encountered affects what they can perceive, how they categorise it and which interpretations become available.

This matters for the Information Wheel because diversity is not exclusively a feature of the information landscape. A system can contain many outlets and perspectives while the people producing or processing information share similar knowledge, assumptions and blind spots. Conversely, a cognitively diverse group may discover different implications in the same body of evidence.

A narrow source mix and a narrow experiential mix can reinforce each other. Familiar sources confirm familiar categories; familiar categories make unfamiliar signals easier to dismiss. Triangulation is therefore not only the mechanical consultation of several sources. It is an attempt to expose an interpretation to perspectives capable of seeing different things in the same material.

Where does information become meaning? · CONTINUED

Should governance respond by trying to manage citizens' attention directly? That would create obvious paternalistic risks. A more defensible aim may be to preserve conditions in which people can discover why something has been selected for them, encounter genuinely independent perspectives and recognise the limits of their habitual source mix.

I would therefore acknowledge the behavioural layer without simply inserting it into the Information Wheel. The Wheel maps functions in the information ecosystem. Behavioural and strategic inquiry ask how a situated person engages with what that ecosystem produces. The distinction protects the scope of the model while preventing the citizen from disappearing into an endpoint labelled "judgement".

Can knowledge ever be first-hand?

From unmediated truth to accountable mediated perspectives

QUESTION 03

I am not convinced that “new, first-hand knowledge” exists in the strong sense the phrase can imply. All sources are already engagements with reality. Observation is selective; testimony is framed; data reflect decisions about what to count; journalism transforms experience into communicable form. This does not mean that all accounts are equally good. It means that quality must be located in the strength and transparency of mediation.

In my field, I would speak of mediated perspectives. One useful way of expressing the underlying problem comes from Charles Sanders Peirce’s triadic account of the sign: the object, the representamen through which it is perceived, and the interpretant—the meaning produced through interpretation.[3] We do not simply transfer the object into another mind. We encounter a representation and make meaning from it.

Consider a municipal council meeting. A reporter may be physically present, but what is observed includes the perspectives and performances of councillors. The reporter then selects exchanges, supplies context and interprets their significance. Without selection and interpretation there may be a recording, but not yet journalism. Presence offers access to material; it does not abolish mediation.

Can knowledge ever be first-hand? · CONTINUED

The question is therefore not whether journalism can escape perspective. It is whether the chain of mediated perspectives through which a claim is produced is sufficiently strong, diverse and open to criticism. Helen Longino's work on knowledge as a social and critical achievement is relevant here: objectivity is not secured by pretending that observers have no standpoint, but through conditions in which assumptions and evidential relations can be exposed to meaningful criticism.[4]

If I understand the intention of the Five Readings correctly, "new, first-hand knowledge" may be intended to identify the public production of evidence about a phenomenon. That function remains important. Otherwise interpretation risks circulating without renewed contact with documents, testimony, observation and data.

QUESTION 03 · CONTINUED

I would nevertheless keep the formulation open. Is the distinctive journalistic contribution new and first-hand material? Or is it the disciplined effort to make an interpretation transparent and plausible through good evidence? These alternatives overlap, but they are not identical.

Can knowledge ever be first-hand? · CONTINUED

"Original reporting" may be a more serviceable institutional category than "first-hand knowledge," provided originality does not become a claim to epistemic purity. Its value lies in extending the publicly available evidentiary network: making documents accessible, testing testimony, recording observations, revealing methods and identifying uncertainty.

The strength of such a network cannot be judged only by the number of links. Ten sources repeating the same account do not provide the same epistemic diversity as several independent perspectives with different access, knowledge and incentives. Quality depends on provenance, independence, critical interaction and the possibility of correction.

This position does not entail relativism. Mediation does not mean that reality places no constraints on interpretation. It means that claims become more trustworthy when the route between object, representation and interpretation is visible enough to be examined. The alternative to naïve objectivity is not arbitrary opinion; it is accountable knowledge-making.

Can knowledge ever be first-hand? · CONTINUED

Perhaps the language of the Wheel should therefore remain deliberately dual. Journalism produces new public evidence, while acknowledging that evidence becomes knowledge only through mediated and contestable interpretation. I leave open whether this is best expressed by refining "first-hand knowledge" or replacing the phrase altogether.

OPEN QUESTION
Does "first-hand" clarify journalism's originating function, or does it obscure the mediated character of every publicly meaningful account?

How should governance deal with technologies that are still evolving?

Is this an incident, a phase, or a stable structural development?

QUESTION 04

I do not think generative AI can yet be regarded as a mature component of the information landscape. Historically, we remain in an experimental phase. The technology exists, but institutions, norms, quality mechanisms and social practices are still developing—much as they did during the early periods of print, radio and the internet.

This has at least two consequences. First, we do not yet know the stable role AI will occupy. Will it primarily become a search system, editor, personal assistant, journalistic intermediary or infrastructure on which other systems depend? At present it performs aspects of all these roles. Governance formulated around one interface may miss the function in which power eventually settles.

Second, many behaviours currently attributed to AI may belong to the present generation of models rather than to the technology as such. Hallucination, weak source transparency, excessive confidence and dependence on existing material may change substantially. Rules that imagine the future as a continuation of the present can become legacy problems before they are fully implemented.

This does not mean that governance should wait. Experimental technologies can produce real consequences, and early institutional choices can become difficult to reverse. But governance should distinguish temporary technical manifestations from durable relationships of power and responsibility.

How should governance deal with technologies that are still evolving? · CONTINUED

Principles such as accountability, contestability, provenance, pluralism and redress may remain relevant across changing implementations. Technical standards will need faster revision. The model may therefore require several tempos: stable normative commitments, revisable institutional duties and highly adaptive technical requirements.

DESIGN CHALLENGE
How can governance act under present uncertainty without turning today's model limitations into tomorrow's permanent regulatory categories?

QUESTION 04 · CONTINUED

Strategy under uncertainty is not the prediction of a single future. It is the disciplined examination of assumptions, alternatives and signals that may indicate a changing environment. Governance should be capable of the same discipline. It needs mechanisms for review, independent evaluation, incident learning and revision when the role of the technology changes.

The distinction between an incident and an ongoing development is important. A dramatic failure may reveal a structural weakness, or it may reflect a transitional implementation. Conversely, a quiet change in source access, recommendation or infrastructural dependence may appear temporary while creating a durable shift in power. Governance needs criteria for learning which is which.

How should governance deal with technologies that are still evolving? · CONTINUED

One way to reduce legacy problems is to formulate obligations around what institutions do rather than around what a particular product is called. An organisation that deploys a system to shape consequential visibility should be able to explain its purpose, monitor effects and provide meaningful contest—whether the system is described as an algorithm, model, agent or editorial assistant.

At the same time, principle-based governance can become empty if principles are not translated into observable responsibilities. Adaptability should not mean vagueness. It should mean that the purpose of an obligation remains stable while its technical implementation can change.

The Five Readings are strongest when they are understood in this way: not as a permanent classification of today's information technologies, but as a normative account of the capacities a trustworthy information environment should preserve. Whether the present Wheel achieves that balance is a question worth continuing to test.

What does an informed citizen actually do?

The citizen as an active information processor

QUESTION 05

The strength of the Information Governance model lies in placing the citizen at its normative centre. Yet the citizen can still appear mainly as the recipient of information produced and organised elsewhere. My reading suggests that an important link remains underdeveloped. Between information and judgement stands a human being who selects, categorises, interprets and gives meaning.

A high-quality information landscape is therefore a necessary condition for independent judgement, but not a sufficient one. Citizens can inhabit a diverse and accountable environment while consulting a narrow source mix, avoiding cognitively expensive material or failing to recognise the relevance of an unfamiliar development.

The informed citizen does not possess all relevant information. That ideal is impossible. An informed citizen is better understood as someone capable of reflecting on the limits of a source mix, asking what has been excluded, comparing perspectives, recognising uncertainty and revising an interpretation when stronger evidence or a different frame becomes available.

What does an informed citizen actually do? · CONTINUED

This active role includes restraint. Independent judgement may require delaying a conclusion, distinguishing what is observed from what is inferred, or recognising that several interpretations remain plausible. It may also require action: seeking another source, asking who can know, testing an assumption, participating in institutional processes or changing a decision.

Governance cannot guarantee these practices without undermining the independence it seeks to protect. It can, however, reduce obstacles. It can make provenance visible, sustain spaces for criticism, prevent manipulation of attention, support information literacy and ensure that consequential decisions can be contested.

The citizen should neither be romanticised as perfectly autonomous nor reduced to a passive victim of systems. Human behaviour is shaped by environments, but people retain agency within them. The governance challenge is to design conditions that make reflective agency more possible.

QUESTION 05 · CONTINUED

What does an informed citizen actually do? · CONTINUED

This suggests that the informed citizen is not the final box in a linear process. Information and judgement interact recursively. A judgement changes what a person subsequently notices; new information can expose the categories through which an earlier judgement was formed. Discussion with others may reveal knowledge that no individual source supplied.

The behavioural layer is therefore dynamic. Selection affects exposure, exposure affects interpretation, interpretation affects future selection. Professional roles and institutional cultures shape this cycle, as do interfaces and recommendation systems. The quality of the public environment and the capabilities of citizens continually influence each other.

Should that cycle be incorporated into the Information Wheel? I remain cautious. Adding every relevant process can make a model less useful. It may be better to place two models in dialogue: one describing the public functions through which trustworthy information is produced and circulated, and another describing the human processes through which information becomes meaningful and actionable.

What does an informed citizen actually do? · CONTINUED

What matters is that the boundary between them is not mistaken for an explanation. The Wheel cannot assume that trustworthy information automatically becomes informed judgement. Behavioural science cannot assume that individuals can compensate for poor institutions through vigilance alone. Strategy cannot assume that more data dissolves uncertainty.

The citizen-centred ambition of Information Governance is strengthened when these limits are made explicit. Its task is not to produce correct beliefs. It is to protect the conditions under which people can form, test and revise judgements of their own.

CAPABILITY QUESTION
Can citizens discover relevant information, inspect its provenance, encounter genuinely different knowledge and revise a judgement without disproportionate cognitive or institutional cost?

Keep the questions open

Dr. Barbara van Veen

CONCLUDING REFLECTION

My response does not lead me to reject the Information Governance model. Its normative case remains strong. An abundant and technologically mediated information environment requires public arrangements that preserve evidence, plural perspectives, accountability and the possibility of independent judgement.

The questions concern focus. Human beings and institutions remain responsible for technologies they create and deploy. Citizens filter information before consciously assessing it. Knowledge is mediated through representation and interpretation. Diversity concerns the people who produce and process information as well as the formal variety of the landscape. Generative AI is consequential, but its stable role remains uncertain.

These propositions complicate governance in productive ways. They advise against treating AI as an autonomous cause, against equating availability with attention, against confusing proximity with unmediated knowledge, and against regulating temporary model behaviour as if it were a permanent feature of the future.

Keep the questions open · CONTINUED

They also caution against making the Information Wheel carry every adjacent theory. Behavioural science, semiotics, social epistemology and strategy reveal processes at its boundaries. Their contribution may be strongest when those boundaries remain visible and the models are allowed to question one another.

The result is not a closed synthesis. Several issues should remain open: whether AI requires a distinct place in the Wheel; whether "first-hand knowledge" can be retained without epistemic overclaim; how governance can distinguish transient incidents from structural developments; and how far public institutions should intervene in the conditions of attention.

A framework is not weakened by such questions. It becomes more intellectually responsible. Information Governance is strengthened through dialogue with behavioural science, strategy and epistemology—not because those disciplines confirm every choice, but because they help us see what the model illuminates, what it presupposes and what remains to be learned.

Sources and points of departure

ENDNOTES

1. Nic Newman et al., Digital News Report 2025, Reuters Institute for the Study of Journalism, 2025. The report documents declining engagement with traditional news sources and growing dependence on social media, video platforms and aggregators.

2. Barbara L. van Veen, Straws That Tell the Wind: Top-Manager Perception of Distant Signals of the Future, doctoral dissertation, Delft University of Technology, 2020, DOI: 10.4233/uuid:d26842a4-8f72-44d4-8952-cd17988d18d8; and B. L. van Veen, J. R. Ortt and P. G. Badke-Schaub, “Compensating for perceptual filters in weak signal assessments,” Futures 108 (2019), 1–11.

3. Charles Sanders Peirce’s triadic account distinguishes the object, representamen and interpretant. The accessible point of departure supplied with the original feedback was “Charles Peirce’s Triadic Model of Communication,” Media Studies.

4. Helen E. Longino, Science as Social Knowledge: Values and Objectivity in Scientific Inquiry, Princeton University Press, 1990; and The Fate of Knowledge, Princeton University Press, 2002. See also the lecture “Perspectives and Pluralities,” referenced in the original feedback.

Sources and points of departure · CONTINUED

5. The approximately 200-person executive-education exercise, including the categories of time, relevance, cognitive cost, negativity and familiarity, is described in Dr. van Veen's written feedback on the Five Readings dated 28 July 2026.

6. Barbara van Veen's research profile concerns the anticipation of emerging technologies, weak signals, strategic decision-making and strategy under uncertainty. See the TU Delft Research Portal and futuristbarbara.com.

7. The distinction in this essay between stable principles, revisable institutional duties and adaptive technical standards develops the reasoning of the original feedback; it should not be treated as a direct quotation from that feedback.

8. URLs consulted for this publication:
reutersinstitute.politics.ox.ac.uk/digital-news-report/2025;
research.tudelft.nl/en/persons/bl-van-veen-2;
futuristbarbara.com/barbara;
media-studies.com/triadic-model-semiotics;
youtube.com/watch?v=631gObE7ctA.

ABOUT THE AUTHOR

Dr. Barbara van Veen

Barbara van Veen is an independent researcher and futurist working at the intersection of emerging technologies, strategic decision-making and systems thinking. Her particular interest is how assumptions, perceptual filters and anticipatory practices influence what decision-makers notice and how they act under uncertainty.

She earned her doctorate at Delft University of Technology with the dissertation Straws That Tell the Wind: Top-Manager Perception of Distant Signals of the Future. Her research addresses weak signals, emerging technologies, mental models and the role of cognitive diversity in strategic interpretation. She has published in Futures and Futures & Foresight Science.

Barbara teaches and works with decision-makers confronting technologies whose long-term effects cannot yet be known. She is certified in the Harvard case method and uses real-world cases to examine how organisations can develop strategic responses without simplifying uncertainty prematurely.

Her current work connects technological emergence with the human practices through which futures are anticipated, interpreted and shaped over time.

https://www.futuristbarbara.com

POLDERS.MEDIA · AMSTERDAM · 2026