

The Information Wheel as a Living System

An Addendum by Anna Pereira to the Five Readings

Anna Pereira

Information Governance · Scientific Perspectives

This is preliminary and ongoing work. [Email me](#) for comments or conversation.

A scientific lens, not a substitute

This addendum asks what becomes visible when the Information Wheel is read as a living system—adaptive, contested and inseparable from the citizens who inhabit it.

Introduction	3
A Brief Scientific Foundation	4
A Scientific Foundation for the Information Wheel	5
The citizen as an active participant	8
Entropy as one supporting lens	10
Concluding reflection	12
References	13

Introduction

One of the things I find most compelling about Five Readings of the Information Wheel is that it brings together two ways of understanding the world that are often kept apart.

Human beings reason through emotion and felt experience, but also through observation, comparison, and scientific analysis. When these forms of understanding are allowed to complement one another, we may become better able to see not only what is happening, but also what might help us respond.

Guido's Five Readings begin with a practical concern: democratic societies do not necessarily suffer from too little information. They suffer when the conditions under which information is produced, tested, preserved, encountered, and connected become difficult to see and difficult to contest.

The Information Wheel offers a clear response. Original reporting, verification, context, diverse perspectives, and synthesis describe five functions through which information can become useful for independent judgement. The framework then considers these functions as ingredients of an information environment, capabilities needed by citizens, responsibilities distributed across institutions, and conditions that can be observed over time.

A scientific perspective does not replace this civic and institutional argument. It offers another way to ask why these functions belong together, how they interact, and what happens when one becomes weak or dominant. At a broad level, this addendum applies the same underlying inferential logic that Karl Friston and colleagues use to understand how biological systems form and update models of their environments to a different domain: how people and institutions exchange evidence, coordinate, and revise their understanding.

I am not suggesting that democratic judgement can be reduced to a scientific equation, or that one scientific framework can determine the right political outcome. Rather, the literature may help make some of the dynamics surrounding judgement more visible.

Although the discussion begins with democratic institutions because that is the setting of the Five Readings, the underlying science is not tied to one political model or intended to prescribe a particular form of governance. It may be useful to any institution or governance structure seeking to better understand, coordinate with, and serve the people and communities affected by its decisions. The relevant capacities are general: observing change, preserving meaningful difference, forming shared understanding, coordinating action, correcting error, and supporting human agency and wellbeing.

A Brief Scientific Foundation

People do not encounter the world as a complete, objective picture. We interpret what we experience through our bodies, memories, expectations, and mental models. When new information conflicts with those models, we have an opportunity to revise our understanding. Sometimes the conflict reveals an error; sometimes it shows us something genuinely new.

Active inference offers one formal account of this process. In simple terms, people and living systems are continually trying to understand what is happening well enough to remain oriented and able to act (Friston, 2010; Friston et al., 2017). The free-energy principle is not simply a drive to eliminate all uncertainty or disorder. It describes, in part, how a system balances having a model that is accurate enough to guide behaviour with one that is simple enough to remain usable.

Uncertainty is therefore not only an obstacle. It can show us where our understanding is weak and where more information might be useful. Active inference describes action not only as changing the world, but also as seeking information that can resolve important uncertainty (Parr & Friston, 2017; Parr & Friston, 2019). Attention matters as well: not every signal receives equal weight. Information environments influence what people notice, what they trust, and what they consider worth exploring.

People also update their models together. Communication allows us to share observations and develop enough common ground to coordinate, without requiring everyone to hold identical beliefs (Friston et al., 2020). Research on collective intelligence similarly suggests that group performance depends not only on individual ability, but also on how people listen, participate, and respond to one another (Woolley et al., 2010).

This creates an important balance. Too little variation can hide error and make a shared model brittle. Too little coherence can leave people with different perspectives but no meaningful way to understand or act together. A healthy information environment therefore needs to preserve independent observation and relevant difference while helping those differences become intelligible and usable.

Artificial intelligence increasingly participates in this process (Friston et al., 2024). It can accelerate the production, distribution, comparison, and synthesis of information, but it can also influence which signals receive attention and how confidently they are presented. The important question is not simply whether AI provides more information. It is whether the larger environment remains capable of observing, understanding, coordinating, and revising when its model of the world is wrong.

A Scientific Foundation for the Information Wheel

Rather than introducing a separate framework, this scientific lens helps illuminate how the five elements Guido identifies in the Information Wheel work together to support the way people and institutions form, share, and revise collective understanding.

Original reporting

Original reporting introduces new observations into the shared environment. Someone attends the meeting, reads the court file, interviews a witness, examines the data, visits the affected community, or asks a question that would otherwise remain unasked. From a scientific perspective, this is a form of active sampling. A system cannot update its model using information that it never encounters.

If original observation weakens, the information environment may continue to produce large quantities of material while becoming increasingly dependent on inherited assumptions and repeated accounts. Generated volume can circulate an existing model more efficiently, but struggles to reliably introduce what that model has failed to notice.

The first function of the Wheel is therefore not simply content production. It is maintaining routes through which something genuinely new can enter public understanding.

Verification

Verification tests whether an observation or claim deserves confidence. It is not the absence of error, but the presence of practices that can detect, disclose, and repair it.

This also connects to calibrated belief updating. Fluency, familiarity, emotional force, and repetition can all make a claim feel more credible, though none of these signals is necessarily sufficient on its own. More calibrated confidence may be supported by considering the quality of the evidence, the reliability of the method, the source's basis for knowing, and the availability of correction.

This does not mean that citizens are expected to verify everything personally. We all depend on other people and institutions to investigate, preserve, and check information. The important question is whether the surrounding environment makes the basis for that reliance sufficiently visible.

Taken together, this points less toward maximum trust or permanent suspicion and more toward calibrated trust: the ability to assign different levels of confidence to different claims for reasons that can be explained.

Context

Context helps transform an observation into an understanding of what it means. It connects events to time, history, institutions, causes, consequences, and affected groups. In the language of generative models, context helps identify the larger structure within which an observation becomes intelligible.

A fact can be accurate while still producing a misleading picture if the relationships surrounding it are omitted. A short-term change may look different when viewed against a longer trend. A court ruling may appear different when connected to the legal history that produced it. An isolated complaint may acquire a different meaning when connected to the institutional pattern in which it occurs.

Context therefore gives information a form that can be used for inference rather than merely remembered as a collection of fragments.

Diverse perspectives

Diverse perspectives allow a system to encounter observations and interpretations that no single participant possesses. They enlarge the range of hypotheses through which an event can be understood and make it more difficult for one partial model to become invisible as the whole.

This does not mean that every claim deserves equal weight, or that disagreement is automatically informative. A perspective is valuable when it reveals relevant experience, evidence, consequences, or assumptions that would otherwise remain absent.

Nor does the number of outlets necessarily measure meaningful diversity. Many publications may rely on the same original report, use similar sources, or reproduce the same underlying frame. Conversely, a smaller number of independent accounts may reveal important differences in experience and interpretation.

The relevant question is not simply how much information is present, but whether the system has preserved enough independent variation to notice when its existing model is incomplete.

Synthesis

Synthesis allows people to form a provisional whole from many partial observations. It can reveal relationships across events, sources, institutions, and perspectives that would otherwise remain disconnected. In this sense, synthesis resembles a necessary form of compression. People struggle to hold every detail in mind at once. They work best when summaries, explanations, and models make complexity usable.

ut compression can either preserve or destroy structure.

B

A good synthesis reduces unnecessary complexity while retaining the relationships, uncertainty, provenance, and disagreement needed for a person to evaluate it. A poor synthesis creates a smooth surface by concealing what was selected, omitted, or inferred.

The value of synthesis is therefore not that it produces one final account. It is that it helps people see how accounts relate while remaining able to move back toward the evidence beneath them.

Taken together, the five functions form less of a production line than a recurring process. Reporting introduces observations. Verification tests them. Context gives them structure. Perspectives expose what any single account may miss. Synthesis makes the resulting field more navigable. New observations and citizen responses then reopen the process.

The Wheel does not need to produce permanent agreement. It needs to help disagreement remain connected to evidence, intelligible to others, and open to revision.

The citizen as an active participant

The citizen is not simply the final recipient of a completed information product. Citizens interpret, compare, trust, question, ignore, remember, share, act, and provide feedback.

This means that access is only the beginning of the information pathway:

**access → exposure → comprehension → calibrated belief →
willingness to update → action → feedback**

A system can fail at any stage. A person may never encounter useful information. They may encounter it without noticing it, understand it without trusting it, trust it without knowing how to act, or act without receiving feedback about the result.

Research on information access illustrates this distinction. In a large field experiment, providing uncensored internet access did not by itself lead participants to seek politically sensitive information. Additional encouragement increased information acquisition, and those who acquired more information subsequently changed in knowledge, beliefs, attitudes, and intended behaviour (Chen & Yang, 2019).

The same information may also affect different people differently. People begin with different prior knowledge, levels of trust, cognitive bandwidth, and tolerance for uncertainty. A technical explanation may be clarifying for one person and inaccessible to another. A correction from a distrusted institution may be interpreted as evidence of further manipulation.

This does not mean that truth is relative to the audience. It means that evidence becomes usable through a pathway that depends partly on the starting state of the person encountering it.

The literature also suggests that the goal should not be generalized scepticism. A digital media-literacy intervention improved people's ability to distinguish mainstream from false news, although it reduced perceived accuracy for both categories rather than simply increasing trust in mainstream sources (Guess et al., 2020). Other work suggests that directing attention toward accuracy can reduce the likelihood that people share misinformation (Pennycook & Rand, 2022).

The design objective is therefore calibrated trust: beliefs strong enough to support action, but open enough to change.

This is why one useful principle for AI-mediated information may be: Personalise the scaffolding, not the truth.

Systems can adapt explanations, language, pace, or context to help people engage with evidence. The more difficult question is how to do this without quietly changing the evidentiary standard or narrowing what a person is able to encounter. Seeing citizens as active participants also does not mean placing responsibility for the whole information environment on them. No individual can personally report, verify, contextualise, and synthesise everything. A more supportive information environment distributes this cognitive work while leaving meaningful routes for citizens to question, challenge, and correct the institutions around them.

Entropy as one supporting lens

Entropy offers one possible scientific lens for examining these dynamics, provided it is used carefully.

In information theory, Shannon entropy describes the uncertainty or distribution of possible signals within a defined communication system (Shannon, 1948). This is different from thermodynamic entropy, and it is not a direct measure of social disorder. It cannot tell us which claim is true, which perspective matters most, or whether a society is functioning well. Used carefully, however, it offers one way to think about the relationship between uncertainty, variety, and understanding.

An information environment with little visible difference can appear orderly while becoming less able to learn. A single account repeated consistently may reduce visible variation without reducing the possibility that the account is wrong. By contrast, independent observations, relevant disagreement, and alternative explanations may temporarily increase uncertainty while helping a system notice error.

This suggests a useful distinction between productive and unproductive entropy.

Productive entropy may emerge when relevant evidence, independent observations, and genuinely different perspectives enter the system. It can unsettle premature certainty, reveal limitations in an existing model, and create opportunities for learning.

Unproductive entropy is uncertainty that does not readily become understanding. It may take the form of duplicated noise, obscured provenance, strategic contradiction, overwhelming volume, or endless circulation without correction. In these cases, the system may contain more signals without gaining a clearer picture of what those signals mean.

This distinction may also offer a more precise way to think about bias. Entropy-like measures would not remove human bias or determine which perspectives are correct. They might, however, help reveal where information sampling has become unusually narrow, where sources are highly correlated, where repetition is being mistaken for independent confirmation, or where relevant perspectives are absent.

They could also help describe whether an apparent abundance of information reflects meaningful variety or merely repeated material. A high number of sources does not necessarily represent a high level of independent information. A low level of visible variation does not necessarily represent clarity or agreement.

Used in this way, entropy could provide a distinct analytical tool alongside human judgement. It has the potential to help make certain patterns visible without determining what those patterns mean or what action follows from them. Interpretation, values, and responsibility remain human questions.

A key point then is not to reduce entropy at any cost. A system can reduce visible disorder by suppressing people, information, or difference. The more useful aim rather is to preserve enough meaningful variation for a system to detect error, while supporting enough coherence for people to understand and coordinate.

Entropy is therefore one lens among several. It has the potential to support the Information Wheel by helping us ask whether difference is being transformed into learning or merely multiplied into confusion.

Concluding reflection

The scientific literature offers one additional way to understand how the five functions of the Information Wheel can support one another.

O riginal reporting supplies new observations. Verification calibrates confidence. Context gives observations structure and meaning. Diverse perspectives preserve the variation needed to reveal what any single account misses. Synthesis supports a shared but revisable understanding.

The citizen remains at the centre—not as the output of an information system, but as an active participant in its continual updating.

Seen in this way, Information Governance is not primarily about controlling conclusions. It is about maintaining the conditions under which people and institutions can observe, compare, understand, coordinate, and revise.

A scientific perspective adds a modest but useful insight: adaptive systems need enough variation to detect error, enough coherence to coordinate, and enough feedback to change when their understanding is wrong.

References

- Chen, Y., & Yang, D. Y. (2019). The impact of media censorship: 1984 or brave new world?. *American economic review*, 109(6), 2294-2332.
- Friston, K. (2010). The free-energy principle: a unified brain theory?. *Nature reviews neuroscience*, 11(2), 127-138.
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: a process theory. *Neural computation*, 29(1), 1-49.
- Friston, K. J., Parr, T., Yufik, Y., Sajid, N., Price, C. J., & Holmes, E. (2020). Generative models, linguistic communication and active inference. *Neuroscience & Biobehavioral Reviews*, 118, 42-64.
- Friston, K. J., Ramstead, M. J., Kiefer, A. B., Tschantz, A., Buckley, C. L., Albarracin, M., ... & René, G. (2024). Designing ecosystems of intelligence from first principles. *Collective Intelligence*, 3(1), 26339137231222481.
- Guess, A. M., Lerner, M., Lyons, B., Montgomery, J. M., Nyhan, B., Reifler, J., & Sircar, N. (2020). A digital media literacy intervention increases discernment between mainstream and false news in the United States and India. *Proceedings of the National Academy of Sciences*, 117(27), 15536-15545.
- Parr, T., & Friston, K. J. (2017). Uncertainty, epistemics and active inference. *Journal of the Royal Society Interface*, 14(136), 20170376.
- Parr, T., & Friston, K. J. (2019). Generalised free energy and active inference. *Biological cybernetics*, 113(5), 495-513.
- Pennycook, G., & Rand, D. G. (2022). Accuracy prompts are a replicable and generalizable approach for reducing the spread of misinformation. *Nature communications*, 13(1), 2333.
- Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, 27(3), 379-423.
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *science*, 330(6004), 686-688.

Anna Pereira

PHD · MS · RESEARCH SCIENTIST · AUTHOR · ADVISOR · SPEAKER

Anna Pereira explores how we navigate uncertainty in a changing world. With a background in human factors, research and design, her work bridges perceptual science, embodied cognition and the physics of life.

After more than a decade leading human-centred research in Big Tech—including Google, Microsoft and Valve—she stepped back to investigate the deeper patterns behind how we sense, adapt and thrive. Guided in part by informal advisership from neuroscientist Karl Friston, her current work translates Active Inference and the Free Energy Principle into accessible insights for everyday life.

Anna holds a PhD in Human Factors from the University of California, Berkeley, completed postgraduate studies at the MIT Media Lab, and lives in Lausanne, Switzerland.

actincycle.com

anna@activeinferencecycle.com