



Histories

■ of the

Fact

in the West

FACT OR VALUE

‘**N**ow, what I want is, Facts. Teach these boys and girls nothing but Facts. Facts alone are wanted in life. Plant nothing else, and root out everything else. You can only form the minds of reasoning animals upon Facts; nothing else will ever be of any service to them. This is the principle on which I bring up my own children, and this is the principle on which I bring up these children. Stick to Facts, sir!’

...

The speaker, and the schoolmaster, and the third grown person present, all backed a little, and swept with their eyes the inclined plane of little vessels then and there arranged in order, ready to have imperial gallons of facts poured into them until they were full to the brim.

DICKENS, HARD TIMES (1854)

KORLING, BIOLOGICAL CHEMICAL LABORATORY AT THE
DOW CHEMICAL COMPANY (MICHIGAN, 1955)



Histories of the Fact

in the West

FACT OR VALUE

The texts in this volume are the edited transcripts of lectures delivered for Fact or Value.

The illustrations in this book are purely ornamental and bear no relation to the texts.

Texts © Contributors, 2026
Design © Fact or Value, 2026

ISBN 978 93 5717 743 6

This work is licensed under
CC BY-NC-ND 4.0.

You may copy and redistribute the material for non-commercial purposes, with attribution, and without modifications.

For permissions beyond the license (including commercial use), please obtain prior written permission from Fact or Value.


Some rights reserved.

FACTS are the invulnerable constants in a world full of speculation and rhetoric. They are descriptions lifted off the surface of reality. Once established, they remain true forever. Or do they?

In the sixteenth century, 'scientists' in Europe believed that garlic could weaken the power of magnets. By the eighteenth, however, they had grown so sceptical that they dismissed the existence of meteorites as absurd.

What explains this upsurge in incredulity? Were these attitudes rooted in the same source: an inability to ground beliefs in evidence?

In the two centuries that divide these perspectives, a significant shift occurred in how the fledgling scientific communities of the West approached the world. The scientific advancements of the seventeenth century prompted governments, scholars, businesses and individuals to



FLUDD, UTRISQUE COSMI . . . HISTORIA (1617)

take data seriously. It sparked a revolutionary but painstaking effort to gather facts—tables filled with seemingly insignificant information that, over time, led to more accurate understandings of reality and overturned centuries-old theories rooted in tradition and authority.

In a sense, the West ‘atomized’ reality and ‘discovered’ facts. Reports of new phenomena—of talking plants and new stars—were considered or dismissed based on how much they deviated from the established body of facts—an epistemic filtering process that pre-scientific Europe simply could not perform.

However, as this ‘culture of fact’ permeated the political foundations of European society, it also served as a tool to justify empire, slavery and prejudice.

The findings of pseudoscientific disciplines, such as phrenology, commanded the reverence and certainty of factual knowledge, and were invoked to justify invasion and oppression.

Today, we regard facts as the indivisible building blocks of knowledge, distinct from values and theories, and immune to human subjectivity, bias or prejudice. Anyone who questions an established fact is often seen as delusional or acting in bad faith, in dire need of ‘fact-checking’.

But if some ‘facts’ have later been proven untrue, what does that say about the nature of facts? Are they objective descriptions of the world? Or are they inferences—informed guesses—drawn by observing the opaque surface of reality? Are facts found, or are they made?

HOOKE, MICROGRAPHIA (1665)

Contents

A detailed black and white micrograph of a spider web, showing the intricate structure of the spiral and radial threads. A small letter 'D' is printed near the center of the web, marking a specific point of interest. The background is a solid red color.

D



1

Steven Shapin

The Soft World of Judgement

53

Daniel Rosenberg

The Age of the Keyword

95 *E*

Lisa Gitelman

Objectivity

119

Matthew Stanley

Solar Eclipses in History

175

Randolph Head

Archives and
the Information State

227

Paul Duguid

On the Social Life
of Technology

267

Richard Firth Green

A Crisis of Truth, Revisited

291

Daryn Lehoux
Sergio Sismondo

Epistemic Corruption

PHOTOGRAPHIC LABORATORY,
UNITED STATES NATIONAL MUSEUM, 1880s



Steven Shapin

The Soft World of Judgement

An Interview

How was testimony managed in the seventeenth century? If empiricism was a collective activity, and if people had to rely on reports of new phenomena, new objects and events from elsewhere, what filtering or sorting mechanisms did early modern scientific institutions use to accept some testimony and reject others? How did they go about gathering data?

I have a seemingly simple answer to the question ‘How was testimony managed?’ With difficulty.

One has to appreciate, understand, and, I think, respect the choreography in which people negotiate their way around testimony of different kinds—about different things, about how testimony relates to their own experience, and about what their own experience even *was*—because ‘one’s own experience’ is itself a historically situated actor’s category. As the cognitive psychologists have shown, people may be clear what they mean when they *contrast* personal experience to testimony, but that experience may itself be a hybrid of personal perception and prior expectations—those expectations delivered by others’ testimony and by individuals’ passage through institutions.

Back to your question: how did they negotiate their way through the world of testimony? If the short answer ‘with difficulty’, a longer answer involves *maxims* about testimony and the artfulness with which they were applied.

For example, one maxim that made sense to people then (and still does): ‘believe things that are plausible; be sceptical of—or disbelieve—things that are wildly implausible.’

Another maxim: ‘believe things testified by many people, especially when the testimony is consistent across a population.’ In other words, if many people are saying the same sort of thing, that lends it credibility.

‘Believe testimony that is immediate’—something I saw on my walk around the neighbourhood this morning counts for more than something I report about a walk from fifty years ago.

‘Believe testimony that is related in a manner that inspires confidence.’ For example, don’t believe blustering politicians; believe testimony that is presented with due modesty—without visible signs of self-interest or emotional investment.

And finally—and this leads to another kind of discussion—‘*believe trustworthy people*’. That, as I have tried to indicate, is both enormously powerful and hugely variable. Everyone has a sense of who is a trustworthy person. And that sense may differ depending on the kind of knowledge

involved and the nature of the knowledge community.

Now, these maxims aren’t infallible. For each, you can pose an exception, so to speak, a counter-maxim.

For example: ‘believe multiple sources’—but ‘don’t believe the crowd’. Or: ‘believe immediate sources’—but those might be more prone to passions, interests and momentary disturbing concerns than the testimony of, say, an established and disengaged historian writing at a temporal distance.

None of these maxims, with the possible exception of the last one about trustworthy people, solves the testimony problem by itself. And even that last one contains grounds for dissent.



You mentioned one of the maxims was to trust trustworthy people. Could you talk a bit about who, in seventeenth-century English society, was generally regarded as trustworthy by virtue of their social position?

The question is legitimate and deep—and the answer, again, must not be simple.

If one is looking at the community of practitioners of natural history, natural philosophy, experimental natural knowledge, one is already looking at a very small subset of society. The vast majority of people in early modern England were illiterate; they were not part of scientific society, not part of learned society, not part of a world in which they could, under any conditions, participate in this realm of knowledge-making.



So it is legitimate to ask: how did *this* circle of people identify trustworthy sources? How did they decide who might or might not be trusted? How did they do that, why did they do that, and what were the contexts around their notions of trustworthiness?

Now, that's not quite the same thing as asking who was regarded as trustworthy in seventeenth-century England as a whole. There's no reason to suppose—although if you're a peasant, gainsaying or contradicting a gentleman is a costly thing to do—that the illiterate classes necessarily regarded the literate class, or especially the aristocracy, as specially trustworthy.

Nor is it necessarily the case that the lower orders of society—church-going, pious people—regarded priests as inherently trustworthy. There's all sorts of folk literature depicting priests as liars and hypocrites. Likewise,

there was an emergent popular literature portraying the aristocracy as dishonest.

But these two things can coexist. It can also be the case that the code of gentility was known to people who were not themselves gentlemen. Take, for example, the audience at the Globe Theatre in London. A number of Shakespeare plays dramatize the institution of *giving the lie*, or the *honour of a gentleman's word*. And yet among the groundlings of the Globe Theatre were people who might not be gentlemen—they might not even be literate.

That's one example, among many, of how a code that powerfully regulates behaviour within one sector of society—say, the gentry and aristocracy—can be known to, and even used as a resource by, people who are not themselves gentle.



BESSON, THEATRUM INSTRUMENTORUM
ET MACHINARUM (1578)



Most of my book, *A Social History of Truth*, deals with the coordination of activity within the circles of people making formal natural knowledge—that is, a subset of society in which the presence of the gentry and the aristocracy was extremely salient.

Scholars like Barbara Shapiro have made a strong case for the importance of *legal* ways of settling matters of fact. But the circle of people who used and were aware of the codes of fact-making, negation, agreement and consensus-making in gentle society was much larger than the

circle of those familiar with formal legal procedure. Though, of course, the two groups do overlap.

In other words, if you want to understand fact-making, fact-bounding, inference-making in the context of early modern natural knowledge, then you have to understand how gentlemen went about things. That's not saying you had to be a gentleman to participate. It's an attempt to describe how this particular world *worked* and what resources were mobilized in knowledge-making in that world.



How did trust and testimony shape knowledge formations in sixteenth- and seventeenth-century England, and how did they influence the rise of proto-scientific disciplines during this period?

Stepping back a little from the material in *A Social History of Truth*, I suppose one of the things that set me off on this track was a confrontation with a kind of rhetoric you find in early modern science—some version of which persists to the present. It's a kind of empiricism and individualism.

You see it in the motto of the Royal Society, *nullius in verba*—often translated as 'on no one's word' or 'take nobody's word for it'. In other words, the authenticity, the value and the power of new scientific practices were meant to be vouched for by setting aside authority—especially

the authority of antiquity, of testimony, and of the practices of the Schools and the medieval universities. Instead, one was to rely on one's own senses—a very strong form of individualism.

I remember quite vividly having an awkward conversation that led absolutely nowhere with a distinguished scientist who claimed he didn't believe *anything* in the literature of his discipline: he took *nothing* on trust. His firmly maintained and—he thought—proper stance vis-à-vis claims made by the contemporary scientific community was to *disbelieve everything*. I gently tried to point out that this couldn't possibly be true. The very resources he would use to be sceptical of one bit of knowledge would have to come from other bits of knowledge that he had accepted on trust. I got nowhere on the subject—but it taught me, many years ago, that

this kind of individualism and empiricism is alive and well in much of modern science. And elsewhere, especially in widespread *images* of how science works.

There are two voices running through the *Truth* book—not just mine, but those belonging to many historians.

One voice tries to recover what people in a certain setting *meant* when they said something. For example, when the Royal Society said, we're not going to meddle with affairs of church or state, or we're not going to rely on testimony, what did they *mean*? From the point of view of the interpretative historian, that's a legitimate and valuable thing to do—not to *correct* past thinkers, but to try to *understand* how they thought, how they made sense of the world, how they made sense of themselves.

So one is confronted by this historical fact—the individualism and empiricism expressed by members of the Royal Society and other early modern practitioners of science. That’s one mode of sorting knowledge, of assigning value to different kinds of knowledge. That’s one question the book tries to explore—in what some commentators considered to be excessive detail: how did people negotiate their way around testimony and individual experience?

But there’s something else going on in the book that isn’t interpretative. It links to normative matters: how might we think *differently*? How might we think differently about the role of the individual, the senses, testimony and authority? How might we think differently about the nature of fact, theory and inference?

So when I say much of the book is set against individualism—and is aimed at revealing or uncovering the role of trust and testimony in making knowledge—it’s against the background of early modern thinkers in the seventeenth century who insisted quite otherwise: that they would rely on no one else’s word but on their own experiences.

One thing I try to show in the book is that early modern actors were—on a practical level—well aware of the role of testimony and trust, and they tried to manage their way through it. That hybridity, that messiness, in the categories of knowledge and knowledge-making is something I wanted to address in understanding early modern science. But it’s also, I think, a useful attitude to take towards our own predicament and our own knowledge-making practices.

In the interaction between these two seemingly conflicting tendencies—individualism and trust—how is the discourse of objectivity shaped during the emergence of early modern science?


Objectivity is not, as historians say, an early modern actor's category. I don't know if the word *objectivity* actually appears in the book. It possibly does, but it's not an actor's category. In fact, I probably should have said more in the book's introduction about the terms designating different qualities of knowledge, including the category of *truth*.

What's at issue in much of the book is how matters of fact were established. How could a register of facts be constructed and maintained amid the welter of testimony, authority and experience?

The notion of Truth does indeed appear—but in the past, as I think in the present, it tends to be a special category notion—like pounding the table to insist on a special quality of knowledge-items. That there are icebergs in the North Atlantic is a matter of fact. Truth is invoked, for example, when testimony is doubted—say, when there might be reasons to doubt particular testimony about icebergs—or in institutional contexts—say, when one swears to 'tell the truth, the whole truth, and nothing but the truth'—or in a school exam where you're asked to choose between *true* or *false*.

I believe that the book's title—*A Social History of Truth*—was useful at least as a provocation. The questions that are actually central to the book include: How does one establish facticity? How are different kinds

CIRCINVS NOVVS SVPERIORI CONTRARIVS,
VTPOTE DESCRIBENDIS FIGVRIS RECTILINEIS, APTVS,
EX ORDINE CERTARVM PYRAMIDVM CONVEXARVM,
QVAE PRO BASI HABEANT PRECEDENTES
FIGVRAS PLANAS ET CVRVILINEAS-



of facts produced in different settings? How does one evaluate testimony so that other people's experiential world can become part of one's own?

In that sense, the notions of *objectivity* and *subjectivity* don't figure very largely in the book. You do have the notion of *passions* and *interests*, which might distort or disturb a confrontation with reality or with facts. You have Bacon's Idols, culturally and psychologically rooted disturbing influences. But I think it would be facile to equate early modern categories with present-day sensibilities about *objectivity* and *subjectivity*.

This is a case where the notion of objectivity should not be unproblematically projected onto the past and then used to ask: how did early moderns sort out the objective from the subjective? Their sorting mechanism might bear a family resemblance to ours—but it's not the same.

Towards the end of *A Social History of Truth*, you turn to those who worked alongside gentlemen, the laborants—assistants and technicians—the people who conducted experiments but were not credited in print. That chapter comes with your strongest methodological caveats. You emphasize how difficult it is to reconstruct their histories due to the lack of evidence. Could you talk a bit about the absence of the technician in the scientific literature of the time?

Five years before *A Social History* appeared, I wrote a very short article published in the journal *American Scientist* called ‘The Invisible Technician’. To my surprise, it turned out to be perhaps the most cited thing I’ve ever written.

Though the article starts with Robert Boyle’s laboratory and laborants—operators or, in our current usage, technicians—it is more explicitly addressed to the question: who was a scientific author? And who is a scientific author *today*? It used the historical material to invite readers to think about the modern laboratory and about systems

of modern knowledge. I've been told that this paper has even been used in legal cases about who counts as an author.

The author—one or, increasingly, many named individuals—is right there at the top of a scientific paper. But when you go into a knowledge-making setting—as an observer or anthropologist—you see a booming, buzzing hive of social actors: coming in, going out, doing things, communicating, building—some physically present, some digitally accessed at a distance.

And yet, by convention—not a convention I want to brand as illegitimate, but a convention nevertheless—those named at the top of the paper are treated as *the* authors. That is, the people responsible for the claims in the paper are assumed to be adequately bounded by those named authors.

Again, this is—and I suppose this is something I return to often—an anti-individualist invitation to consider how knowledge is actually made, sustained, and modified.

The sociologist Howard Becker, for instance, in his 1982 work *Art Worlds*, offered a good analogy: We value paintings for their authenticity, their reliable attribution to a named artist. But what is the nature of the studio? What is the context that houses the paint-makers, the canvas-makers, the critics, the entrepreneurs, the salespeople, and so on? It's a *world* made up of different sorts of people, doing different sorts of things.

In a sense, by describing the people who are not visible, not named in the scene, I'm trying to encourage a sensibility—which, I think, is even more necessary for past science than for present science—about

knowledge-making as a collective enterprise.

Now, many readers of ‘The Invisible Technician’ responded by saying, in effect, ‘technicians *should* be acknowledged—they should get credit’. And as a personal prejudice, I rather agree with that—partly because I have a very great respect for instrumental skill and for those who create and sustain the *conditions* for scientific authors to do what they do. But both that paper and the book-chapter were trying to ask a different question: What were the grounds of seventeenth-century legitimacy by which laborants, operators—who did so *much*—were *not* regarded as authors of that work?

That leads to a broader political sensibility: about who counts, who is recognized, who has a consequential political voice. It’s tied up with things like the vote, with gender, with education, with class. These are all sortings that we may now deprecate, but which functioned as legitimate sorting mechanisms in the past.


I don’t think, for example, that Denis Papin—or the invisible Hooke—thought that Boyle appearing as the *author* of their work was illegitimate. That’s a supposition, of course. But at the very least, we should entertain the thought that this was then a to-some-extent shared scheme of legitimacy.



On a related note, Felicity Henderson, in her recent biography of Robert Hooke, writes about how he straddled both the world of the gentleman and that of the technician. He was, I believe, the first person to receive a stipend from the Royal Society. He also played an active role in institutions like Gresham College, which were aimed at communicating technical knowledge to practitioners and technicians.

If you want to look for the tensions between the gentle world and the non-gentle world in early modern knowledge-making scenes, and if you had to point to a contemporary figure who most embodied those tensions, the answer would indeed be Robert Hooke. He was at the heart of debates about intellectual property, secrecy, the intense concern for priority—all of which stand in contrast to the smooth, aristocratic sprezzatura that characterized someone like the Honourable Robert Boyle. Hooke was someone who could inhabit and move between gentle and non-gentle worlds—sometimes manifesting the friction between them, sometimes smoothing it over.

ALIUD DENIQUE TORNI GENVS NONDVM VISVM, PER POLIENDO
ET SECANDO MARMORI, AVT LAPIDI CVIVSVIS SPECIEI
ETIAM DVRIORIS, IN FRVSTA, AD ORNAMENTA
SYMPTVOSORVM ADIFICIORVM-



BESSON, THEATRUM INSTRUMENTORUM ET MACHINARUM (1578)

You have argued, in your essay ‘Is There a Crisis of Truth?’, that the non-expert’s awareness of something like climate change doesn’t come from scientific training, but from ‘social knowledge’—one’s sensibilities about who, and what institutions, to trust. There are affinities of this idea in the *Capital*. Marx argues that wage labour conceals the social relations that structure capitalist society. Do you see an analogy there, between your argument about trust and testimony, and Marx’s point about how relations of production are obscured?

In the essay, as well as the *Truth* book, what I wanted to point out was: If you think you can circumvent the issues of trust and testimony, you just *cannot*. That is, I meant to challenge those who think that the problems we’re facing in today’s crises of knowledge can be resolved simply by pumping out more facts, more evidence, more scientific truth.

In that essay, I couldn’t point to a single stable solution to this problem. As opposed to the seventeenth century, where, in many parts of society, the figure of the gentleman, and the accompanying codes of civility, offered

a powerful solution to the problem of how to make and warrant knowledge, that’s clearly not the case anymore, for all sorts of interesting reasons.

So the question is: What is the code that we use now? And just by framing the question that way, we begin to see the shape of the anxious muddle we’re in.

What I called a ‘crisis of social knowledge’ is this: we don’t trust gentlemen—who are they anyway? Many of us bridle at the notion that men—males—are to be trusted: as ‘the patriarchy’, as ‘the boys’ club’, *masculine* knowledge-making forms generate far

more scepticism than trust. We're on the way to not trusting *experts*, because they're assumed to be in the pay of ExxonMobil, or of some environmental movement. Elite institutions, notably now including Harvard, are subject to powerful modes of distrust.

So whom do we trust—if anyone?

These are, indeed, among the great questions of our time. And, again, I'd be wary of anyone who claims to have a simple answer. And I can't give a single definitive answer either—but it's evident that one way of viewing our problem is not the *absence* of recognized trustworthy sources but their *proliferation*.

People often blame this on technology—on the internet—the internet here is, I think, an *accelerant* of this proliferation and fragmentation, not the ultimate *cause*. In some ways, you

could trace fragmentation back to the printing press, the telegraph, the telephone, the television. I don't think there's something *qualitatively* new about it.

I remember being in California, in the early 1990s, when lots of people I knew in the tech community, and even in academia, were saying that the internet is a uniquely powerful force for good. It will make the world fully democratic. Everyone will have a voice. Authoritarian regimes will crumble. We will all live in peace, tolerance and mutual understanding.

What's happened instead is that the internet has delivered a kind of fragmentation and an authoritarianism that we had not then imagined. And it has delivered regimes of credibility that we have scarcely begun to understand.

The institutions that used to serve as more or less stable ways of distributing and warranting knowledge—like schools—have become part of a feedback loop with the internet. I’ve seen this happen myself while teaching at Harvard. You say something in class, and a student checks their phone in order to check *you*.

One of the most charming things I’ve experienced: I say something; the student checks her device, raises her hand and says, ‘You’re right’. That, I think, is more telling than if the student had said, ‘You’re wrong’. In other words, the professor has been *confirmed* by something found on the internet—even if that

‘something’ turns out to be something the professor himself wrote.

To go back: I don’t like the idea that the internet is what has caused all this. But you cannot tell a story about the decentring of knowledge, about the fracturing of the credibility schemes we use to make sense of the world, *without* talking about technology.

I don’t offer a solution. But I do think we need to get better at naming the *problem*. We need to talk more clearly—and more intelligently—about what this crisis actually is. And only then can we begin to think about what a solution might look like.



In 'Is There a Crisis of Truth?', you talk about the growing tendency to assume that experts have ulterior motives. You trace this, in part, to the post-war emergence of the military–industrial–academic complex. Do you think this suspicion is exacerbated by large-scale institutional failures—like the 2008 financial crisis? And is it, to some extent, irreversibly exacerbated when those failures have global consequences?

There are two things going on that relate to your suggestion.

One of them is what one might call an *effect of success*. Think of Francis Bacon and the early Royal Society. What they wanted was to firmly embed legitimate natural-knowledge-making regimes into government—into the practices of statecraft. Fast forward three centuries, and that's been achieved—spectacularly so. It's

now impossible to conceive of the state—or of the market—absent scientific expertise.

Sometimes it's useful to remind people that most trained scientific expertise is found not in universities but in government and industry. The political and business worlds have been substantially *scientized*. Science lives in practically all modern institutions; it's not confined to institutions of higher education.

And that means that when business fails, or when government fails—scientific expertise is bound to be implicated. Science is no longer ‘pure’—if it ever was. The *idea* of purity lingers, but the fact is that science has become an integral part of statecraft, of war-making, of profit-making—and, I should add, of *self-making*.

Take the so-called ‘soft,’ or human, sciences—often thought to be less powerful than the ‘hard’ sciences like physics or chemistry. But there’s an argument to be made that they may be *more* powerful, because—as they enter into the vernacular—they are resources for self-understanding; they become present in everyday social practice. That is, they make *us*. The *us* who asks these questions about the nature of science.

The credibility problems science now faces arise partly because science is no longer

distinct from, or set apart from, the state or business or the military. It has become part of those institutions, and, to an extent, its claims to credibility are not superior to, nor specific to, those of the institutions in which it now lives.

And we’ve had generations of scientists and cultural leaders insisting that the power of science must be properly appreciated. Well, that power *has* now been appreciated. What we’re seeing is the *collateral damage of success*.

The second, ancillary, point is something that was central to Max Weber’s concerns. Writing in 1917 about the *vocation* of science, Weber was describing a setting in which science was not yet a job but a calling—a way of life from which one could expect neither lasting honour nor wealth.

But now science *is* a job. And as they say in Scotland, *he who pays the piper calls the tune*. So even though we accept the presence of expertise—in service to the university, to a discipline, or to the state—as a natural state of affairs, the same condition creates grounds for suspicion.


So that critics might say in response to a scientific claim—‘Oh, you would say that, wouldn’t you? Because you want environmental science to be more recognized and more powerful than

it is now. Because you want your discipline or preferred theories to enjoy more prestige and to secure more financial support.’ That kind of response, suspicious of ulterior motives and self-interest, has become increasingly common.

While, in another mood, this could just as well be described as evidence of the great success of modern science, yet, it has compromised the credibility of science.







LAVOISIER, TRAITÉ ÉLÉMENTAIRE DE CHIMIE (1798)

Would you say that today, objectivity and the human are in some sense in conflict? That is, the ideal of scientific objectivity—particularly as shaped in Christian or Western contexts—differs from how knowledge and science are conceived in non-Christian or non-Western societies. And those

One thing your question asks me to do—something I definitely didn't do in the book or in the 'Crisis of Truth' essay—is to take a more global perspective on the knowledge economy. I didn't do that, and perhaps I should have. I have very little competence in that area, but your intervention does prompt me to tell a story that probably reflects my age and background.

There was a period—at least in Western thought—around the Cold War, especially between the 1950s and the 1970s, when it was a feature of much Anglo-

differences seem to be surfacing more clearly now.

Would you say we're past the 'age of objectivity' when it comes to addressing complex global problems? And that perhaps placing the human—rather than abstract objectivity—at the centre might be a better way forward?

American social thought to believe in the coming triumph of secularism.

It was widely assumed that there would be no more fundamentalist religious believers—not just in the West, but anywhere in the world. This was a powerful prejudice, strongly held at least by mid-century academic commentators.

It hasn't worked out that way. And I don't think I've read much commentary that really grapples with how these Western expectations were not met—were, in fact, contradicted and thwarted.

One key belief was that the advance of Western technology—telephones, refrigerators, electricity—would bring scientific ways of thinking in their wake. As Max Weber had intimated, these technologies would figure in creating a de-magified, disenchanted world. The technologies and the secular worldview were expected to coalesce. But that hasn't happened—a spectacular failure of predictive sensibilities.

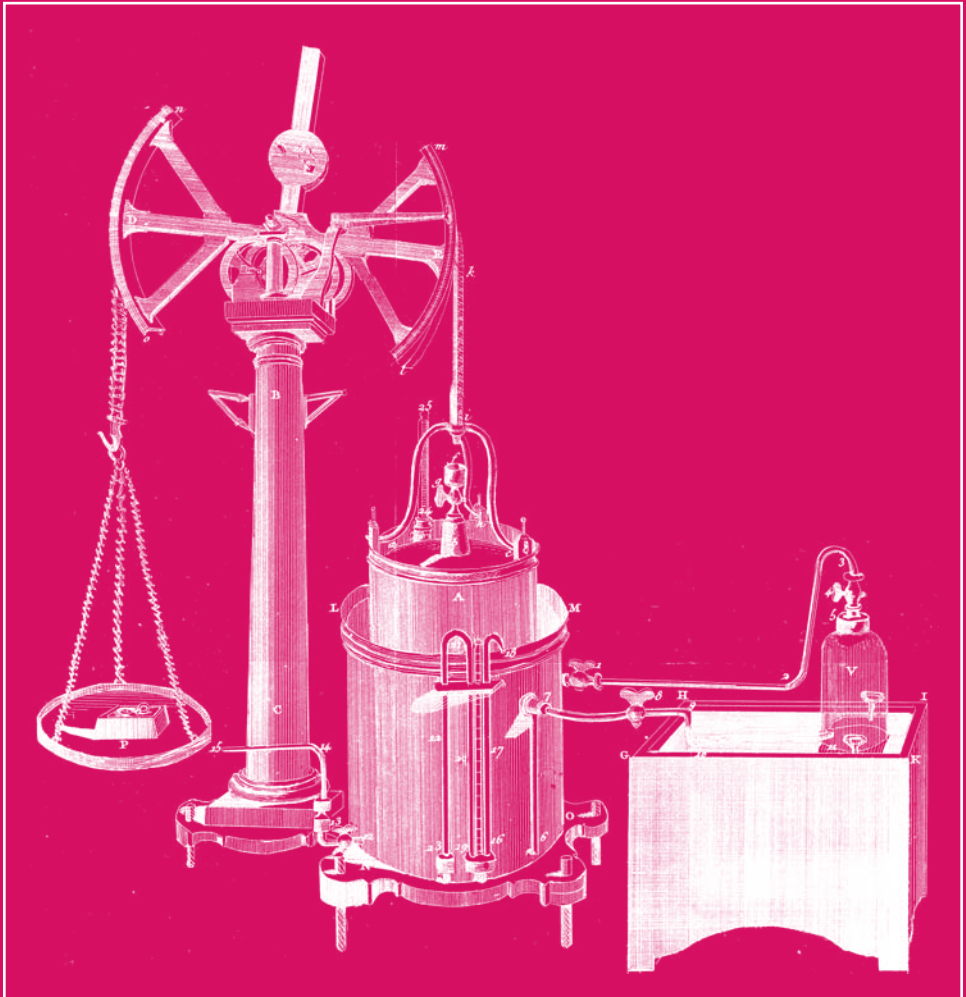
We now have fundamentalist religious groups who are

internet whizzes. And of course they are—why would one expect otherwise?

So it is interesting to confront that moment in the twentieth century when some people thought a secular worldview would inevitably follow from material technologies.

And I think the disconnect between the furniture of a modern scientific-technological world, and a worldview that was supposed to match it, is one of the defining features of the modern global disorder.





LAVOISIER, TRAITÉ ÉLÉMENTAIRE DE CHIMIE (1798)

In reading your work—especially your essays on dietetics, food, and wine—what strikes one is the idea of custom as second nature in the early modern period. You describe how diet once had strong moral associations: moderation, balance, not eating in excess. But now, as you argue, we increasingly understand diet in chemical terms—as a composition of molecules, proteins, carbohydrates, nitrogen and so on.

You suggest that in this transition from moral to

That's a question I do address in *Eating and Being*. I found myself, in this new book, trying to find the right language to describe historical change.

For example, it would be too glib to say that nutrition science—the language of proteins, calories, and so on—*triumphed*. That would be like

scientific framing, the affective and ethical force of dietary custom has weakened. A child might once have been taught not to overeat as a matter of virtue, but now they're more likely to be taught about calories and macronutrients.

But isn't the scientific, chemical understanding of diet just a layer on top of the older moral sensibility? Has that moral language really disappeared, or is it still present—just expressed differently?

using the 'death of magic' narrative of historical explanation, when the old ways vanish completely—and I don't think that's remotely right.

What I wanted to describe was a world in which, if you like, there are sediments of knowledge—accumulated layers of thought about food and its

effects. But like geological strata, you sometimes get intrusions erupting through the layers, or, at least, becoming intermittently more visible. The language of moderation *hasn't* disappeared—it's not dead—but it no longer has the authority it once did.

That loss of authority has a lot to do with the movement away from Galenic medicine—the language of the qualities of warm, cold, moist and dry and of the four bodily humours—and towards a language of chemical constituents and calories, which tends to have an amoral quality. I want to talk about this substantial historical change—but not in terms of triumph, and it's not my purpose to celebrate it or to criticize it.

So, the voice of moderation is still with us—if you listen for it or if you know where to look for it. I'm old enough to remember that my mother told

me to eat *in moderation*—neither too much nor too little. But the doctor doesn't tell you that sort of thing anymore. The doctor tells you: here are the number of calories you should consume, here's how much sodium you need, here's a pill to adjust your intake of vitamins or micro-nutrients: we've substantially moved from a *moral* vocabulary to a *technical* one.

In the 1960s, there was even a school of thought, known as *orthomolecular medicine*, that held that the quantity of, say, a vitamin needed to avoid disease—if taken in mega-doses—could make you *better than well*. That's not moderation as the way to health; it's optimization.

Many of the so-called fad diets today are explicitly set against moderation. You could even argue—though it's admittedly contentious—that the cultural centrality of a condition

like anorexia arises just because it is seen as a form of *immoderation*. The Atkins Diet—hugely popular in the late 20th century—advertised the innocuousness of dietary *excess*—so long as carbohydrates were avoided. And as fashionable people have been known to say: *You can never be too rich or too thin*.

So, yes, this is the world we live in now—a new age of extremes. But if you're asking: is morality gone? Is there no moral guest at the dinner table? The answer is: yes, there is. But the nature and modes of that moral presence have changed.


One way they have changed is in our attitudes towards animals. Vegetarianism in the seventeenth century is not the same thing as our contemporary consciousness of animal suffering. Whether religiously framed or not, we've developed a

heightened sensitivity to animal suffering—for example, in connection with factory farming, and even with philosophical tendencies to regard human beings not as masters of nature, but as an integral part of it.

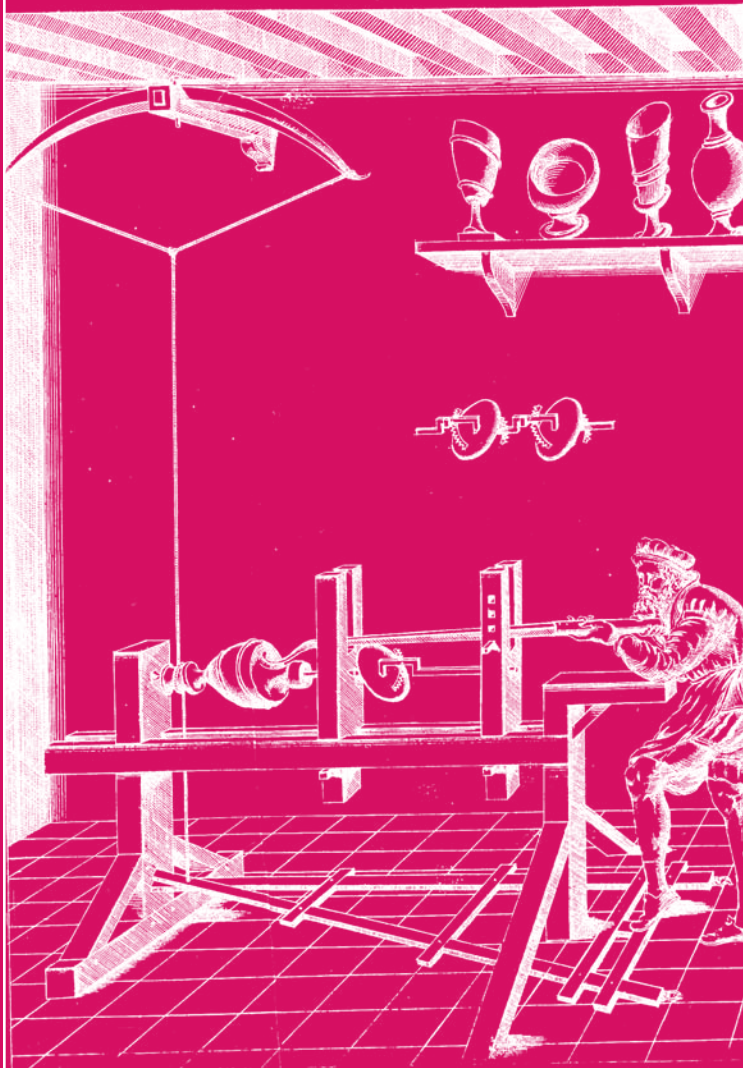
And then there's the climate crisis. Our awareness that our eating practices affect the health of the planet, the stability of ecosystems and the futures of generations yet to come—that consciousness has grown very strongly in recent decades.

Eating remains intensely moral—for these and other reasons. But the morality is no longer about crafting the self through the intrinsic qualities of food. Proteins, calories, carbohydrates—these are amoral entities. What's happened is that the moral presence at our dinner table has shifted its attachments.

TORNVS NOVVS ET GEOMETRICVS, AD REDVCENDVM
IN FORMAM OVALEM QVEMLIBET CYLINDRV, ET
CONVM, CVM SVIS ORNAMENTIS, EX OMNI MATERIA
TORNATILI-



ALTERVM TORNIGENVS, EX PROXIMO ORTYM, AD
TORNANDVM, EXCAVANDVM, ET ORNANDVM
IN FORMAM OVI PATERAS, ET VRCELOS, EX
MATERIA QVALIBET FERRI PATIENTE-



BESSON, THEATRUM INSTRUMENTORUM ET MACHINARUM (1578)


The shift in the moral aspect of diet also has to do with the direction of morality. In the early modern Galenic view, the regimen—diet, exercise, balance—was a means to living well. It wasn't morally good in and of itself. But now, with the success of nutrition science, the scientific view of nutrition has taken on a moral aura. If someone isn't eating 'correctly', they're often seen as being unscientific. Science, in this framing, has become a kind of moral image of itself.

Yes—what we're really talking about here is virtue and its place in modern dietetics or nutrition science.

If you assume that the goal of the good life today is to be healthy and to live as long as possible—and we think that everyone would want that—then, yes, nutritional advice carries an intensely moral character.

But that's also being questioned now. Do we *really* want to live forever? Is it *good* to live forever? Is it good for *everyone* to live to a hundred?

Many of our institutions—and many of our senses of moral accountability—are shaped by the possibility that scientific diet and medical science could deliver on this assumed goal of extended longevity. But all our societies are now grappling with the burden of aging—which has intensely political, economic, and, indeed, moral effects.



KIRCHER, *MUNDUS SUBTERRANEUS* (1665)

There's one sense in which, as you're suggesting, nutritional advice exists within a moral field that assumes our goal—our sense of the *good*—is to be healthy and to live as long as possible, maybe even forever.

Elon Musk wants to live forever. Many Silicon Valley tech people are seriously working on it, investing large sums in it. Which raises another moral and political question: *who* gets to live forever, if such a thing becomes possible?

Continuing the discussion on the shift from a moral view of food to an instrumental one, it seems this transformation has also shaped broader lifestyle trends.

In earlier times, there were moral codes that guided how to live a good life. But now, with the proliferation of self-help literature and the influence of online figures, we see two overlapping modalities. On the one hand, there's scientific advice—presented as empirical and rational—telling us how to live longer and more healthily. On the other, there are influencer-driven maxims, often capitalist in tone, like those associated with Elon Musk. These become quasi-truths for some people.

Take, for instance, the maxim that 'it's better to drop out of college if you want to succeed'—a kind of anti-institutional truth claim that coexists with scientific lifestyle advice. So how do you understand this contemporary moment, where empirical truth and influencer maxims coexist—and even compete—as guides for how we live?

The answer I recurrently give in other contexts is: it's a mess—fractured and complicated, or, put more properly, offering a coherent account of this moment would be not just hard but wrong. It's very difficult to describe the kind of fractured moral–scientific world we're now living in—where multiple, often conflicting, authorities shape how and what we think about living well.

There's clearly a backlash—a reaction to nutrition science itself, to the *scientificity* of dietary advice. And to the draining not just of morality, but of *affect*, from the experience of eating.

In the West, this backlash has been led by writers who are not themselves scientists—a notable example is the journalist Michael Pollan. There's been a rediscovery of moderation, which you'll find very strongly in Pollan's writings, and in his critique of *nutritionism*—the idea that we should conceive of food purely in terms of chemical constituents, and that the 'correct' mixture of those constituents will yield a healthy body.

Pollan draws heavily on the Australian sociologist Gyorgy Scrinis, who seemingly coined the notion of nutritionism. And another hugely important figure in the pushback against 'scientific eating' is Carlo Petrini, the Italian founder of the Slow Food movement, which may be one of the most consequential social movements in the world today (possibly excluding explicitly religious ones).

The power of the Slow Food movement lies in its message: don't eat omega-3 fatty acids and proteins—*eat food*. Eat real food, eat food in convivial settings with other people. A moderate pleasure in food and drink is part of what it means to be civilized—to be human. That's

the voice of the 19th-century French gourmand Brillat-Savarin, for example. And it's an explicit counter to the idea that food is just fuel to power the body and build its fabric.

There's also been a reaction to the *foodies*, supposedly obsessed with food. The reaction suggests that people should 'stop going on and on about food,' 'just shut up and eat'. In Britain, especially, critics see an equivalence between the foodie and the glutton—as if those who talk endlessly about food must also eat endlessly.

So, we can observe—not just on the margins, but more

centrally—some very important reactions to this world of nutrition science and to the fractured relationship between *diet* and *virtue*.

Historians are not good at predicting futures—who is?—but it is possible that the world in which nutrition science provides the categories through which we understand our food—and ourselves—may come to an end or at least lose much of its authority. As the climate crisis worsens, as inequality grows, as reactions to Western dominance deepen—these shifts may also alter how we think about eating, health and human flourishing.



Building on the Slow Food movement and the critiques of nutrition science, I wanted to ask about your idea of the hegemony of scientific language—specifically, the model of science and how it gets applied more broadly.

One thing you seem passionate about is wine, and in one of your interviews, you spoke about changes in how wine is written about—how experts describe it. You used the phrase ‘a ritual of objectivity’ to characterize the way people now write about wine.

Have you thought about how this ritual of objectivity relates to the distinction between the hard and soft sciences? In particular, how soft sciences sometimes seem to approximate the kind of rigor that hard sciences are perceived to have?

It was once thought—and I think still is, in many circles—that the ‘soft’ sciences were on their way to becoming ‘hard’ sciences. The hard-soft distinction, the association between hardness and high-value, and the sense of historical direction seem to be particularly Anglo-American preoccupations.

The Continental sorting of the *Naturwissenschaften* and the *Geisteswissenschaften* is not quite the same as the largely Anglo-American contrast between hard and soft sciences, and there are many reasons for this. One of them is that in the Anglo-American world, hard sciences are seen as *better* than the soft sciences; the hard is objective and the soft is cursed with subjectivity; the hard is capable of yielding technologies, the soft impotent. The world has, to a considerable extent, been Americanized in that respect—but not completely.

In the French, German and Italian academic systems (at least as far as I know), it's not so easy to say that the human sciences are just failed versions of the natural sciences. Nevertheless, keeping in mind that this is primarily an Anglo-American concern, the objectification of *taste*, the attempt to find out what is *really* good has become a significant mode of judgment. It has seemed to trump the idea that what is really good is whatever gives you pleasure, where pleasure-giving is irredeemably subjective—nobody else's business and not something to be assessed in a laboratory.

This characterizes much of the modern language of taste, particularly in the Anglo-American world, though increasingly elsewhere. And the idea that 'something is good because it gives you pleasure' has been

submerged beneath a different idea: that what is really good is measurable—scored with a number or traced to the presence, absence or proportion of specific chemical constituents—for example, in a wine.

That's part of my broader interest not only in the category of objectivity but also of subjectivity—and in how the two relate—as an everyday matter and not just in academic philosophy. I'm interested in the questions about how subjective experience—'I like this painting', 'I enjoy this wine'—fares in a world increasingly dominated by objectified ways of declaring what is 'really' good. I want to try to understand how subjective experience 'travels'—if it does; what are the conditions of 'intersubjectivity' and how that is managed.



CHEMISTRY LAB, SMITHSONIAN (1900s)

One way we objectify aesthetic value in wine is through the language of numbers and chemical constituents. Another way is through the operations of the market.

So if we shift from wine to painting, to a large extent the modern art world subjects aesthetic judgment to the logic of the market.

And here, again, you get a ‘looping effect’. If the market says that ‘This painting is worth five million dollars’, then it is—and paintings worth that much are—tautologically but interestingly—more highly valued than paintings for which people pay less. The beautiful and the expensive substantially coincide. But that five-million-dollar price involves a whole network of attribution and judgment. Is this an authentic work by this specific artist? Is the authorship confirmed, how,

Stanley Fish has voiced strong critiques of how the language of objectivity now pervades parts of the humanities.

What your reflections on wine and food provoke is the fact that we increasingly try to create objective references—even technical jargon—to describe experiences like taste or aesthetic perception. But these are domains where we have very little hope of fully conveying meaning, and yet we still attempt to pin them down with precision.

and by whom? The relationship between value, authorship and fashion is constantly in flux.

And it’s a substantial fact that both art and wine have become investment products—like many others: you needn’t ever drink the wine or hang the painting on your wall to get value from it. In this way, these aesthetic objects have escaped the domain of ‘giving pleasure’.

You've written about the shift in how we experience and describe things like wine, whisky or cigars—from being purely subjective experiences ('I like it, therefore it's good') to attempts to establish objective, even universal standards of judgment, whether through chemical analysis or market value.

One is reminded of Kant's distinction: he says that when something is merely agreeable, it's a personal pleasure. But when we call something beautiful, we're not just describing our own feeling—we're suggesting that others ought to share it. If everyone's idea of beauty were completely different, the term would lose meaning. Does that idea help explain what's going on in our modern attempts to objectify taste?

I haven't really begun to deal with these questions—and maybe I never will. But a few notes: we have different modes of interaction with one another. For example, imagine we're in a museum, looking at Dutch Golden Age paintings (which I happen to like very

much). Maybe there are a few still-lives—some cheese, some bread, some wine. And I'm with someone who says, 'There's nothing happening here. I'm not interested. It's just some food on a table.'

What do I do?

Well, I might ask whether my companion has *noticed* the way the light is used. Do you see how the white wine picks up the reflection of the oysters over there?

And the person I'm with might say, 'Oh yes, I hadn't noticed that.' And in that moment we may create a degree of consensus: there's a noticeable feature of the painting that *might* be liked, might even be thought beautiful.

Or, just as easily, the companion might just say, ‘Does nothing for me. Still don’t like it.’

What do I do then?

Well—I may say, ‘Fair enough. You like this, I like that. Let’s go look at something else.’

We have two modes at play here, and I think everyone will recognize them—whether it’s about painting, or film, or novels, or personal beauty.

On the one hand, when you say something is beautiful, you’re inviting someone else to see what you see. You’re saying: ‘There is something in *this*—in the thing itself—that should be available to you as it is to me. And I think it should prompt a similar judgment in you as it does in me’—even though there’s no certainty at all that this noticeable thing *will* become subject to shared aesthetic judgment.

That’s what we mean when we call something beautiful. So,

even without talking about universals in a philosophical sense, even without being an aesthetic theorist, I’m pointing to something in that Dutch still-life that I assume is available to you as it is to me. If you look at it—both intellectually and even physically—from the right angle, you may say, as I do that ‘It’s beautiful’. And I suppose that in any such scene the acknowledged *authority* of the pointer will have something to do with whether features *are* noticed and whether those noticeable features prompt a changed aesthetic judgment.

But we also know this may not happen. My companion isn’t swayed and doesn’t change her opinion about the painting. And so we do something else. We move on. That’s a token of aesthetic judgment in everyday action, not in philosophic discourse.

That sort of scene, I think, captures much of our social life. We've been talking a lot about *objectivity* and *fact*—but much of our actual social world, our intimate social world, is governed by *judgment*. I suggest we coordinate much of our social interaction through shared judgment, not by shared factual knowledge.

Is this professor saying interesting things? Is this bottle of wine worth buying? Is this painting as good as people say? Is this novel better than some other one?

This is the hybridity of the world we live in—the soft world of judgment.





HOOKE, MICROGRAPHIA (1665)

On the topic of subjectivity, you hint at the idea that, when someone says a painting is beautiful and someone else says it's ugly, they might not be talking about the same thing. They might be responding to different elements within the same work, meaning they're not really contradicting each other—they're just attending to different features.

That brings back to the early modernity, and to *A Social History of Truth*, in which one of the most surprising revelations was the idea that in the seventeenth century, there was a widespread belief that perception itself—our perceptual apparatus—was conditioned by class. That people of different social classes might literally see the world differently. Not metaphorically, but physiologically. Could you elaborate on that?

This is not a feature unique to the seventeenth century; it becomes especially characteristic of what's called the *age of sensitivity* or the culture of *sensibility* in the eighteenth century and the Enlightenment.

And perhaps it persists.

Take the language of *refinement*. It may no longer be grounded in supposed neurological facts or theory, but we still say that some people have refined sensibilities—that they can *see*

things, or respond to things, in ways others cannot. Think of the idea of wine connoisseurship—from the French *connaître*, to know. So, to be a connoisseur is to know better and to perceive more sensitively, to have refined senses.

The language we use to talk about people who can hear, see, smell or taste subtle differences—whether in paintings, in wine, or in perfumes—is still very much a feature of our society.

Today, we tend to regard the connoisseur as someone who has undergone a *regime of training*—and, in principle, we might believe that anyone at all could acquire those sensibilities if they went through similar training.

But every time this kind of meritocratic idea is voiced there's always someone who counters: *No, it's innate. It's biological. It's about your nervous system. It's about your genetic makeup.*

And I think that language—although no longer so explicitly grounded in class, ethnicity or gender in the way it once was—is not absent from modern culture.

There's still, I think, a quietly circulating sense that some people are constitutionally more sensitive than others. It's just been significantly stripped of the more overt language of class, race or gender.

But I would urge caution—and attentiveness. Listen carefully when judgments of this sort are being passed. When people talk about *popular taste* versus *refined taste*, or *potboilers* versus *serious literature*, or *bestsellers* versus books for *connoisseurs*.


We shouldn't assume that some form of group-based notion of sensitivity isn't still circulating, implicitly rather than explicitly.

It's no longer politic to say that the 'lower orders' are crude or unrefined. Nor is it acceptable to describe groups as constitutionally different in their innate capacity for aesthetic judgment.

But I'd be very careful before assuming that those older assumptions have entirely disappeared from our systems of attribution and judgment. They're still circulating—perhaps less visibly, perhaps more subtly—but they're still with us.



CHEMISTRY LABORATORY,
UNITED STATES NATIONAL MUSEUM
(1880s)



EDWARD CHARLES PICKERING'S
HARVARD COMPUTERS (c.1890)

Daniel Rosenberg

**The Age of
the Keyword**

My aim is to discuss the history of words, and of one word in particular—the word *keyword*—a history that, I will argue, tells us a great deal about contemporary digital culture and about how, in it, words count.

When quantification enters the area of culture, as in any other area, it requires countable objects. In the realm of text, quantitative analysis is applied without any ado. It is used on the book, as with bestseller

lists; on published articles, as with the h-index, which the dean wants professors to worry about; on the page, as with the PageRank score, the foundational analytic behind Google; and even at the level of the letter, which matters in data compression.

Yet, it seems to me, among these various countable textual units, the *word* has held a special status for a very long time.

The word is, in part, a very intuitive textual unit, particularly since the end of the medieval *scriptio continua*, in which words were written without spaces between them. The widespread and unmarked use of the word as a counting unit supports the intuition that words are well suited to that role. That is to say, when we see things like word clouds, they suggest to us not only that words are susceptible to counting, but also that we've culturally trained ourselves to expect that if you've got a word, you've got something worth counting.

We can point to many everyday examples of word counting. When I was drafting this paper in MS Word, a little counter was ticking away at the bottom of the screen. As I arrived at the start of this sentence in my draft, my word count had reached 280. By the end of this sentence, I was at 298.

There are many useful examples of word counting in the so-called digital humanities. If you've thought at all about digital approaches to humanities research, as I have, you know the first and often most devilish problem is establishing a clean, functional corpus—or, to use a term of art I love, a *bag of words*—you can count on.

Popular digital humanities tools like word cloud generators, the Google Books Ngram Viewer, the MALLET topic modelling platform and the Voyant Tools package are all, at root, word-counting machines.

The bag-of-words approach is broad but distinctive. Looking back over the last half-century of linguistic processing, you can see a contrast between word counting, on the one hand, and attempts to model natural language through syntax and structure on the other. Most scholars agree that with the rise of big data, word counting gained a significant lead over syntactic approaches to natural language processing (NLP).

In important respects, AI is the ultimate expression of that development. From one angle, this is intuitive: large language models are prediction machines that essentially do nothing but word counting.

At the same time, the patterns AI detects and replicates so thoroughly embody the kinds of structures that generative linguistics envisioned. It may be that with AI, we're finally coming full circle, where predictions based on concatenated words resemble linguistic structures in ways the old NLP couldn't achieve two decades ago, right when word counting took the lead.

It remains to be seen whether textual AI marks a decisive change in this balance of power, or whether it's simply another iteration of a familiar tension: that between numbers and structures. Either way, it's become abundantly clear that words are very good things to count with.

This is less obvious if you focus only on the culture of quantification; if you also consider the changing cultures of words, it comes into sharper relief.

In my view, over the past half-century, a certain culture of words, one that dates back to the 18th century and its great dictionary and encyclopaedia projects, was newly enabled and intensified by these modes of word counting. And that period now seems to be coming to a close, precisely because of its success. Robust data sets and vast processing power make it possible to recognize patterns in much larger, more sophisticated, and more ambiguous units than the individual word in ways that you're familiar with in word counting.

What matters to me in all this is not just the rise of quantification in textual analysis, though that's interesting; nor simply the tension between quantity and quality in these practices, which is also interesting; nor even the character of the analysis, though as a sometime practitioner, I do have a stake in how we do these things.

What matters to me is the way quantification reinforces a tendency in textual studies to focus on the word—rather than, say, the sentence—as the thing to be counted. I'm interested in how, by the end of the 20th century, the word had come to seem the countable object *par excellence*.

What I want to do now is examine a couple of figures, Raymond Williams and Hans Peter Luhn, and a couple of their texts, which form the core analysis of this paper.

To understand where word counting fits into the broader history of words, I'd like to step back about 60 years to the emergence of a concept central to how we handle words today: the *keyword*. I want to look at two foundational pieces of writing on this subject that I think helped define the terrain on which the problem of quantified words developed over the following half-century.

The first is Hans Peter Luhn's article, 'Keyword-in-Context Index for Technical Literature', published by the German émigré computer scientist in the technical research journal of the IBM Corporation. The second is *Culture and Society*, the well-known book by the Welsh literary critic Raymond Williams.

In the late 1950s, when Luhn and Williams were developing these ideas, neither was aware of the other. Each proposed a distinctive theory of the keyword. In what follows, I'll outline how these two texts laid out competing conceptions of the nature and significance of the keyword as a unit of analysis—and how these contrasting approaches eventually merged into a powerful hybrid idea.

It shouldn't take much to convince you of the importance of the keyword today. Dozens of books, articles, and online projects appear each year under the *keyword* rubric. Here are a handful of the many keyword titles published in recent years: *Modernism Keywords*, *Keywords for American Cultural Studies*, *Keywords for Modern India*, *Keywords for Asian American Studies*. These are just a taste; every year, dozens more appear. Some of my own research appeared in the recent volume *Information: Keywords* from Columbia University Press.

I must say, I enjoyed immensely the *mise en abyme* of writing an entry on *keywords* for a keyword book.

Another convenient, quantitative way to illustrate the proliferation of the keyword concept is with the Google Books Ngram Viewer. If you plot an Ngram for the word *keyword*, you'll see that a sharp upward curve in the twentieth century, suggesting the growing importance of the term, especially since the 1950s and 60s, relative to other terms and ideas in the English language.

But as you may already have sensed, hidden behind the one-line graph produced by the

Google Books Ngram Viewer is more than one sense of *keyword*. Roughly speaking, there's the sense of a *keyword* as an important word, exemplified in the work of Raymond Williams. In contrast, there's the sense implied in the phrase *keyword search*, which we find originally in the work of Hans Peter Luhn, where the identity of the keyword is momentary and situational, emerging only during and through a process of query.

These two senses are so distinct as to be almost antagonistic. And yet, somehow, in our current epistemological environment, they reinforce each other

in a sort of *collaborative contradiction*. That is, when you hear the word *keyword* used in one sense or the other, you instinctively understand that the meanings are very different. And yet, neither undermines the other. In fact, they strengthen each other and make the idea of the keyword, as we now use it, even more compelling.

This phenomenon—what I've just called *collaborative contradiction*, though you might also call it *dialectic*—is precisely what Raymond Williams sought to explore in his trailblazing literary-historical study *Culture and Society*.

In that book, Williams examined cultural and intellectual change in Europe between 1780 and 1950, through a broad analysis of language. Within that domain, he focused on just five words: *culture*, *industry*, *democracy*, *class*, and *art*. He argued that each of these words had acquired important new meanings during the period in question. These were words that not only changed, but that exemplified, as he put it, 'a general pattern of change in life and thought in the modern world.'

These were, in the terminology Williams would go on to develop over the next two decades, *keywords*.

As intellectual descendants of Williams, we often think of keywords precisely as he characterized them: words of special importance—like *culture*, *industry*, *democracy*, *class*, and *art*. That's exactly what's implied in the book titles I just listed. You don't buy a book on architectural keywords expecting just any words at all—you expect words that are central to the field of architecture.




But the keyword has another meaning in our time, just as fully naturalized: the meaning we associate with it in the context of digital search, using tools like Google or Spotlight. In that context, the keyword is just another name for any word at all.

So how do we resolve this apparent contradiction? Or maybe better: what do we do with it?

As it happens, Williams had some ideas on this subject. The keywords he was most interested in, like *culture*, had a stressed character similar to what we now associate with the term *keyword*.

In Williams's own words, 'the development of the word "culture"'—one of his key examples, along with others like 'democracy' and 'industry'—'is a record of a number of important and continuing reactions to changes in our economic, social, and political life, and may be seen, in itself, as a special kind of map by means of which the nature of the changes can be explored.'

Part of Williams's point was broadly methodological, and part was a specific characterization of the historical period he was studying. It's up to us to decide whether we still belong to that period or not. In *Culture and Society*, that period ends in the time Williams was writing,



in the 1950s. ‘The mood of England in the Industrial Revolution is a mood of contrasts,’ he wrote. And yet, Williams asserted, ‘these contrasts undergird the emergence of a common discourse.’

Here is Williams on the subject of discursive contrasts again:

‘Edmund Burke has been called the first modern conservative, William Cobbett the first great tribune of the industrial proletariat, yet to put the names of Burke and Cobbett together is important not only as a contrast but because we can only understand this tradition of criticism of the new industrial society if we recognize that it is compounded of very different and at times even directly contradictory elements.’

This is the very phenomenon I referred to earlier as *collaborative contradiction*.

Now substitute the word *information* for *industry* in Williams’s passage, and you might find that it describes not just the language and culture of the Industrial Revolution, but also our own language and culture—the language and culture of our Information Revolution.

And this brings us back to *keyword*.

In 1958, Williams was already thinking beyond his initial five keywords.

At that time, he had hoped to discuss additional words in a long appendix to *Culture and Society*, but his publisher declined. So the more extensive discussion of *keywords* he envisioned had to be deferred for what turned out to be 18 years, until 1976, when that appendix was finally published as a standalone book titled *Keywords: A*


Vocabulary of Culture and Society. By then, Williams's list of special words had grown to 110.

If you take nothing else from this paper, I would be happy if you came away hearing the title of Williams's earlier book *Culture and Society* in that of the later one, *Keywords: A Vocabulary of Culture and Society*, understanding that the second book was meant to offer a glossary for the first.

A second edition of *Keywords* added 21 more terms. A later volume by Williams's followers raised the total to 231. And while it's hard to argue with many of the later additions: terms like *desire*, *diaspora*, *emotion*, *everyday*, *fashion*, and *home*, it's also hard not to notice that as the scale of the work grows, the list begins to drift. It drifts toward another sense of keyword—toward the 'anywordness' of the term *keyword* as we often use it today.

This is also part of a digital humanities project I'm working on mapping concepts in the corpus of Raymond Williams's works. I mention this partly to reveal myself as a member of the word-counter tribe, and partly as a gesture toward another argument I hope to make elsewhere: that the roots of a quantitative approach to culture have too casually been labelled *digital*. What I want to argue through my diagrams of Williams's *Keywords* is that there's already an implicit project of structuring and counting words that underlies many analogue humanities projects in the early 20th century, which ultimately culminates in Williams.

WATCH and CLOCK WORK.



Keyword is ubiquitous today, and like so much of the furniture of everyday life, it often passes beneath notice. That's true even in a technical sense. Try doing a Control-F word search on a screen that displays the word *keyword*. There's a good chance your search will return zero results even though the word is visibly there. The reason? In this situation, the browser treats the term *keyword* not as text but as an image.

This present-absent character isn't unique to *keyword*, but the flickering presence of *keyword* on a web page like this is indicative of a structural phenomenon worth attending to: text that operates simultaneously on several different and discontinuous planes—or *panes*, in this case. Text that functions as a space for interaction, and text that's simply presented visually.

ENCYCLOPAEDIA BRITANNICA
(1768–71)

And what happens when you search *keyword* on Google? Well, the wormhole opens: a search for *keyword* throws you straight into the world of commerce. I won't go into detail about Google AdWords and other such systems, or about the business of gaming them, called search engine optimization (SEO), except to point out that commodification is another crucial vector in our larger story about the quantification of words.

For our purposes, the emergence of the word *keyword*—as a linguistic object—is the same process by which the words *key* and *word* are denatured, or re-cultured, depending on how you look at it.

Visit an SEO website and follow one of the AdWords links; you'll find services designed to help you pick the words that will generate more traffic for your webpage.

Or try a searching the term *keyword* in Google Trends. What catches the eye there is the world map showing where people are searching for a given word. It is relevant to our discussion today that, when I searched *keyword* in Google Trends, India, Pakistan, and Bangladesh were the geographic places where searches for the word *keyword* were most common.

According to Wikipedia, a search on the string ‘1958’ will tell you that it was the year Sputnik crashed into the ocean. It was the year Nikita Khrushchev became Soviet Premier. It was the year 14-year-old Bobby Fischer won the U.S. Chess Championship and the year Momofuku Ando first sold instant chicken ramen in Japan. It was also the year that the first video game—*Tennis for Two*—was introduced at the Visitors’

Day exhibit at Brookhaven National Laboratory on Long Island.

For our purposes, 1958 is the birth year of the *keyword*.

It was a quiet debut for a consequential word. When Raymond Williams outlined his keyword methodology in 1958, he didn’t yet know that what he was describing were keywords. For several years afterward, he used the term *keyword* only rarely and inconsistently. Remember, the subtitle of *Keywords* is *A Vocabulary*, not *A List of Keywords*.

Meanwhile, the computer scientist Hans Peter Luhn was also searching for the right term, trying out *keyword*, *key term*, *significant word*, and several others. In those early years, neither Williams nor Luhn felt entirely comfortable with the portmanteau quality of the word *keyword*, and both often separated the two parts with a space or a hyphen.

Indeed, as late as 1976, when he published the book *Keywords*, Williams still didn't seem to have settled in with the term. The book's entries on *alienation*, *dialect*, *improved masses*, *mediation*, *originality*, *psychological*, *sensibility* and *utilitarian* all use the two-word phrases key word or key term. In the entry *unconscious*, we find the compound word all by itself. And remarkably, in the entry for *structural*, we find both the phrase and the compound. I'm of course tempted to suggest that there's something unconscious and structural about those two variations taking place first in those two particular keyword entries, but that would just be for fun.

In fact, Williams did not fully commit to his own neologism for several years, even after *Keywords* was published. It seems that the book *Keywords* itself did much of the work of codification. That's something you can confirm quantitatively by analysing Williams's own corpus.

At the same time, I'd argue that part of what was happening during these years—both linguistically and epistemologically—is that the principles of importance and complexity behind Williams's approach in 1958 were already eroding by the 1970s. And in important ways, *Keywords* was successful because the term *keyword* was already becoming something it hadn't been in 1958—and couldn't yet have been.

Okay, so we've had our Cobbett. What about our Burke?


In the late 1950s and 1960s, while Williams in England was still only flirting with the term *keyword*, across the Atlantic on the IBM campus in Armonk, New York, Hans Peter Luhn was making a strong case for a different notion of what a keyword is and what it's good for. From 1957 to 1960, Luhn produced a burst of influential articles on electronic full-text search, automatic indexing, optical scanning, and selective dissemination of information, all foundational to our modern experience of electronic text.

On the surface, it might seem that there's little connecting the cultural critic Raymond Williams and the computer scientist Hans Peter Luhn. But following Williams's own method, I'd argue that it's precisely because of this difference that the relationship between the two usages is so notable.


As I've said, Luhn arrived at the language of keywords between 1958 and 1959, just as Williams was publishing *Culture and Society*. Yet Luhn's argument about *keyword* could hardly have been more different from Williams's.

In his 1958 article *The Automatic Creation of Literature Abstracts*, Luhn takes a radical stance on the anywordness of the keyword, proposing that the measure of significance for any chosen keyword should be statistical rather than semantic.

Here's a quote from Luhn: 'Keywords need only be defined as those which characterize a subject more than some others. To derive them, rules have to be established for differentiating the significant from the non-significant.'



THE ENIAC
Electronic Numerical Integrator and Calculator
Developed by the University of Pennsylvania
Completed in 1946





BETTY JENNINGS
(MRS BARTIK) AND
FRANCES BILAS
(MRS SPENCE),
ELECTRONIC
NUMERICAL
INTEGRATOR AND
COMPUTER (1946)

Any word can be treated as a *keyword*; its significance is determined by its relative prominence within the corpus, subject to whatever constraints the searcher applies, as with the Ngram Viewer we looked at earlier.

One rule that Luhn applied early on, and quite influentially, was a relative frequency test that eliminated words occurring either too frequently or too infrequently in the corpus to be informative. The contrast could hardly be stronger. For Williams, *significance* is a matter of importance. For Luhn, it's a ratio of signal to noise, and yet, in both Luhn and Williams, we also find some intriguing congruences.

Luhn resisted the dominant debate in information-retrieval studies in the 1950s around category systems and controlled vocabulary. To him, pre-given taxonomies were, by their very nature, foreign to documents, to texts, to corpora.

His own approach was statistical and, as he put it, *native*. In a wonderfully poetic turn, he described the advantage of this method as follows:

‘Since it is born from the collection, it reflects most naturally the spirit of the collection.’

As distant as his approach was from that of Williams, there is a kind of hermeneutic in it.

Another area in which Luhn's treatment of keywords intrigues me, as a humanist, is in his application of keywords to readers. This goes back to that earlier problem I mentioned about how we talk to computers.

This application is developed in another extraordinary article, also from 1958, titled *A Business Intelligence System*, in which Luhn outlines a method for what he called ‘the selective dissemination of information’—the great-grandmother of your X feed, Threads feed, Bluesky feed, or whichever you’re using these days.

In brief, Luhn’s system worked through feedback loops that assigned keywords not only to texts but also to readers. These loops used matching algorithms to determine which subscribers would receive which articles.

Just as documents could be searched, once bags of words were associated with readers—who, like documents, Luhn treated as factors internal to the system—those readers too could be searched, measured, and sorted.

So what’s happening in *Culture and Society*, at the moment where Williams first uses *key* and *word* in his own work?



The context is quite different from Luhn’s, but it’s related in many intriguing ways.

Very briefly, here’s what’s going on. At this point in *Culture and Society*, Williams is discussing Coleridge: specifically, Coleridge’s argument against the economism and utilitarianism of Jeremy Bentham. Williams’s strategy here, as throughout the book, is to seek a common cultural idea dialectically, through the juxtaposition of contraries. And for our purposes, it’s notable that the controversy he’s focused on at this point in the book has everything to do with quantification, and everything to do with words.

COTTON MANUFACTURE.

PLATE XIV.

Sections of one of Messrs Strutt's COTTON MILLS atelper in Derbyshire.



In this passage, Williams argues that for Coleridge, Bentham's utilitarianism was capable only of addressing, quote, 'the business part of the social, spiritual interests of society.'

'It is not uncommon,' Coleridge writes, 'for 100,000 operatives (mark this word, for words are in this sense things) to be out of employment at once in the cotton districts . . . and thrown upon parochial relief, to be dependent upon hard-hearted taskmasters for food.'

It's a big thematic point in Coleridge, but Coleridge wants you to notice the word *operatives*, and so does Williams. Here, Coleridge is taking issue with Bentham's position on unemployment. At the same time, in parentheses, he is setting out his linguistic methodology: one that is evidently influential on Williams.

In the passage, Coleridge observes that the word *operatives*, as commonly deployed in his time, performs a kind of objectification that is connected to (and perhaps necessary for) the calculating practices of utilitarianism on the one hand, and of industrial capitalism on the other. In other words, Coleridge here accuses Bentham of creating both a dehumanizing analysis and a dehumanizing neologism that facilitates it.

The results, Coleridge suggests, are devastating:

'It is this accursed practice of ever-considering only what seems expedient for the occasion, disjoined from all principle or enlarged systems of action, of never listening to the true and unerring impulses of our better nature, which has led the colder-hearted men to the study of political economy.'

Coleridge objects to what he calls ‘glib economists,’ and Williams certainly does too. There’s a story here about a particular political attitude that Williams very much wants to validate. But for our purposes, the crucial lineage is the one connecting Coleridge’s treatment of words as cultural units to that of Williams, and to our own.

In this respect, it’s notable that the passage also contains, so far as I can tell, Williams’s first usage of these two words together: *key* and *word*. And I should add that the only reason I feel confident making this claim is that I scanned Williams’s books and sorted and counted the words.

So where does all of this leave us?

I think it’s clear that we should understand *keyword* as one of the keywords of the second half of the 20th century, in the sense proposed by Williams. It’s a term that, for our period, has a force not unlike that of *culture* as studied in the book *Culture and Society*. *Keyword* is a genuinely important term for interpreting our contemporary period.

And this is not because it’s a word with a single dominating and clear significance, quite the opposite, in fact. It’s because, like *culture*, it is a term that is both emergent and deeply divided against itself, in ways that, I believe, reveal something larger about our cultural situation at the dawn of the age of AI and its expediencies.

The more recent generalization of quantification as a condition for interacting with language has thrown into sharp relief the cultural logics to which the keyword, already in the 1950s and 60s, attested, and through which modern culture appears iteratively to pass.

How does the modern concept of keywords, as markers to identify and retrieve information, differ from earlier projects like 18th-century encyclopaedias? I know those stressed alphabetization, but do you see roots of the keyword in that era's encyclopaedic thinking?

Yes, and among its most important theoretical and structural innovations, something that now seems so natural it barely registers, was its alphabetical organization. The articles were definitions of words, presented in dictionary format.

Encyclopaedias had existed for centuries by the time of Diderot and d'Alembert in the 1750s and 60s. Isidore of Seville, in the 7th century, had written an encyclopaedia called *Etymologiae*, meant to cover everything in the world. And all

through the Middle Ages and early modern period, there were many others.

If you used a printed *Encyclopaedia Britannica* any time after 1974, you'll remember its division into *Micropædia* and *Macropædia*. The *Micropædia* follows the Diderot model—a dictionary of brief entries—while the *Macropædia* contains long and even very long articles on large topics. The longest entry, at 310 pages, was for the United States, while India weighed in at 175.


Before Diderot, encyclopaedias were mostly organized by theme or system, and using them required a fair amount of prior knowledge. You had to know how to navigate lengthy treatises stitched together under thematic headings. Let's say you wanted to research birdsong: you'd likely have to read everything under *Ornithology* or *Birds*, plus another section on *Music*, *Harmony*, or *Mathematics*, and then try to connect them yourself.

Diderot asked, what if all this could be made indexable—easily accessible? What if we could just bring *sparrow* and *song* together and start from there?

That's the moment, I argue, when the *word* comes to be seen as a fundamental element in a new kind of information mechanics.

As I said at the beginning, the quantification of words in our period is epiphenomenal to this larger identification of the *word* as a kind of component within the information machine. While this talk focuses on quantification today, it's really about reaching the endpoint of that arc, where words are objectified as operational units in the information system.





MIT DYNAMO ROOM (c.1890)

As an editor, I often write book synopses for websites—text that makes the book searchable for customers and booksellers alike.

There's a tension here: authors often *don't* want keywords in the synopsis; they see them as clichéd. For example, a crime novelist might not want the word *crime* in the blurb. They want the synopsis to reflect the strangeness or originality of their language. But publishers and distributors push for keywords to be included—sometimes we even hide them using SEO techniques or meta-data.

Do you think the centrality of keywords to searchability is fossilizing our vocabulary? Are we being forced to reuse the same words over and over, making it hard to evolve new ways of talking about things?

That's really a set of embedded questions—at least three. One concerns the competing priorities of the publisher and the author. Another is about cliché. A third is about whether language is being frozen or prevented from evolving.

I'd answer each one a bit differently.

I don't think what we're doing with keywords or online language is freezing language at all—we couldn't even if we tried. That's partly to do with the long history of language, which I mentioned earlier. Language always

slips through our fingers. But it's also a reflection of the fact that the way words function and matter is changing, even as we speak.

And we have no idea what protocols AI will apply to the webpage you've constructed. It's likely (this was implied in the talk) that AI *doesn't* read your keywords as keywords. Actually, likely is the wrong word: I *know* it doesn't. Google does, at least in certain circumstances. That's where SEO comes from. But that's part of an older model.

I've written about this in a piece called 'Stop, Words'—about the idea of the *stop word*. I referred to it briefly earlier. Hans Peter Luhn made this point when computers were slower. He noted that in English, the word *the* is far more frequent than any other word. *The* is used twice as much as *a*, which is twice as frequent

as the next word, and so on in an asymptotic curve.

So Luhn, in his early keyword work, proposed a list of stop words—articles, pronouns, prepositions—that should be excluded from search because they didn't yield sufficiently useful results to justify the processing demands required to handle them.

But that world, the world of stop words, where you had to decide what *counts* as a keyword, essentially ended around 2009 or 2010, when big data got fast and cheap enough to say: 'Give us all the "the". We love the "the".'

And now with AI, we're in a new phase. AI is going to read your keywords as *strings*, not as isolated words.

So on the question of language evolution, we can't stop it. It's not only going to continue, it's likely to accelerate.

As for cliché: yes, certainly. It's the business of publishing, it's the business of *business*, to turn creative work into cliché. That's what advertising is. That's what SEO is. But that's just what it is. And that brings us to your first question, about the tension between author and publisher. In a funny way, the author is echoing something also present in Luhn. In his first version, Luhn said we can't search for words that are *too* frequent, but we also can't search for words that are *too characteristic* of the specific text.

So if a crime novel is about crime, sure—that's fine for the publisher. But the author and the reader both already assume it's about crime. The real question is: what happens next?

And that gets to something I said in our earlier conversation about the computational tools I've developed. Each is designed

to help me answer second-degree questions about the corpora underlying popular out-of-the-box DH tools. For the Google Books Ngram Viewer, for example, I built a simple tool called the gNormalizer (gnormalizer.uoregon.edu) to help users peek under the hood of the Ngram Viewer corpora in a few different ways. It normalizes curves, allows for comparisons across the different Google Books corpora, and gives you an easy way to download the data behind Ngram curves. Another tool, Correlation Explorer, not currently online, essentially runs Ngram searches backwards, finding terms with similar frequency curves.

These are examples of how, by interacting with computers and understanding their logic, we gain insight into how we use language natively as well.

Do you think the organic quality of keywords—the way Luhn described them as emerging naturally from a corpus—is now disappearing?

Since people involved in SEO or website design know how keyword systems work, they can game them. Even LLMs can now generate content that includes target keywords multiple times just to optimize search results. This makes it hard for search engines, even those using techniques like TF-IDF (Term Frequency-Inverse Document Frequency), to surface content meaningfully.

Do you think we're approaching the end of the keyword as a functional concept?

Yes. The title I gave this talk was *The Age of the Keyword*, but the title I'd originally planned was *The Vanishing Word*. I went with *The Age of the Keyword* because I wanted to emphasize the specificity of the historical shift through which we so recently came.

Embedded in my framing is the idea that this was a moment in time. I do believe the *age of the keyword* begins roughly with Diderot—let's say around 1750—and ends around 2010.

It ends right at the moment you're describing: when TF-IDF and other more sophisticated methods for filtering or gaming keyword search begin to fail, when they can no longer meaningfully resist SEO manipulation.

At that point, we can say the age of the keyword is over, because a completely new logic for generating keywords has been established.

Big data and AI mark that change. The *keyword* won't disappear entirely: we'll still have

important words, we'll still use keyword searches, but the role they play is very different now.

For example, when I was using AI to help me build the gNormalizer, my normalization tool for the Google Books Ngram Viewer, I wasn't searching with single terms like 'D3' to find help with code. Instead, I'd write something like, 'It looks like the D3 code is broken on line 15—what do you think?' That's a different model of interaction entirely.

So yes, something important is changing. And what I'm trying to understand is not just how the machine interprets me, but how *I* now speak to the machine, how my own patterns of language are shifting.

That's a great point—how our language changes when interacting with machines. But what also worries me is that, with the demise of the keyword, we now enter a space where the algorithms are so complex that hardly anyone understands what's going on behind the scenes.

When I read your article on search, I could follow your explanation of PageRank and how early Google algorithms worked. But today? What is a vector? What's a transformer?

These systems have become black boxes. It feels shamanistic: we're speaking to something opaque, like a woodland spirit, hoping it gives us the right answer. We don't know what's happening behind the curtain.

We also don't know why Google surfaces some results and not others. That's where conspiracy theories start, like 'Google's pushing a liberal agenda' or 'a conservative one.' Back when the algorithm was at least somewhat understandable, it could still be critiqued. Now it just feels unknowable.

Yes, I couldn't agree more.

Even in 2013 or 2014, when I was first working on stop words and starting to think more seriously about search, I met with a senior engineer at Google. This was a much simpler time. I asked about a well-known controversy over the high ranking that, for a time, Google was assigning to an antisemitic hate site, and I said, 'What can you tell me about how this is happening?'

And his answer was: 'I can't tell you anything.' Not because he was being evasive, but because, as he put it, 'These aren't algorithms

we write: we've written algorithms that write these algorithms.'

There's not a single engineer at Google, and not even a whole team, who can give a full account of what's going on in these systems.

So yes, *shamanistic* is exactly the right word.

That doesn't mean we should throw up our hands analytically. I do think AI may provide tools for helping us analyse what's happening in these opaque systems—but those tools will be different. They'll be inductive rather than deductive. We can't read the algorithm and trace its logic—we'll have to examine its *outputs* instead.

It makes you think in new ways. It reminds me of early moments in the so-called Scientific Revolution—where experiment and empiricism

became central, and Francis Bacon talked about submitting nature to 'trials and vexations'. Maybe now we have to vex Google and ChatGPT.

I can imagine people experimenting with search algorithms: just running tests to see what they surface.

Yes, that's something I do quite a bit. I'm not a computer scientist, so the tools I use are pretty rudimentary, but that's actually where the gNormalizer idea comes from.

The original question (and I could show you, but you can probably envision it) is based on that curve I presented for the word *keyword*. The term barely exists until the 1950s or 60s, then it shoots up on the graph. But if you noticed, sometime after (I think around 1985) it peaks, and then the curve drops.



Art Militaire, Exercice.

That's totally counterintuitive. You'd think *keyword* would be more relevant in 1995 than in 1985.

But this kind of peak in the mid-80s followed by a decline is something I've seen across many terms relevant to my research. It happens with *data*. It happens with *computer*.

And I don't believe for a second that these words were less important in 1995 than in 1985. I'm confident it's an artefact of the corpus.

So the real question is: what's in the Google Books corpus?

My working hypothesis is that the composition of the corpus changes for works published after the mid-80s. I think that

shift is probably related to how books enter the corpus: newly published books are processed more quickly than books acquired through other channels, like library scanning.

Words like *graph*, *chart*, *table*, *figure*—they behave similarly. I think those terms are especially prevalent in grey literature: instruction manuals, technical documents, project reports. And I suspect those take longer to make their way into the corpus.

So I built the gNormalizer as a way of exploring those patterns and trying to better understand what's actually in the corpus, very much in the spirit of what you're suggesting.



This is slightly off-topic, but I think you'll appreciate it. I mostly work on early modern literature, and I stumbled upon a funny quirk in the Ngram Viewer.

If you search for the word *fuck*, there's a visible spike around 1650—which seems . . . unlikely. Turns out, the OCR misreads the long medial *s* as an *f*, so you get sentences like 'French mothers are obliged to give *fuck* to their children.' It keeps happening again and again.

Yes! It's a beautiful thought, isn't it?

That's a fantastic example—and the way you put it really gets to the heart of the issue. When we use something like the Ngram Viewer, we tend to assume it reflects what people were writing or saying.

But what it actually reflects is what the OCR recognizes.

In this case, with the medial *s*, OCR errors are especially significant. But more generally, understanding the corpus, and how the tool interacts with the corpus, is essential.

To me, that's where the interesting work is right now.

Sure, it's fun to use the Ngram Viewer for patterns, like when I showed you how the curve for the ngram *Agent Orange* peaks 15 years after the curve for the ngram *Vietnam*. You can get some genuinely revealing results. But in my current research, I'm more focused on how the Ngram Viewer interacts with the corpus, and what that tells us about our culture of words, than on whether French mothers were really giving fuck in the 17th century.

This is slightly tangential to the discussion of keywords, but still related. You've written about the impulse toward classification and organization of knowledge in the 18th, 19th, and 20th centuries.

I'm wondering how we should situate the development of library science as a discipline, especially given that Dewey's Decimal System appears in the late 19th century. Is there a specific place where it intersects with the history you've mapped out?

Excellent question, and very relevant to Hans Peter Luhn, who comes out of that tradition of documentalism.

Historically, library science develops from the documentalism of people like Paul Otlet at the turn of the 20th century, and

Dewey a little earlier. So that's the intellectual lineage, and it's closely tied, especially in Otlet's case, to the encyclopaedic tradition.

But I'd argue that library science, more than anything, reflects *institutional* change rather than purely intellectual change. It really emerges out of the Second World War—specifically, out of problems in military intelligence. A key issue was: how do you process all of these captured documents?

New tools, new computers, had been developed through the course of the war. And those tools, built for the war machine, started being applied to these problems of information organization. So right after the war, you get this hybridization: the encyclopaedic dreams of Otlet and Dewey meet the institutional demands of military and intelligence agencies.

From what I've seen in the archives, Luhn is right at the centre of this. He's talking to people in libraries, and he's also talking to the CIA. And I think that tension drives the development of the field.

One interesting and often overlooked aspect of this history, which I discuss in the ‘Stop Words’ article, is its gendered dimension. Library science was a field where women were highly prominent.

Again, if you look at who Luhn was working with, who he was talking to, it’s often women. And you see the same thing in early digital humanities labs, these great male figures were running labs where the actual computing, typing and coding were being done by women.

You also see this in the intelligence world. It’s one of the government and military contexts in the mid-20th century where women were flourishing professionally.

That’s a side point, but an important one.


To come back to your question—where does library science fit?

Luhn is a perfect example of where library science stood in the 1950s. At the time, the big debate was between structured category systems—controlled vocabularies, following Otlet and Dewey—and new, corpus-based approaches emerging from the world of computing.

And a lot of those computer scientists were coming out of the military-industrial complex. Douglas Engelbart, for instance, who developed the graphical user interface we’re using right now, was working on an Air Force grant at Stanford.

So yes, 1958 is a key moment. Library science is central to it, and it gets caught up in exactly the tension this paper is about: the tension between controlled vocabulary and the native vocabulary of the corpus.

PROJECT MERCURY
ASTRONAUTS (1959)



**Lisa
Gitelman
Objectivity**

Objectivity has a history.

One could focus on ‘short’ histories of objectivity: within journalism, for example. This would include the creation of journalism schools at the end of the nineteenth century, the emergence of professional associations for journalists and the development of norms like the separation of editorial from publishing, or opinion from reporting. Eventually a whole infrastructure formed around journalism as a profession.

Journalism may be in crisis today—professional journalism certainly seems to be in eclipse—but I think it’s hard to overstate the enduring influence of some of these norms. The idea of balance, of ‘both-sides-ism’, still carries weight, and carries it beyond the narrow confines of whatever counts as ‘real’ journalism today.

But these short histories aside, I think the development of objectivity in the *longue durée* is extremely significant: it’s a concept that has been at once highly valued and deeply contested. Lorraine Daston and Peter Galison, in their book on the subject, call it an ‘epistemic virtue’—an orienting ideal. One should note that calling objectivity an ‘orienting ideal’ isn’t meant as praise. From some vantage points, objectivity has a rather dubious history.

The broader question I want to pose is about our ‘ways of knowing’: a phrase that can be frustratingly vague.



The *Meno* is set a generation before Plato's time, just after Athens' defeat in the Peloponnesian War—a moment of political and intellectual uncertainty—and features two figures who were already dead by the time Plato wrote. One concern in the dialogue is Meno's paradox—or the quibbler's dilemma. Meno, frustrated by Socratic questioning, asks: how can we even begin to know anything, if we don't already know what we're looking for?

Related to this is the central question: *Can virtue be taught?* While the dialogue seems to end in the negative—virtue is a gift from the gods—there are ways to read the dialogue as offering an emphatic *maybe*. Perhaps virtue *can* be taught—if there were someone like Socrates to teach it.

But Socrates sidesteps or misrepresents Meno's quibble, and yet together they frame a real problem about knowledge—one that's still useful for us.

The *Meno* marks Plato's turn to the idea that learning is a form of remembering. What Plato is trying to get at is *a priori* knowledge: knowledge that comes from reason alone, from within.

This is in contrast with empirical knowledge—knowledge from experience. And this plays out in a famous scene where Socrates questions an enslaved person, who then becomes a kind of paradigmatic knower. Guided through the Socratic method, he works out the square root of two, or the hypotenuse of a triangle.

Thus, the dialogue frames *reason* as a way of knowing, even if, in practice, reason is very hard to isolate—you can't only know by reason,

despite the way Plato sets it up. Even Socrates and Meno have an interesting exchange about how virtuous fathers often have less-than-virtuous sons—an argument drawn from experience. Reason is perhaps never entirely unmixed, but the *Meno* still gives us a way into that discussion.

Thucydides' *Peloponnesian War* offers an immediate contrast—philosophy and history from the same world. Interestingly, Thucydides was held up in the nineteenth century as a model of historical objectivity. Today, it's easier to spot his biases.

Most obviously, he calls his work *The History of the Peloponnesian War* because he's Athenian. If he'd been from the Peloponnesian Peninsula, I suspect he might have called it *The Athenian War*. But Thucydides lets us ask questions about historical knowing—as distinct from philosophical knowing.

He lays out his method very clearly in Book 1. He's writing prose, he insists, not poetry—he wants to serve the truth, not please the ear. He goes out of his way to distinguish himself from Homer. And there are plenty of contrasts between this account of war and Homer's earlier rendering of the Trojan War.

Thucydides claims his statements are supported by evidence. There's an empirical sensibility here—firsthand reports, use of documents—rather than pure reasoning. He explains events in human terms. There's no *deus ex machina*, as in Homer. He praises his own painstaking, laborious efforts to assemble a useful history.

But any quick survey of the *History of the Peloponnesian War* reveals that its narrative is punctuated by orations—most famously, the Periclean funeral oration in Book 2. How did Thucydides faithfully reproduce these speeches? Thucydides tells readers that he reconstructed these speeches with reasonable likelihood, based on the occasion and what would have been *appropriate* to say.

The fascinating thing about Thucydides is how historical knowing emerges as a kind of interplay between events and occasions. If the nineteenth century saw him as an exacting chronicler, it's now more accurate to think of him as an artful reporter.





Alberti's *On Painting*—a fifteenth-century treatise on painting from the moment when linear perspective was being developed as a technique—is another watershed moment. What Alberti gives us, in the simplest terms, is the idea of *point of view*—a key concept in the history of objectivity. Linear perspective depends on a single, monocular point of view to render the world via graphical illusion. Alberti offers an account of a nonverbal way of knowing, centred on the eye—both the 'I' and the literal eyeball. In that sense, he gives us an entry into how, especially in the Western intellectual tradition (but not only there), seeing becomes a paradigmatic metaphor for knowing.

I tend to think of linear perspective as a technology of the room. You can see the rays Alberti describes—how the architecture of a three-dimensional world gets rendered in two dimensions. It has symbolic resonance for someone like Leonardo da Vinci, certainly, and other painters of the Renaissance. And once linear perspective became dominant as a representational and perceptual technique, it would open the door for artists to push against it: think of Picasso or Braque around 1912—cubism is riffing on linear perspective, breaking it apart in various ways.

Perspective isn't a neutral transcription of reality. It's a representational programme—and a deeply cultural one. I try to push students to reflect on that: What's the difference between a world *before* perspective and one *after*? What does it mean to see the world—or to render the world—in this way? Is this a realist or positivist method? Is it responding to the world as it is, or producing a way of seeing?

Alberti writes about pyramidal rays shooting out from the eye—a visual geometry that underpins his account of perspective. And the kind of knowing he’s invested in is quite distinct from the philosophical or historical modes. He wants his ideas to be *practical*.

Good painters, in Alberti’s view, must think—but they must also be trained. The hand knows, in Alberti. And painting, as well as knowledge, for him, is grounded in comparison. Among the Spanish, he writes, many girls appear fair who, among the Germans, would appear swarthy. Painting doesn’t uncover essential truths, as Plato might seek, or narrate events, as Thucydides does. Instead, it produces truthfulness as a visual *effect*. Alberti, then, is all about appearances, not essences.



Montaigne wrote his essays during a time of intense political and religious upheaval in the sixteenth century. A devout Catholic, he eventually retired from public life to his château near Bordeaux.

His father had him learn Latin as his first language, and his essays are punctuated by quotations throughout. Montaigne read using the technique of commonplacing—collecting fragments, maxims, excerpts—and those commonplaces make their way directly into his writing.

Like Thucydides' tissue of events and occasions, Montaigne gives us a kind of tissue too—one made up of reflections and what I like to call 'flowers': the commonplaces he gathers from his reading.

You could argue that Montaigne invented the modern essay. *Essay* means 'trial' or 'attempt', as well as the literary form. And if we're trying to pin down essaying as a way of knowing, Montaigne's crucial contribution to the history of objectivity is that he created the form. But with that, he also helped to invent something like modern *subjectivity*, a recognizable sense of selfhood.

This is a watershed, since from Montaigne onwards, we can think of objectivity as an epistemic virtue that takes shape in contrast to subjectivity. The slippery language of objectivity gives way to a kind of proxy. Rather than having to define objectivity up front, we can sneak up on it by exploring different varieties of subjectivity. And with every new kind of subjectivity that emerges, we get a corresponding modulation in what counts as virtuous knowledge.

We can't say Montaigne is fully modern, but he *is* recognisably a self. In the *Essays*, he uses himself as an organizing logic—arranging topics like idleness, cannibals, coaches—all anchored in his own perspective. He also uses himself as the anchor for all the excerpts and quotations he includes. So even if he's not quite 'modern' in our sense of the word, selfhood (i.e., subjectivity) has real purchase in these texts.

And in retrospect, I think Montaigne helps us explain some of the sheer difficulty of Plato and Thucydides. With Plato, it's often hard to separate Plato from Socrates. In Thucydides, it's hard to separate the author from Pericles. And part of what makes these texts challenging might be the absence within them of recognizably modern selves—at least in the form the texts have come down to us.





The *Philosophical Transactions*, the Royal Society's journal, marks early modern science as a distinct way of knowing. Unlike Socratic reasoning, which is deductive, early modern science is inductive and empirical. It's also emphatically *useful*—directed toward the collective advancement of knowledge. That collectivity, in this case, is the scientific elite: the fellows of the Royal Society, who function as witnesses and correspondents, building knowledge together.

Early modern science represents a radical departure from religious explanation. It depends on a secular discourse, and it's explicitly *non-rhetorical*—or claims to be. It strives for a plain, transparent language, which later scholars have pointed out remains a core part of the scientific ethos. That idea of transparency is still baked into how we think about information today.

The scientific journal—like Montaigne's essay—was a new knowledge genre. And one of the most interesting things to note, thanks to the work of historian Adrian Johns, is that the authority of science and the authority of *printedness* were co-produced. They evolved together. As an aside, a question to consider is: what's happening to the authority of printedness today? And is that related, in any way, to the authority of science?

This question can serve as a provocation for a presentist interlude.



In her essay ‘Ground-Zero Empiricism,’ Lorraine Daston creates what I’d call a kind of wormhole—a link between 2020 and, say, 1665. In the first issue of the *Philosophical Transactions*, we learn that it was *very* cold that year—so cold, in fact, that Robert Boyle’s ‘New Experiments and Observations Touching Cold’ was delayed in the press—apparently the ink was too stiff to print. 1665 was also a plague year. And Daston draws a parallel between that moment and our own.

COVID threw science into crisis. A crisis about *how to know*. What is this new phenomenon? What even counts as a phenomenon? The first *Transactions* posed similar questions. Is it a monstrous calf? A spot on Jupiter? How do we define what we’re trying to know? How do we study it? Measure it? Experiment on it? What’s happening? Why now? And most urgently: what do we *do* about it?

Daston writes that in 2020, as in 1665, none of these basic questions had an agreed-upon answer. We always lack knowledge—but in moments like these, there’s no settled script for how to proceed.

By contrast, in ‘Beware the Magic of Metrics,’ Arjun Appadurai and Paula Kift offer a meditation on metrics—and an opportunity to think through the language we use during such crises: terms multiply, and they start to shift registers. Metrics, such as those we saw on COVID dashboards, have an aura of transparency, even comparability. They draw us in. But while facts are, by definition, true, metrics are not—they’re rhetorical.

Daniel Rosenberg argues that facts are ontological—they make claims about what *is*—whereas data are rhetorical. (Your data can be incomplete, insufficient, wrong.) Data are part of an argument, or at least something an argument depends on. And metrics? Metrics are deeply rhetorical—they shape what we can see. The dashboards didn't represent the world *as it is*, but as it can be rendered knowable—through quantification.



For many of us today, sociological knowing is a kind of default. It has become naturalized. W. E. B. Du Bois, in *The Philadelphia Negro* and *The Souls of Black Folk*, helps us see how this tendency emerged, and how contested it could be. Du Bois presents a powerful articulation of sociological knowing and also exposes its tensions, its discomforts. As an investigator—but also as a potential subject—Du Bois calls himself a *problem*, a social problem: the problem of Blackness in America.

In the US, sociology has a complicated origin—entangled in its early days with eugenics, but also with reformist aims. It takes on a new knowledge object: *society* as an organized collectivity. Nation, class, race—all become structuring formations that locate the individual in dynamic ways.

Du Bois described sociology as attending to the ‘great near’. He was drawing a contrast with the anthropologists of his generation—those travelling to distant lands to study ‘exotic’ cultures. Du Bois brought the focus home, to American racial capitalism and to questions of reform.



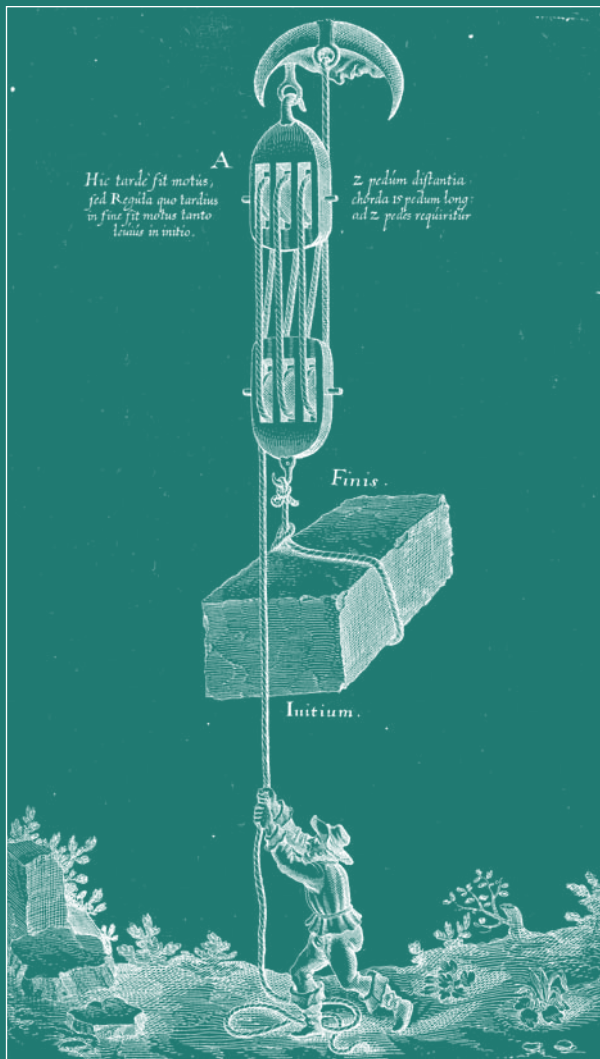
If Montaigne gives us selfhood, and Du Bois gives us double consciousness, then Freud and Borges open up the psyche.

In Borges's 'Funes the Memorious', we are presented with a riddle: how can an account of someone with perfect memory be written or narrated by someone *without* it? The narrator, in fact, presents himself as fairly forgetful.

What is presented here and in Freud's 'The Mystic Writing Pad' is a layering—a two-level model of the mind. In Freud, the conscious and the unconscious produce a new understanding of memory and repression. 'The Mystic Writing Pad' is a child's toy with multiple layers. Freud uses it to illustrate how traces of memory accumulate below the surface—even after the surface appears blank.

Whereas in Plato, you *remember* what you already know (that's what reason is), in Freud, you *know* things that you *can't* remember. The unconscious holds what has been repressed—shaped by libidinal impulses, blocked desires, and so on: a radically different model of memory, and one that invites us to reflect on the present moment as well.





FLUDD, UTRISQUE COSMI . . . HISTORIA (1617)

Alfred Loos' *Ornament and Crime* is a strange, unpleasant document.

An artist and designer, Loos helped to shape the modern world in some very tangible ways. He introduced a sense of minimalism to an ornate world. And in his time, the spartan minimalism was scandalous. Clean lines, an emphasis on function over form—Loos has left his imprint on all our rooms, furnishings, and facades.

Loos imagines civilization progressing along an evolutionary path—and for him, ornamentation is a crime because it disrupts *function*. Ornament is excess. For Loos, it is a regression. Loos was a developmentalist, though in ways that are often unsettling. Consider Darwin's concept of vestigial organs—like the appendix I still carry with me. Or Freud, for whom neurosis is simply normality taken to an extreme. And Marx, who saw surplus value as something extracted from labour for the benefit of capital. Loos seems to twist up all of these developmental frameworks—Darwinian, Freudian, Marxian—and apply them to culture; it was an ugly philosophy that underwrote what proved to be a lasting modern aesthetic.



One way to structure reflection across these very different texts is to think about *paradigmatic knowers*. In Plato: the philosopher—but also perhaps the enslaved person or Meno, the bewildered interlocutor. Then we have the historian Thucydides, the painter Alberti, Montaigne the essayist, the scientist Boyle, the sociologist Du Bois, the psychoanalyst Freud and their subject. And then Loos: is he a knower? A designer? An architect? He was, by all accounts, a horrible person—and also a convicted paedophile—so he troubles any neat moral or epistemic categories. What kinds of *knowledge objects* are being examined across these texts? We've seen essences, events and causes, occasions, Alberti's appearances rather than truths, the self, social groups, statistics, psychic repressions, surplus functions and modernity itself.

Another axis for analysis might be *genre*. What are the forms through which these different knowledges are expressed? In several cases, we're dealing with newly invented forms—Montaigne's essay, the scientific journal, the infographic, the manifesto. And finally, we might look at *epistemic virtues*—the different ideals of knowing embedded in each case. Objectivity, if we think of it that way, becomes domain-specific. It's not a fixed trait, but a constellation of virtues assembled for particular purposes.



Keiji Nakazawa's *I Saw It* offers, as it were, an unhappy ending for this long history of objectivity. Because of the Holocaust, because of Hiroshima and the bomb and what events like that require witnesses.

As a comic book, *I Saw It* offers us a chance to reflect again on visual representation. We might say that an analogue photograph is an *index* of what was in front of the lens—light bouncing off objects causes the image. So it has an evidentiary claim. Photography becomes a fascinating site to think about the border between objectivity and subjectivity.

But even if we grant that, drawing or painting—or cartooning—is also a kind of index. It's just not an index of the external world. It's an index of the artist's hand. In that sense, it's *autographic*.

Photography, too, can be more than an index. It can act as a shield, a buffer against experience. So we might ask similar questions about the cartoonist: what kind of witnessing is possible in this form?

And that brings us to the question of *eras of witness*. Thucydides tells a hair-raising account of the plague in Athens—again, resonant with 2020. 'I caught the plague myself,' he writes, 'and witnessed others suffering from it'. That's one form of witness.

The early modern scientist, observing laboratory proofs, offers another. Du Bois—with his sociological clarity and his awareness that Black America is seen as a 'problem'—represents another kind of witnessing. Later on, photojournalism and documentary photography carry forward that ethic, shaped by professional norms.

But in the post-war period—after the Holocaust and the bomb—we arrive at a new urgency around the figure of the *surviving witness*. Eyewitnessing, of course, is notoriously subjective. And yet: the survivor of trauma is a figure to whom we grant a special kind of authority. A kind of epistemic privilege. That, too, is part of our history of knowing.

ASTRONOMERS AT THE NAVAL OBSERVATORY IN DC
VIEWING ECLIPSE OF SUN (24.01.1925)



**Mathew
Stanley**

**Solar Eclipses
in History**

We have special words for talking about the future. We talk about ‘prediction,’ ‘prophecy,’ and ‘divination’ because the future is difficult. The future is yet to be made. It’s uncertain and mysterious.

The past is different. It’s no special mystery, and talking about it is usually straightforward. After all, the events have already happened. But astronomers disagree with this. The laws of nature they use to predict the motions of celestial bodies have no preference for the direction of time. The future is

as set as the past, and the past is as uncertain as the future. In principle, predicting a solar eclipse for next year is exactly the same as describing one that happened a thousand years ago. For astronomers, then, the past sits in this peculiar mix of certainty and confusion. The past and the future yield equally to their mathematical techniques, but the past is also as concealed as the future.

This essay is about a series of nineteenth-century controversies that stemmed from this odd situation. The point of contention was something that should have been simple: when and where certain solar eclipses happened. Some argued that historical records settled the matter clearly. Others believed the equations of astronomy did. And when the chronicles and the calculations disagreed, which side should give way? Should the histories be changed to match the equations, or should the equations be tweaked to match the histories? The answers weren't obvious, even for astronomers, and they relied heavily on deeper assumptions about the nature of science itself.

The story begins with a battle of no particular consequence. About two and a half millennia ago, two powers warred in Asia Minor, a conflict narrated for posterity by the Greek historian Herodotus. The

Medes gained many victories over the Lydians, and the Lydians many victories over the Medes. Among their battles was one particular engagement. As Herodotus tells it, just as the battle was growing warm, day was suddenly changed into night. This event had been foretold by Thales, the Milesian, who warned the Ionians in advance, fixing the very year in which it would occur. When the



JOHN OF HOLYWOOD,
DE SPHAERA MUNDI (c.1230)

Medes and Lydians saw the change, they stopped fighting and agreed to seek peace.

Astronomers became interested in this passage because of its apparent description of a solar eclipse: day suddenly turning to night. And if Thales truly predicted it, this would be the first recorded successful prediction of a solar eclipse. That's a historical novelty. But what really mattered to historians was that the eclipse could be used to date the battle precisely. An eclipse is just a peculiar result of the Moon's motion, governed by the laws of gravity. Astronomers could use those laws to predict future eclipses with remarkable accuracy. Perhaps more surprisingly, the same laws work just as well for the past. So a clever chronologer trained in astronomy could simply calculate the date on which the Moon would have caused an eclipse over that battlefield, and thus fix the date of the battle. But this technique didn't give the certainty they hoped for. Isaac Newton dated the battle to 585 BCE. Scaliger said 583. Calvisius said 607. Ussher said 601. Bayer calculated 18 May 603 BCE, while Volney argued for 3 February 626 BCE. So, no certainty was to be found.

The project was taken up again in 1811, with great confidence, by the British astronomer and famous eclipse-hunter Francis Baily. He insisted this was a problem only astronomers could solve, not historians. He said previous attempts had failed because their lunar tables were 'exceedingly defective and incorrect.' Baily's access to Laplace's newer theory, he claimed, put the 'visionary speculations of the antiquarian' in competition with 'the immutable laws of nature.' The result

of Baily's calculations was that the battle took place on 30 September 610 BCE. He insisted historians should be grateful for this calculation. He suggested their numerous errors were easily detected by our knowledge of astronomy, and that the unintentional errors of the historian were soon rectified and adjusted. So this set a particular hierarchy of knowledge: astronomy was reliable; history was not. Only astronomers, he claimed, could speak accurately about the past.

Baily went on to try dating other eclipses mentioned in Greek chronicles, but encountered a puzzling lack of success. His calculations couldn't quite place the path of the eclipse in a likely location at a likely time. This suggested to him that the equations used to predict the Moon's motion might need to be altered. Citing a lack of time, he declined to make those changes himself. But that offhand comment actually pointed to a very difficult technical problem: the technique of dating past events by eclipses assumes that the Moon's motion doesn't change with time. The tables describing its current motion still apply thousands of years into the past or future. Sometimes that's fine, as with Baily's confident calculation. But sometimes, modern tables place an eclipse somewhere completely different from where it was reportedly observed, or at a different time altogether.

The easiest fix is to add a small correction term to the Moon's motion, which shifts the eclipse to the right place and time. This change is called the secular acceleration. The existence of this acceleration was first suspected by Edmond Halley in 1692 and confirmed by eighteenth-century astronomers. Getting an accurate value for it is crucial, since without it, it becomes very difficult to predict the Moon's

motion more than a century or two away from the present. Unfortunately, calculating this value is tricky. You need accurate data on where the Moon was in the distant past. And here's the irony: because of how the equations work, the further back in time the observation, the more useful it is. So historical eclipse reports are exactly what's needed. But using those records proved difficult. There was significant disagreement about how best to do it.

Baily's task of working through the Greek eclipse reports was taken up in 1851 by the Astronomer Royal, George Biddell Airy. He saw himself as building directly on Baily's work, although he stated, quite plainly, that there could now be no doubt he had fixed on the wrong eclipse. What exactly he had done wrong became the heart of a controversy that lasted the rest of the century. Airy argued that the problem wasn't Baily's mathematical skill or astronomical insight. The problem was that he wasn't a good enough historian.

Airy set out to demonstrate the right way to use Greek eclipse reports. He began with an eclipse mentioned by Diodorus Siculus, said to have been seen by Agathocles, the tyrant of Syracuse, in 310 BCE, as he fled a Carthaginian blockade.

Unfortunately, Baily's lunar equations showed no eclipse anywhere near Syracuse at the recorded time. To fix the equations and get the eclipse to line up with Agathocles meant adjusting the secular acceleration. But to do that, Airy needed to know exactly where and when Agathocles saw the eclipse. And that was a problem. Diodorus tells us that Agathocles escaped the blockade and sailed to raid Carthage but doesn't specify his route. He might have gone north or south around

Sicily, and the route he took would produce different values for the secular acceleration. There was no way to solve that astronomically. So Airy decided this was fundamentally a historical problem. Baily had assumed Agathocles took the southern route, which was the standard view. Airy, instead, consulted Captain William Henry Smyth of the Royal Navy, an expert on ancient naval tactics. Smyth's insights helped Airy identify, with what he called 'perfect certainty,' a very different locality.

Looking more closely at Diodorus, Airy noticed that the party passed by some quarries en route to Carthage. He became convinced that these were the same quarries from which Utica and Carthage had been built, which led him to place Agathocles' landing at Al Haouaria. This bit of classical detective work pointed to the northern route. Airy adjusted his equations accordingly.

Airy went on to reanalyse the eclipse of Thales in a similar fashion, consulting military historians to determine where the battle had taken place. He concluded there were only two practical routes for armies crossing Asia Minor: Melitene and Issus. Issus had been the site of a number of other notable marches, including those of Cyrus, Alexander, Valerian, and the First and Second Crusades. The trick, Airy said, was getting inside the heads of the military commanders of the time. And it was precisely this kind of historical insight that he claimed Baily lacked. Airy was making the case that it was reasonable to base astronomical equations on historical texts, arguing he could do so with certainty.




GRACE AND CALVIN COOLIDGE,
ECLIPSE OF SUN (24.01.1925)

Why was he so confident? Well, like many men of his generation, he had been steeped in the classics from a young age. During his time at Cambridge, he spent as much time reading Greek and Latin as he did studying mathematics, and proudly listed in his autobiography all the classical texts he had conquered. For the rest of his life, he continued studying the classics and ancient history daily. One can imagine his delight at discovering a project like the historical eclipse problem, which called on both his astronomical and historical skills. The problem of secular acceleration was critical to astronomy, yet solving it, in Airy's view, depended utterly on historical methods. To him, and to many other classically trained Victorians, it was no surprise that ancient chronicles still had vital lessons to teach the modern age.

The decades after Airy's first foray into ancient eclipses saw a number of debates over the mathematics of secular acceleration. But the belief in the reliability of ancient chronicles, and the practice of relying on historical records for astronomical data, remained widespread.

That is, until 1878, when a major challenge was issued by an upstart American named Simon Newcomb. And if one were looking to contrast the scientific worlds of Victorian Britain and industrializing America, you could hardly do better than Airy and Newcomb. Where Airy had been shaped by a classics-infused grammar school and a Cambridge tripos education, Newcomb was a self-taught astronomer from the backwoods of Canada. Known for his manic energy, he eventually made his way to Cambridge, Massachusetts, and then to the US Naval Observatory in Washington. He gained a reputation for extraordinary skill in positional astronomy, specifically calculating the motions of planets. And for Newcomb, the sheer intensity of computation and rigour required for such work represented the pinnacle of human achievement.





TYCHO BRAHE, DE NOVA STELLA (1573)

He once wrote that an astronomer must ‘weigh the bodies, predict their motion in all future time, compute their orbits, determine what changes of form and position these orbits will undergo through thousands of ages, and make maps showing exactly over what cities and towns on the surface of the earth an eclipse of the sun will pass fifty years hence, or over what regions it did pass thousands of years ago. A more hopeless problem than this could not be presented to the ordinary human intellect.’

He went on: ‘There are tens of thousands of men who could be successful in all the ordinary walks of life, hundreds who could wield empires, thousands who could gain wealth, for one who could take up this astronomical problem with any hope of success. The men who have done it are therefore in intellect the select few of the human race, an aristocracy ranking above all others in the scale of being. The astronomical ephemeris is the last practical outcome of their productive genius.’

So, for Newcomb, the power of astronomy lay in its ability to speak of the past and future equally. At the time, it was widely believed that the motion of the Moon would also yield to this power, and that its movements, past and future, were fully determined. But in 1870, Newcomb concluded that this was not the case. The existing lunar tables, he believed, had serious flaws, and the key issue in determining the Moon’s motion remained how to calculate the secular acceleration.

Airy and most of his peers believed that one must calibrate the equations using ancient eclipses. Newcomb disagreed. He declared, ‘The ancient total eclipses of the sun, which have been so much discussed during the present century, are here thrown aside.’ He simply was not willing to take historical records at face value. In doing so, he implicitly criticized earlier investigators, stating that he, unlike them, approached the subject ‘without any bias whatever, unless a general distrust of the scientific precision of ancient authors may be regarded as a bias.’ His main critique of using documents like Herodotus’s was that passing references in historical chronicles were not the same as astronomical observations.

We cannot even be sure those ancient passages were actually documenting eclipses. Newcomb wrote that the statements in the chronicles were so vague they might well refer to other, less rare phenomena. He warned: ‘It must never be forgotten that we are dealing with an age when accurate observations and descriptions of natural phenomena were unknown, and when mankind was subject to be imposed upon by imaginary wonders and prodigies.’

In a time of mythology and epic tales, Newcomb argued, we should not expect writers to adhere to modern standards of evidence. He then launched into an impassioned critique of each ancient eclipse previously used to calculate the secular acceleration. Unlike Airy, though, Newcomb could not simply sit down and ponder the original Greek and Latin. He had to call in help. His first target was the eclipse of Thales, the most discussed case. According to Newcomb, there were three clear facts:

There was a battle that ended with a sudden onset of darkness.

On 28 May 585 BCE, the shadow of the Moon did pass over Asia Minor (according to astronomical tables).

Thales predicted eclipses.

Most astronomers had simply concluded that these were obviously connected. But not Newcomb. He wrote: ‘That these propositions all refer to one and the same event, I see no sufficient reason for holding.’ Their connection might be real, he conceded, but its reality was not well enough established to base changes to the lunar equations on it.

He continued: ‘It seems to me that most commentators on this eclipse have not sufficiently distinguished between the phenomenon as seen by the contending armies and the conclusions drawn by the Ionians that that phenomenon was what their favourite philosopher had predicted.’


Newcomb was drawing a critical distinction here. Herodotus’s account might be true. Airy might even have been right that it described a real eclipse. But that kind of historical truth was not enough to build astronomy on. The equations of astronomy were tools for divining the past and future with almost superhuman accuracy; the kind of certainty offered by Herodotus just did not meet that standard. As Newcomb put it, considerations of ancient eclipses show that ‘there is no one of these eclipses we can feel reasonably confident was total at a definite point.’ Even the best were not good enough to use for calculations.

But this left him with a problem. If he rejected all ancient eclipses, how could he calculate the secular acceleration? The entire point of using ancient eclipses was that even crude observations, if far enough back in time, were more valuable than very precise ones from the present. Newcomb refused to accept that dilemma. He believed there had to be another option.

In 1870 he took a trip to Europe, ironically, to observe an eclipse. He used the opportunity to search for suitable unpublished records in the continent's observatories. A visit to the most promising location, the Paris Observatory, had to wait until after the Franco-Prussian War. Once in Paris, he had to pick his way through barricades and carts, consulting his map alongside the barricade itself, just to reach the observatory. He spent his days inside, mousing among the old records, within hearing of the besieging guns. He could see the flashes of artillery from his window.

His efforts paid off when he discovered a trove of raw eclipse and lunar-occultation observations from the seventeenth and eighteenth centuries. But deciphering them was a nightmare. The notebooks were written in idiosyncratic shorthand, with no names, titles, or explanations. It took him six weeks just to collect the notebooks, and nearly four years to turn them into usable data. Observations that did not meet his standards were discarded immediately. The rest became the foundation of his lunar theory.

What made these observations better? They were not ancient; they were only two centuries old, lacking the benefit of a long time



FLUDD, UTRISQUE COSMI . . . HISTORIA (1617)

baseline. But Newcomb argued they were far more precise, and, crucially, they were made by astronomers, not Mediterranean chroniclers, that is, by scientists, not historians. He also presented a quantitative argument. The increased precision of modern observations, he said, cancels out the benefit of longer time baselines. As he put it: 'The secular acceleration of the moon can be determined with nearly the same order of accuracy from the modern as from the ancient observations.' The emphasis, 'nearly the same order of accuracy', was Newcomb's own. And this was not strictly true. In fact, the estimated error in the secular acceleration using ancient data was 0.4, while for Newcomb's more modern observations it was 0.6. So the quantitative precision of the modern data was only slightly worse.

But for Newcomb, precision was only a small part of what made the modern observations more reliable. What really mattered was that they came from scientific observers: people with the training and skill to qualify as part of that aristocracy ranking above all others. He even discarded observations made by Edmond Halley, yes, that Halley, who, as Newcomb put it, 'seems to have been extremely unskilled in every branch of the art of practical astronomy.' This was despite the fact that Halley had organized hundreds of observers across England to record the 1715 eclipse. Newcomb dismissed their contributions too, on the grounds that those individuals were unknown as astronomers.

It is not entirely clear from Newcomb's published work why he invested so much effort debunking the value of ancient eclipse records, even when it seemed counterproductive. The answer lies in a much

larger project to which he devoted much of his life: defining a rigorous scientific method and fighting for cultural authority based on that method. He believed this way of thinking was not natural to human beings; it had to be instilled through training. The rough and doubtful quality of ancient eclipse data, he argued, was due to the fact that the observers lacked scientific discipline. Human judgement, he said, was fallible and prone to bizarre conclusions, such as connecting the weather to lunar phases. Unfortunately for what he called ‘credulous readers’ like Airy, Newcomb claimed the early Greeks were particularly prone to wild speculation.

Scientific observation, he maintained, was hard. The testimony of a single trained observer, he said, completely outweighs that of untrained or less disciplined ones. But observation alone was not enough either: facts had to be organized under general laws of nature to produce real insight. For Newcomb, those laws, uniform, inflexible, unchanging, were the foundation of scientific reality. He wrote:

All scientific conclusions necessarily rest on the postulate that the laws of nature are absolutely unchangeable, and that their operations have never been interfered with by the action of any supernatural cause; that is, by any cause not now in operation in nature, or operating in any way different from that in which it has always done.

For Newcomb, this absolute uniformity was what enabled science's most astonishing power: prediction. It allowed not only forecasts of the future, but also reconstruction of the past. It is hard to overstate his awe at this capacity. He wrote that a sophisticated calculation could almost rival being there:

As it is, we can now calculate the circumstances of an ancient eclipse, from our modern tables of the Sun and Moon, almost as accurately as any of the ancient astronomers could have observed it.

The scientific tools of the present, he argued, allow us to understand the past more precisely than people who actually lived through it. And that, he said, was the result of two thousand years of progress in astronomy, progress that had recently extended to chemistry and biology as well. He believed it was the duty of the scientist to extend that progress even further:

Every one knows that, within the last two centuries, a method of studying the course of nature has been introduced which has been so successful in enabling us to trace the sequence of cause and effect as almost to revolutionize society. The very fact that scientific method has been so successful here leads to the belief that it might be equally successful in other departments of inquiry.

What Newcomb particularly had in mind was linking his liberal political and social views to the power of scientific method. If he could predict the orbits of asteroids, why not the movement of the free market?


He was confident that he could predict, for example, exactly how many barrels of flour the inhabitants of a city would consume. Free will was not an issue because humans, after all, were just as subject to natural laws as the stars.

He even wrote a large treatise called *The Principles of Political Economy*, laying out how this could be done. And he insisted: ‘Nothing is needed to give [economics] a recognized place among the sciences except to treat and develop it as a science.’ Of course, he added, ‘this can only be done by men trained in the work of scientific research.’ Newcomb noted that many elements of economics, such as demand, were intangible and could not be directly seen. So how could scientific method gain traction there? Well, he said, just as with studying the past, we must substitute mental sight for eyesight. The way to do that is through scientific method.

In matters of money, as in all things, he believed humans were highly prone to error unless trained otherwise. But once trained, the possibilities were astonishing. He wrote:

Were our knowledge of the whole world, including every man in it, complete in every particular, and were we able to apply all this knowledge at every moment, we might imagine ourselves to predict all economic phenomena by this method much as the astronomer predicts the motions of the planets.

Of course, our knowledge was imperfect. But the principle, he said, remained sound. In a deterministic universe governed by unchangeable laws, the past and future alike were within the grasp of science.



PETER APIAN,
ASTRONOMICUM
CAESAREUM
(1540)



And for Newcomb, a deterministic universe had other consequences, particularly for religious belief. He wrote: ‘Whenever the subject becomes so well understood that the chain of natural causes can be clearly followed, miracles and final causes cease so far as the scientific explanation of things is concerned.’ The laws of nature, in other words, left no room for the supernatural. Publicly, that was as far as Newcomb was willing to go. Privately, he believed Christianity was a dying and untenable religion. But he was reluctant to say so openly and continued to practise as a Presbyterian for social reasons. He did, however, write anonymous articles criticizing Christianity on these grounds, and was publicly vocal in his attacks on spiritualism. He held out hope that a new, humanistic, naturalistic religion might soon emerge.

So the stakes for Newcomb were quite high. He was not particularly impressed by knowledge for its own sake. He believed science was essential to citizenship and social progress. If the American people could be trained in scientific thinking, the nation could be vastly improved.

And so we find that historical eclipses and their implications have a peculiar property: they blend past, present, and future. Airy, as we have seen, wanted to use the past to predict the future. Newcomb, by contrast, wanted to use the present to predict both the future and the past. Everyone in this story agreed that what gave science its power was the ability to predict the unknown. But the question remained: Was the past known or unknown? Does it help us predict, or is it something that needs to be predicted itself? And what are the implications of being able to predict?

All of them, Airy, Newcomb, and the others, recognized the power implicit in the simple statement: 'On this day in the past, an eclipse happened.' But, as we have seen, the resolution of that statement was anything but simple.

The people involved in this story understood that it did not stand alone. It was entangled at every turn with high-stakes questions about scientific authority, about the relative status of scholarly disciplines, and about the very certainty of our knowledge. Should we use history to understand science, or use science to understand history? There was no clear answer. But the question carried profound implications regarding what we can know about the past, and what kind of authority scientists ought to have.

In the end, the truth about the past seemed to depend on the truth about science itself.



This seems like a good case study in how malleable our sense of plausibility is. Some people say the historical data is implausible because it does not match scientific expectations, while others say the science is not plausible because it does not align with historical records. Was this one of the first cases like that?

That's a good question. I think the short answer is no. It's a long-running issue that has existed since people have been doing astronomy and writing history.

As I mentioned, Newton and some of his contemporaries wrestled with exactly this problem. Newton was particularly obsessed with chronology, and he spent a lot of time trying to sort it out. He had to reconcile not just astronomical calculations and historical records, but also what he saw as the divine authority of biblical texts. So he has a third element in the mix.

You see something similar with medieval chronologers. Take Archbishop Ussher, when he came up with his famous 4004 BCE date for the origin of the world.

What's interesting, and part of why I focus on the nineteenth century, is that this is the point where science becomes professionalized. Its authority starts to separate from other scholarly disciplines. In the seventeenth century, it's not obvious that science should be given more weight than, say, history or classics or linguistics.

But by the nineteenth century, scientists have carved out enough space socially that it becomes a real fight. They can actually argue that their methods deserve priority. And by the twentieth century, and certainly today, that authority tends to dominate without much resistance.

So I think the period I'm talking about is useful because it sits in that transitional zone, where it's not yet clear who's going to win.

You mentioned their reliance on Greek sources, but did figures like Newcomb or Airy consider Chinese or other non-Western astronomical records?

They do, but not during the period I've just been describing. That comes later. And the reason is, frankly, ignorance. European astronomers were not really aware that non-Western sources

existed, or could be relevant, until the late nineteenth century.

Which is surprising, because when the Jesuits began working at the Chinese court in the late sixteenth and seventeenth centuries, they recognized very quickly that Chinese astronomy was highly sophisticated. But the idea that Chinese historical records might be valuable does not take hold for quite a while.

Simon Schaffer has written some terrific work on this subject, exploring how British astronomers start building observatories in India in the mid-to-late nineteenth century. They are confronted with a very long and detailed Indian astronomical tradition, and they have to decide: do we take these seriously? Often, they don't. It becomes a kind of intellectual imperialism, an effort to justify why these non-Western records can be dismissed or ignored.

There is also a literal aspect to this discovery process. For instance, Europeans only learn to read cuneiform in the nineteenth century. So Babylonian and Assyrian astronomical texts effectively come into being for European scholars at that moment, and suddenly have to be dealt with.

It's really not until the twentieth century that Western astronomers begin to take non-Western sources seriously in a systematic way.

Did we ever stop relying on historical records to determine secular acceleration, or are they still used?

We still use them. But the centre of gravity has shifted.

Today, the Moon's tidal acceleration is measured very precisely using lunar laser ranging. Historical records are still

valuable, especially for reconstructing long-term changes in the Earth's rotation, things like ΔT , because those changes affect the timing and apparent location of ancient eclipses.

And, yes, there's a kind of modern consensus that the Greek eclipse reports are of limited use for this, compared with other sources. The most valuable long-run records tend to come from places like Babylon and China. But it's still debated. People will argue about the details, and you can get a good argument going at an astronomy conference by asking what people think the current value should be. Depending on how you define it, the standard value used in eclipse calculations corresponds to something like -26 arcseconds per century squared in the Moon's ecliptic longitude.

FLUDD, UTRISQUE COSMI . . . HISTORIA (1617)



And in the nineteenth century, during the period you described, how accepted was the idea that the Moon had a secular acceleration? I know the idea goes back to the late seventeenth century, and I think the first numerical calculation came in the eighteenth. How widespread was it by then?

I think it's safe to say that by the time you get to Airy, it's treated as a real effect, what they would have called a non-trivial value. And by the time you get to Newcomb, in the 1870s, it's generally accepted that secular acceleration is real. The question becomes: what, exactly, is the value?

And did that have any theological implications? I mean, the fact that the Moon's motion is not stable, was there any kind of religious reaction to that?

Not specifically about the instability of the Moon's motion. But this is the second half of the nineteenth century, when cosmologies that suggest the universe is running down or has a finite age start to become popular.

So I'm thinking here of things like the heat death of the universe, or William Thomson's work on entropy.

And, interestingly, a lot of religious figures embraced those ideas. Thomson himself was quite excited by the notion of a heat death, since it implied the universe had a beginning, and so the Christian vision of an apocalypse must, in some sense, be true.

Newcomb took a very different view. His reading was strongly anti-theological. For him, the correct lesson was that the universe is entirely governed by natural law, and that there is no room for supernatural intervention.

And Thomas Henry Huxley, over in the UK, made similar moves. Both of them used these lunar theories to argue that divine intervention in nature simply could not be a thing.

Yes, this is a funny bit of linguistic play he does. When we say ‘prediction,’ we usually mean using equations to figure out what’s going to happen in the future. We might calculate when the next solar eclipse will occur, for instance.

But because the laws of astronomy are symmetrical in time, it’s exactly the same operation to calculate when an eclipse would have happened in the past. The maths doesn’t care.

But we don’t usually call that ‘prediction,’ right? So it raises this odd question: what do we call that?

Huxley says, well, it’s still prophecy. We usually think of prophecy as foretelling the future, but in this case we are using the same mechanism, science, to pinpoint events in the past. So he calls it retrospective prophecy. We are using scientific tools to figure out what happened, but we are doing it with the past, not the future.

And, of course, what Huxley is really interested in here is evolution and how human beings emerged through natural processes. He wants to say: look, Newcomb has shown us that science can make claims about past events with the same kind of confidence it has about future ones. So geologists and biologists should be able to do something similar.

He has these playful, mythic sketches of what it would look like to ‘predict’ the emergence of life in the distant past.

It’s also a great example of something Huxley excelled at: taking religious categories like ‘prophecy’ and reworking them into naturalistic, scientific ones.

Speaking of Huxley, your paper on this subject mentions his idea of ‘retrospective prophecy.’ Could you explain that for people who have not read it?

You mentioned disciplinary tensions earlier: was someone like Newcomb sceptical of historical data because he thought history as a discipline was less rigorous, or was it more about the specific sources, like Herodotus, which are known to include things that are not reliable?

As you say, a lot of historians in this period are sceptical of Herodotus because these texts are manifestly not self-consistent. And yet one of the funny things is that someone like Airy, writing at a time when that is already well understood, still wants to take Herodotus seriously.

Part of what helps us make sense of that is thinking about the different cultural locations at play, which I only touched on briefly. In Victorian England, if you went to Cambridge to study astronomy and mathematics, you actually spent most of your time studying classics. So it's not 'history' in the modern professional sense, but rather classics. These texts are to be studied and taken

seriously as a body, as part of your education.

Whereas in nineteenth-century North America, Newcomb starts out in Canada, nobody is studying classics. It is not a well-respected field at all. If you call yourself a historian in 1870s America, you are probably studying the European Renaissance. That is 'modern' history.

So part of what is going on here is a kind of slippage between disciplines. People like Airy, when they say they are doing historical work, really mean they are doing classics, specifically Greek and Roman texts. Those are the ones worth taking seriously. And the idea that there might be other texts worth consulting is almost beside the point.

I guess the kind of confidence required to dismiss historical data would only emerge once astronomy had reached a certain level of precision. By Newton's time, and even more so later, there were still plenty of things the equations could not predict well, like Mercury's anomalous perihelion advance.

Yes, exactly. By the mid-nineteenth century, people like Airy and Newcomb feel very triumphant about the success of astronomy.

They would point to things like the discovery of Neptune as extraordinary proof of how precise astronomical calculation could be. Newcomb, in particular, made his reputation by predicting the motions of asteroids, which were still pretty new objects at the time. The fact that he could say where they would be based on fragmentary observations, that was a big deal.

There has even been a long-standing suggestion that Newcomb helped inspire Professor Moriarty. Moriarty, as you know, is a mathematician, but he is supposed to have written a treatise on *The Dynamics of an Asteroid*. The point is that this kind of astronomical thinking was seen as the very height of intellectual sophistication, so much so that if you wanted a mind who could rival Holmes, you might give him celestial mechanics.



Was there any similar astronomical work on secular acceleration in other bodies, like the Earth itself gradually drifting, or other planets?

That's a good question. The short answer is: generally not.

The Moon is a very strange and difficult object. It does have a measurable secular acceleration, along with other weird features. People have looked for analogous effects elsewhere, and technically every object in the Solar System has some element of secular change, but it tends to be very, very small.

Where it does come up in a very prominent way is Mercury. Mercury has a strong perihelion precession. Technically it's not called secular acceleration, but it's a similar kind of anomaly.

And the general move in nineteenth-century astronomy, especially after the discovery of

Neptune, was that when you find an anomaly in the motion of a planet, you assume it is caused by some as-yet-unknown body. That's how Neptune was found: Uranus had an unexplained wobble, so you work out where the invisible planet must be, and then you look for it, and there it is.

So when Mercury shows a strange precession, the same logic is applied. Astronomers start looking for a planet even closer to the Sun, and they call it Vulcan. For decades, people search for this hypothetical planet.

What ends up happening, of course, is that Einstein explains Mercury's anomalous precession using general relativity, in 1915. So that's one of those fascinating cases where the effort to explain an anomaly leads not to a new planet, but to a whole new theory of gravity.

HARVARD OBSERVATORY TEAM, SOLAR ECLIPSE (7.8.1869)



With the Moon, there's a similar hope that you might be able to track the anomaly by linking it to the motion of other Solar System bodies. But that's where the calculations become really intense, as you are dealing with systems of equations with hundreds of parameters, and solving that by hand is incredibly difficult.

And one of the reasons Newcomb impressed everyone so much was that he could actually do that.

This might be a technical question, but let's say you're tracking the margin of error between a current model of the Moon's motion and historical records. If, over time, that margin of error increases in a regular, lin-

ear way, fine. But what if it starts growing exponentially? Would that suggest something is wrong with the data or the theory? Or would you just accept that the margin of error could grow exponentially?

Yeah, I mean, this is a deep question about theory and observation more generally: what do you do when there's a mismatch between what your theory says and what the data show?

In this particular case, the expectation is that the anomalies should grow with the square of the time. So if you see that, if your anomalies grow quadratically, you'd say, okay, this is a real effect. That matches the theoretical expectations.

But if you saw exponential growth, then you'd probably say you have one of two problems. Either there's something wrong with your calculations (maybe you've chosen the wrong value for the secular

acceleration, or structured your equation in the wrong way), or there's something wrong with the observations you're using to calibrate your equations in the first place.

And here's the tricky bit: mathematically, you can't tell which one it is. You can't distinguish between an error in the theory and an error in the data based on the maths alone.

That's what I find really interesting about this story. You immediately have to bring in other sources of epistemo-

logical judgement to figure out where the problem lies.

So in the nineteenth century, that meant deciding whether you could trust the classical authorities. These days, as we discussed earlier, the consensus is that Chinese observers are more reliable than Greek ones. But arriving at that conclusion, deciding that one culture is better at astronomical observation than another, is not a purely scientific decision. It's a very tricky thing.

I imagine there's a kind of spillover effect too. Like, if one set of historical observations fits a certain tradition. You get one eclipse right, and next thing you know, people are out hunting for more.

Very much so. And I think, as you're suggesting, if you're a classicist, then these texts hang together as a body. The idea that you would accept Herodotus' weird stories but not Plato's weird stories doesn't really work. You kind of have to take them all as a whole.

Probably the most notorious example of this is Immanuel Velikovsky's theories from the 1950s. I don't know, do people still know that name? This might be totally obscure now.

So: Velikovsky was a psychoanalyst in the '40s and '50s who became convinced that all historical myths are true, that every story in the Christian Bible, in

Chinese tradition, in the *Iliad*: they all describe real events that actually happened. His goal was to line up all of these accounts with astronomical events.

He wrote this wild book called *Worlds in Collision*. One of his claims is that the planet Venus was launched out of Jupiter, flew past Earth, and part of its atmosphere mixed with ours. And that's what the Jews wandering in the desert saw as manna. And we know this is true, he says, because scientists have found hydrocarbons in Venus's atmosphere.

It's completely over the top, but also incredibly detailed. And he insists on starting from the

tions turns out to be correct, people might start trusting other parts of the same thing for Atlantis. That must happen all over the world, right?

idea that all of these texts are accurate historical sources. So for him, Atlantis is destroyed by earthquakes caused by Venus's near approach. It's a whole cosmology.

And the remarkable thing is that everybody read him. Everyone had a copy of *Worlds in Collision* on their bookshelf. This was also the era of Erich von Däniken and ancient aliens and so on.

Carl Sagan gained some early public prominence attacking Velikovsky. That was one of his first big moments in the spotlight: debunking these theories.

So it became this big cultural clash, Velikovsky saying, 'Trust

me, I have all these historical texts on my side,' and Sagan saying, 'That may be, but science is better.' Astronomy tells us the story can't be true, so even if the historians are enthusiastic, you have to listen to the astronomers.

If you want to know more about that whole episode, a friend of mine, Michael Gordin, wrote a wonderful book called *The Pseudoscience Wars*. It's about these debates in Cold War America between scientists and people who thought of themselves as historians, showing how the whole category of 'pseudoscience' emerged from that conflict. Really fascinating stuff.

PETER APIAN,
ASTRONOMICUM
CAESAREUM
(1540)



This is a two-parter. First, when we talk about history, especially through public education, there's also this rise of infotainment from the era of television onwards. You get public science figures with a face and personality, such as Carl Sagan with *Cosmos*, or Neil deGrasse Tyson now, who have these cult followings.

And that creates this situation where what they say, even if it's influenced by their personal or political motivations, gets treated as scientifically valid simply because they're scientists. That's different from, say, an article in *Scientific American*, as these figures have enormous control over how science is understood by the public.

Second, and this is just a factual question: how was Carl Sagan's career affected by the whole nuclear winter controversy?

Those are both good questions, and I'd say they're actually pretty intertwined.

As you said, figures like Sagan and Tyson have been able to use the medium of television, and now video more broadly, including YouTube, to build this kind of cult of personality. Simply presenting something as 'coming from a scientist' gives it a kind of immediate authority. It's taken seriously because of who says it, not necessarily because people understand what's being said.

And that's a bit ironic in Sagan's case, because his whole mission with *Cosmos* and his other programmes was to educate people on how science works and explain why it's reliable, what makes it powerful. He didn't want to be seen as just a guru

dispensing truths. But he was so successful as a television personality that people started treating him exactly that way. What he said was taken as gospel, without much critical thought.

Now, with nuclear winter, just for context, in the early 1980s, Sagan and some other scientists argued that even a so-called 'limited' nuclear war would cause massive climate disruption: global famine, ecological collapse, that sort of thing.

Sagan came forward and said, basically, you shouldn't be listening to the politicians on this; listen to the scientists. He argued that scientists, because

they understood the underlying physical processes, should be the ones guiding nuclear policy.

And that didn't go well for him. He tried to leverage his cultural authority as a scientist into political authority, and it backfired. People turned around and said, 'Wait, you're not a nuclear strategist, you're a TV guy. Why should we listen to you?'

So in a sense, the same public platform that gave him his power also undermined him. The thing he thought would make people pay attention became the reason they dismissed him. It's a complicated legacy.



Huxley in Victorian England and Newcomb in industrializing America are working in very different religious and institutional contexts. In England, Huxley is dealing with an established

In your article, you showed how Huxley drew on Newcomb's scientific determinism to support Darwinian evolution, which ultimately became part of a broader contest with the Church over ideas of creation.

But the relationship between science and religion at the time was complex. For instance, John Heilbron has shown how, in the seventeenth and eighteenth centuries and even into the nineteenth, Catholic churches were used as solar observatories, or as sites for sound experiments. And the reception of Darwinian evolution varied across regions, with Calvinist areas like Belfast showing stronger resistance than places like Princeton.

So in that context, how did the Church respond to Newcomb's brand of scientific determinism?

church, the Anglican Church, which still controls many of the country's educational institutions. That's a big part of the fight for him. He wants to use deterministic science, like Newcomb's, and evolutionary science, like Darwin's, as a kind of bludgeon to argue that science needs its own independent space, ideally in secular universities where it can be taught without theological interference.

So that's one of the main targets of Huxley's campaign: the institutional power of the Church of England.


In America, the situation is different. There's no established church, so Newcomb doesn't face the same kind of centralized resistance. But most American universities at the time are still affiliated with religious denominations. They're just not unified, which means there's less political authority behind any one theological position.

JOSEPH WALKER,
SOLAR ECLIPSE, LONDON
(14.7.1748)

Instead, Newcomb faces a patchwork of local sectarian critiques, from Methodists, from Baptists, and so on. It's a more fragmented landscape.

That actually gives him more room to manoeuvre. He aligns himself with figures like Andrew Dickson White, the founding president of Cornell, which was one of the first explicitly secular universities in the U.S. And that allows Newcomb to help carve out entire regions, like the Northeast, where secular education, grounded in ideas like his, becomes dominant and widely accepted.

So in a way, Newcomb is able to succeed more quickly than Huxley. This is not because the opposition is weaker, but because it's less centralized. Huxley had to fight powerful, unified institutions. Newcomb dealt with scattered resistance, which made it easier to build up secular scientific authority in specific places.



Are there parallels between the debate over the Moon's secular acceleration and the work happening in the late eighteenth and nineteenth centuries on geological deep time, particularly regarding the debates around uniformitarianism?

Yes, definitely. So uniformitarianism, or, more broadly, the idea of uniformity, is really central here. It's the idea that the laws of nature stay the same over time, and that you can use the present to infer the past, and even the future.

As you say, that's crucial for geology in the late eighteenth and early nineteenth centuries. The standard framing at the time was: you should only invoke processes in the past that you can observe happening now.

So the classic example is that you can't invoke a global flood to explain the shape of mountain ranges, because we don't see glo-

bal floods today. But we do see volcanic eruptions. So if you see certain rock formations and you know how volcanoes work, then you're allowed to say, okay, volcanoes caused that.

That's uniformitarian reasoning.

Now, when you apply that to astronomy, it gets trickier. With something like the secular acceleration of the Moon, you're dealing with a motion we can measure now. It orbits about every 28 days, but it doesn't come back to exactly the same spot. So you ask: what kind of variation counts as small enough to project backwards in time?

The problem, as we discussed earlier, is that the secular acceleration shows up in equations as a square function of time. So small changes now become large changes in the distant past, and you have to decide

how far you're willing to extend that logic.

At what point does it stop being the 'same' Moon? As you go further back, the Moon would have been closer to the Earth, and it would have looked much larger in the sky. So is it scientifically valid to talk about that? Or does that violate the idea of uniformity?

The trick here is that you accept a certain level of uniformity, and then use that to find anomalies, departures from your model. These anomalies, ideally, help you discover a new kind of uniformity, a more subtle one. And that gives rise to new anomalies, and then you go looking for the next layer.

In geology, the classic example would be continental drift. Early geologists couldn't explain certain geological features, and when someone suggested that the continents had

moved, the objection was: well, we don't see them move now, so that violates uniformity.

But eventually we did find a small-scale present-day process: seafloor spreading, with those magnetic bands that you can observe. And that gave geologists permission to extend the idea into the past and explain something much bigger, like the shape of the continents.

So you get this interesting iterative process. You start with one level of uniformity, which allows you to find an anomaly, which then leads to a new layer of uniformity to explain that, and so on.

The secular acceleration is kind of two steps deep in that chain. But presumably the process goes on forever. You keep finding anomalies, and then you need new layers of regularity to explain them.

EUGÈNE ATGET, PENDANT L'ÉCLIPSE (1912)



Even the idea of uniformity seems a bit flexible. If something changes in a regular, predictable way, that regularity itself can be treated as a kind of uniformity, just with structure.

That's right. Huxley had a metaphor he liked for this. He'd say even something that looks non-uniform can be explained in uniform terms pretty easily.

So his example was: a clock is a perfectly uniform device; it runs on uniform laws. But the striking mechanism of the clock could be made to blow up a barrel of gunpowder, or turn on a deluge of water. In that sense, the explosion is actually a uniform event. It only looks irregular if you don't know what it's connected to.

And Huxley used that to make the case that even the most surprising or chaotic events should be explained in uniform terms. That's the scientific ideal.

But the philosophical problem lurking behind that is: why assume uniformity in the first place? What justifies treating the universe that way?

It turns out that's actually quite hard to explain. The traditional answer was theological. Newton, for example, thought the universe followed uniform natural laws because that's how God set it up. So you could rely on them because God wanted a stable, orderly creation.

But by the time you get to Huxley and Newcomb, that theological option isn't available anymore, and it becomes much harder to justify uniformity on purely philosophical or axiomatic grounds.

One of the reasons to bring up geology and deep time is because Halley, who first posited secular acceleration, was actually accused of believing the world was eternal. Part of his motivation was to prove that it wasn't.

Yes, that's right. Halley's a really interesting figure in that regard.

And just to broaden things a bit (although I've been focusing on astronomy, since that's my area), this kind of issue still matters a lot today. You see it, for instance, in climate science.

Climate scientists rely on historical records all the time, trying to reconstruct changes over the past few centuries. But they run into the same kinds of questions. If a seventeenth-century farmer wrote down the temperature every day, how do we know how reliable that is? Can we treat it as scientific data?

It's the same basic epistemological issue that Halley and his contemporaries were facing. And now, arguably, the stakes are even higher.

Staying with astronomy, do you think the field has a certain quality that makes it especially prone to these kinds of disciplinary clashes? You mentioned this almost hostile relationship with history, and this makes one think of archaeology. There, you have myths like El Dorado, and people try to verify them with archaeological evidence. Sometimes they do, sometimes they don't, but that doesn't seem to provoke the same kind of disciplinary tension.

Yeah, I think you're right, and I find that a helpful comparison. In archaeology and history, the boundary is already a bit blurry. No one's particularly surprised when archaeologists and historians have to sort things out together.

But astronomy is different. Astronomers often see themselves as doing something far removed from human concerns. They see their work as purely mathematical, objective, and untouched by interpretation or social context.

Of course, that's not actually true. But it's the way many astronomers think about what they do. So when they're suddenly forced to engage with historians, when they realize they need to rely on historical records, it's a bit of a shock. It feels like a breach of disciplinary purity.

So, yeah, the conflict is sharper in astronomy precisely because of how the field has conceptualized itself. That's what makes it such an interesting case . . .



CARD DIVISION, LIBRARY OF CONGRESS (c.1910)



Randolph Head

**Archives and the
Information State**

How is the relationship between documents, or records, and the world structured? And how does it change?

The answer, of course, is that institutions and practices are where that relationship is modulated, where it evolves. Coming at it from that perspective, although I didn't realize it when I began my research in the late 1990s, aligned closely with developments in archival theory in the same period, which aimed to destabilize the idea that the archive is simply where dead documents go to sit, locked away.

Instead, the relationship between documents and the world is dynamic, and subject to historical change. I don't think there are any non-information states, but the kind of information state that emerges in this period is distinctive.

How did records from the past become information for the future? And how was that different from the work that records had done in the late and high Middle Ages in Western Europe?

I'll confine myself to Western Europe, which is what I know, but I'm quite aware that commensurable, or perhaps incommensurable but parallel, systems existed in many parts of the world, especially in regions with agrarian imperial structures.

Contemporary archival theory argues that even without any change to the textual objects themselves, charters, books, bundles of papers, whatever, the meanings and uses of records can change profoundly.

You can get entirely different facts from the same records, depending on how they're interpreted. And those changes are enabled and reflected by new organizational, institutional, and agent contexts.

And the word that I borrow from archival theory as the central concern is context.

Now, context is a very difficult and unbounded term, but it's how archival theorists think about, and try to access, the ways in which records make meaning.

This is going to be a very schematic and quick rundown, but take it for what it's worth. What was a fact in the premodern European dispensation? We could call it medieval, or perhaps late medieval.

There was a very durable definition that persisted long after it had ceased to be true in practice. A fact was what a document conveyed. The work a document did was to provide authentic memory of authoritative action.

It had to be authenticated, or at least authenticatable. The action represented was assumed to be the issuance of a privilege from above to below, and had to be performed by a proper authority. And it had to be recorded for eternal memory, the classic Latin phrase, in some material form that could carry it forward in time.

Now, in actual practice, documents were used in all sorts of ways that don't fit this definition. Medieval historians have done wonderful work in the last generation showing this. These documents were used ceremonially, they were sometimes destroyed, they circulated in ways that don't match the idealized definition.

Still, that notion that a charter was an authentic memory of authoritative action was very strong. It remained prominent not just in the later Middle Ages but among archivists and historians ever since.

Meanwhile, for various reasons, changes in media systems, shifts in political economy, population growth, and other transformations, the contexts in which such records operated changed quite significantly between 1400 and 1700, the range covered in my book.

What happened was a shift: the salience of individual documents and their embedded meanings diminished; their value as part of a larger set increased. These records were detached from their local, immediate meanings and became abstract resources, data, if you like, that existed in relation to other similar resources.

That collocation of authoritative memories is, in a way, the root of the idea of the archive. The archive becomes the place where those memories are gathered and held together. And, in doing so, they create a new context.

Again, that's a key point in archival theory: the archive itself becomes the context of the documents, not the king's breakfast on the day the charter was issued, not the specific political backdrop, not even who signed the document.

All those historical details are, in a sense, wiped away once the document enters the archive and comes into relation with other documents. And those new relationships are productive. They generate new practices, new personnel, new arrangements, new institutions. That's the emerging archival complex of the early modern period.

This transition was slow. One aspect of it was a very slow, late shift from treating documents as memories of particular acts to creating ways of aggregating them, to find facts in the aggregate, not just in any individual document.

This also marked a shift in the key logic of organization. And I'm speaking here at a very abstract level, but my argument is that medieval and most early modern archives understood documents as being about the world.

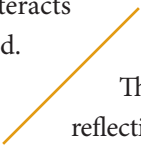
They were about actions, those of rulers, subordinates, subalterns, and those actions took place in the world. So the organization of records focused on the world's structure. The goal of archival management was to find the document that related to a particular site or situation in the world.

NEW YORK PUBLIC LIBRARY (c.1910)



But in the course of their operation, institutional complexes, archives, registries, chanceries, gradually developed in ways where documents became more about other documents, or about the operations of the institution itself.

And in that transition, we begin to see the emergence of what we now call a modern state, what I call an information state. Records become the way the state interacts with itself and with the world.



The organizational models reflecting this are profoundly different from those of the late Middle Ages or the early modern beginnings. In this new model, and here I'll quote myself, 'where the world informed rulers, information became the raw material that entered the registry.'

Information in a registry becomes a substance to be processed by the state.

Through this transformation, moreover, the state becomes intelligible, legible, and capable of action. The outside world still matters, but the relationship between document, institution, and world changes significantly.

That's the core of my theory of records and information states. The essay I wrote for *Information: A Historical Companion* is really where I hypothesized that. And it is a hypothesis, it certainly hasn't been broadly confirmed, but for me, it's a way of thinking about what distinguishes modern from premodern documentary practice in Europe.

In the nineteenth century, the consequence of this transformation was the emergence of something entirely new: archival theory itself. For the first time, really explicit archival theory, built around the principles of provenance and *respect des fonds*.

The rule now was: documents must be kept in the order and context of the institution that produced them. That had happened passively before, but now it became the guiding organizational principle. Institutional bodies of documents had to be kept intact.

And what you see in that is an implicit argument: the context of any document is the other documents around it, not the external world to which it relates.

Now let us shift to a second strand, one that seems relevant to the broader question of European archival practice in a larger world: a world shaped by colonial empires, by resistance to those empires, and by the eventual decentring of European ways of doing things.

Let me begin by returning briefly to the medieval system I sketched earlier: charters of privileges being created and then decontextualized into archives, into treasuries of facts. The key here is decontextualization.

There's some systems theory behind this. I came to it primarily through the work of the Swiss historian Thomas Hildbrand, and also Eric Ketelaar, a central figure, both an archivist and an archival theorist. The idea is that whatever the context of a document at the moment of its creation, and that context is rich, nuanced, the sort of thing historians want to know, once the document enters the archive, most of that context is stripped away.

And what struck me is that this stripping away is actually what makes the document useful. That's a point I first encountered in Geoffrey Koziol's work. Let's say the king gave a particular monastery to a bishop. At the time, that act was embedded in intense local politics. But after a hundred years, almost all of that politics is gone, no longer visible in the details or the silences of the document.

What remains is a simple fact: who owns the monastery. And that's what the charter documents.

Charters are thus very peculiar, highly formalized documents. In classical diplomatics, they're designed to be self-contained. They declare what they are, and in doing so, they also declare what they are not. They leave out the world to concentrate on the fact they aim to convey.

I once tried to think about this model in relation to early European states trying to cope with the new worlds they were encountering,

worlds with their own rich documentary traditions; or large Indigenous populations who managed knowledge without written documents; or whose documentary systems, such as those of the Maya or the Aztec, were destroyed.

How do you cope with that?

And it struck me that the system I just described, where you strip away context and are left with facts, only functions when the sub-structure, the infrastructure of knowledge, is already in place: when the organization of space, the categories of people and places, are more or less understood.

That doesn't mean uncontested. It's always contested, but it's known. And when European systems tried to apply this infrastructure in other parts of the world, they often failed, because that infrastructure wasn't there.

The classic case is the early Spanish Empire. Apparently, in the first generation, Queen Isabella and her advisers genuinely believed they could assimilate the populations they were conquering, especially the Aztec, later the Inca, into Castilian peasantry.

But it didn't work.

They had to radically rethink, several times, not just the kinds of documents they produced, but the institutions that could transmit knowledge, from the periphery to the centre (from their perspective), because they simply couldn't transfer the existing system into these new conditions.



So we begin to see that these processes of subtracting data are what made the European model of the document, the charter, the act, the deed, possible.

At the level of the individual record, it means leaving out the context of production.

At the level of corpora, bodies of records, it means shaping them through the context of use, which changes over time.

In my research for *Making Archives*, that's what I looked for: moments of reorganization. Instances when royal, urban, or ecclesiastical power was confronted with a body of documents that no longer worked, that couldn't be used for the institution's current purposes.

And then you get projects, sometimes small, sometimes ambitious, to reorganize.

It's striking that European archives have preserved so much of the detritus of these efforts. In almost any Western European archive, if you look at documents from the sixteenth or seventeenth century, you'll find hundreds of lists: incomplete, fragmented, unsuccessful.

And then you find a much smaller number of large-scale reorganizations, projects that were carried through, or mostly carried through, but often left unfinished.

Those are the most revealing, because by examining the institutional frame and the practices they embody, you can see what their authors thought records could do, what they should do, and how they might be enabled to do it.

It seems that people have always relied on information and made important decisions of state based on certain considerations, even if they didn't articulate their decision-making processes in objective terms. So how would you define an information state? How is it different from a state that simply makes informed decisions? And what role does a literate political class play in this?

I don't think there are any state formations that don't rely on information. So the question is really about different modes of capturing, managing, and then mobilizing that information.

As my colleague Arndt Brendecke in Munich often points out, most political decision-making


happens outside archives. Formal records aren't the central arena for the kinds of processes you're describing.

So what, then, is an information state?

I think we start using the term for states that self-consciously seek to mobilize information, where records from the past are turned into resources for the future. That's why we wouldn't normally describe medieval monarchies as information states.

They certainly had records, mostly charters, sometimes letters or reports, but they understood those as a treasure, literally. The French crown, once it built its beautiful chapel for the king's relics in the Sainte-Chapelle, kept the archive next to the relics and the treasury.

So it was a treasury of documents, or, as a German archivist once put it, an arsenal, from which one could draw documents like weapons. That's using information for political ends, but it's not really informational, in the sense I'm talking about.





The information being deployed in this older dispensation is highly specific. It's about things in the world: privileges, rights, ownership.

What we begin to see in the fifteenth century, especially with Spain, and Philip II is a key figure here, is a shift. Spain has an empire. Philip II, who ruled from 1556 to 1598, oversaw an extraordinarily large and heterogeneous polity, with possessions in Europe and overseas, stretching from the Low Countries and parts of Italy to Peru and, after 1580, Portuguese territories as far away as Asia, including Macau.

In that context, it became less and less useful simply to have proofs of privilege such as grants, exemptions, legal titles. What became more important was knowing what was out there, to make informed decisions.

That tension is part of what the title of my book, *Making Archives in Early Modern Europe: Proof, Information, and Political Record-Keeping, 1400–1700*, is pointing to. Documents that had evolved to provide proof of very specific things, who owns this monastery, how much tax this mill pays, what court rights a city holds, begin to be retooled to ask broader questions: what do we know about the world? What do we need to know? What do we not know?

One important bridge in this shift is, somewhat surprisingly, the Inquisition, especially the thirteenth-century campaign against the Cathars. The Inquisition took the most advanced intellectual tools of its time, scholastic and university-based methods, and applied them to tracking and eliminating heresy.

To do that, you needed to gather, preserve, and organize information systematically. Of course, a trial might eventually require proof, but what was new was the infrastructure the inquisitors built to track and accumulate information over time.

Those methods eventually transferred to secular states.

Another major site of this transformation was northern Italy in the fifteenth century, particularly in diplomacy. Previously oral exchanges were now recorded on paper for future reference. Venice was a forerunner here, with systematic reports from its ambassadors, clearly informational in purpose.

There was even a secondary market for these ambassadorial reports. They were supposed to be secret, but you could buy them. So, increasingly, states sought to mobilize documents to inform themselves about the world, rather than simply to contest it.

That's the shift to an information state.

A state becomes an information state when it reorients its use of documents from serving as mere links to facts in the outside world, to tools for managing the challenges of daily governance and long-term policy. The circulation of documents becomes primarily internal to the state apparatus.



Richard Firth Green describes how, in the medieval period, performances or symbolic objects could ratify formal agreements. For example, when property was transferred, a branch from the land might serve as the contract, the proof, the charter, so to speak. Some of these practices persist today; when we marry, we exchange rings.

In your paper, you mention that beyond their informational content, documents like charters were also used performatively. Could you give a few examples of what that looked like?

My examples will, of course, go back to Swiss history. I'll cite a couple of dissertations, mostly supervised by my late colleague Roger Sablonier, who was very interested in non-probative documents, what he called economic documents.

Let me tell a few archive stories, if I may.

One of them I got from Simon Teuscher in Zurich. In the early fifteenth century, the city of Zurich was collecting feudal dues

from its subject villages. The city was the sovereign; the surrounding villages were its subjects in a feudal relationship.

A bailiff would go out to these villages carrying a rent roll, a piece of parchment listing, often quite illegibly, the dues owed. The villagers, being Swiss and somewhat assertive, would say, 'Show us where on this roll it says we owe this amount.'

And the bailiff replied, 'Do you think I can read this?'

The villagers scratched their heads and said, ‘Right, fair enough.’

That’s one example.

Teuscher gives another, which really captures the transition I was talking about earlier. There’s a detailed diary from a Zurich bailiff working in the countryside. In one entry, he visits a particular village (sadly I forget the name) and asks for a copy of their book of rules.

By the fifteenth century, many villages had such written books recording things like boundaries, conflict resolution procedures, and rules of internal governance. These were semi-literate communities, just literate enough to begin writing things down to transmit them through time.

The bailiff says, ‘I’d like a copy.’

The villagers reply, ‘We’ll bring it to you, and if you ever need to use it in a legal case or a dispute, we’ll come along to explain it.’

So the meaning of the book, in the peasants’ view, depends on its community of interpretation, that’s Brian Stock’s phrase.

Eventually, the bailiff does receive a copy. And what does he do with it? He puts it on a shelf with books from other villages.

At that point, the document’s context shifts from being embedded in a local community to being just one among many village law codes. Now it can be compared: this village does X, that village does Y.

The modern historian, moreover, encountering these documents sitting side by side in the archive, sees them in relation to each other. And that’s the key point from archival theory: the meaning of a document is always

shaped, indeed, constructed, by the context in which it is presented.

And contexts are constantly changing.

Now, on the ceremonial side, there's another lovely study on annual oath-taking in Swiss towns. Citizens, and I should stress, citizens were a privileged class, not just anyone living in the town, would take an oath of obedience to the city's authority once a year.

In Lucerne, a foundational document known as the *Geschworener Brief* (first issued in 1252) would be read aloud, and then the citizens would swear an oath to obey it.

But it turns out that the city council edited the text over time. So even though they presented it as this unchanging foundational truth, in reality the document was updated to reflect current needs.

What people swore to, in a sense, was the object itself, not necessarily the precise words. That's a ceremonial use.

And yes, it's ubiquitous. Formal records were used performatively in the medieval dispensation. That never entirely disappears, as you mentioned with wedding rings, but it fades, or rather, it becomes accompanied by other forms.

One major shift was the rise of formal adjudication. This is a legal history point I'm sure Richard made: courts transitioned from being sites of community witnessing and oral adjudication to institutions that operated primarily through written evidence.


England is actually an outlier here. Its courts remained more performative. But on the Continent, by 1500, most courts had adopted fully written procedures.

NATIONAL ARCHIVES, DC (1937)

The parties didn't even appear in person; they submitted briefs. Much of the process happened in writing.

And that shift, in which the community is subtracted and the institution is added, profoundly changes the context in which documents are used.

In a community, you get context. In an archive, you get structure.



We've been discussing the asymmetry between how historians approach archives and the status of archival studies as an autonomous discipline. An article we read pointed out that historians tend to speak of 'the archive' in the singular, while archivists talk about 'archives', in the plural and with attention to contextuality, both of the archive and of individual documents.

This also invokes Quentin Skinner's essay 'Meaning and Understanding in the History of Ideas,' where he identifies two broad approaches in intellectual history: one that treats texts as autonomous, self-contained units, and another that sees meaning as overwhelmingly determined by context.

Do you think there's a parallel here with the divide between historians and archivists?

One can't help but also think of the distinction between literary critics and textual editors. People like Honigmann, who worked on Shakespeare's folios in the pre-digital era, did enormous scholarly labour, yet their status still seems secondary to that of critics. Is there a similar hierarchy between historians and archivists?

Yes, Skinner and Pocock were very influential for me when I was working on my dissertation, which looked at political culture in a largely non-literate society, specifically the Canton of Grisons in Switzerland. I was interested in how people struggled to constrain elite power using pamphlets, but also through action, and how to connect those elements.

The text/context issue is indeed a kind of predicament. It's not binary, but more of a spectrum.

What Skinner and Pocock critiqued was the older model of intellectual history, what you might call the 'great chain of texts,' where Polybius wrote, then Machiavelli, then Hobbes, and so on. In that model, each text is treated as autonomous, and if there's a context, it's just the other canonical thinkers.

This approach excluded countless voices, outsiders, subaltern groups, women, and built an intellectual tradition defined by a narrow chain of brilliant men. So the critique of that model is well justified.

The turn to context, and to seeing texts as part of broader discursive or rhetorical formations, opened up space for more inclusive and nuanced readings. And eventually, yes, we reach the more postmodern approaches to text and meaning.

As for the singular 'archive' versus plural 'archives,' that's quite telling. These terms align with the broader differences in disciplinary orientation.

Interestingly, though, context isn't just a focus of contemporary archival theory; it was central to classical archival theory as well, which developed in the nineteenth century.

At the very time Leopold von Ranke was elevating administrative documents as sources of ‘what really happened’, *wie es eigentlich gewesen*, archivists were insisting that documents weren’t just parcels of truth to be mixed and matched at will.

Let me give you an example that illustrates what’s at stake.

Zurich, as a political entity, developed an extensive record-keeping system from the fifteenth century onwards. It accumulated hundreds of boxes of dossiers related to conflicts, bundles of paper recording entire disputes.

But in the nineteenth century, as the archive was institutionalized under the modern Zurich state, all those bundles were taken apart.

Documents were instead filed individually into broad chronological series. So, if a 1750 boundary dispute involved a

document from 1250, that earlier record was pulled from the bundle and filed under 1250.

In doing so, they treated each document as self-contained, as if it encapsulated a single fact, detached from its web of meaning.

And they destroyed the archive.

It’s now much harder to use what had been a rich and coherent body of evidence. A colleague of mine who studies the Zurich Anabaptists points out that this disaggregation has made it almost impossible to trace the systematic persecution of Anabaptists in the seventeenth century.

Worse, there’s now a curated collection of ‘Anabaptist documents’ in the archive, assembled from this process, which omits a vast amount of relevant material, and of original context. But

everyone uses that collection, because it's available and labelled.

So it's fascinating that classical archival theorists, many of them working in the service of the state, were the ones most committed to preserving context. They weren't just bureaucrats of fact.

Now, turning to your question about the status of archivists versus historians, and the parallel with literary critics and textual editors, I do think there's a connection.

That division is also a nineteenth-century product.


As Lynn Hunt and others have pointed out, there's a whole rhetoric in English about archivists being the 'handmaidens of historians.' In French, archives have been described as the 'laboratory of history', as if historians are the scientists and archivists merely the technicians.

There's clearly a gendered hierarchy built into that.

In German, there's the term *Hilfswissenschaften*, the 'auxiliary sciences.' These include archivistics, diplomatics (the science of document authentication), sigillography (the study of seals), and others.

Textual editing, too, belongs to that family. It has roots in philology, which carries a bit more prestige, but even there you find the same division between the 'master' philologist and the assistants who specialize in the technical arts.

It's all captured beautifully in Jean Mabillon's *De re diplomatica* (1681), the foundational handbook of diplomatics. In the supplement he published in 1704, Mabillon tries to explain how you recognize whether a document is authentic.



He lists various techniques such as checking dates, signatures, whether the king was in the right town at the time. But in the end, he says it's like a goldsmith who can feel real gold from fake, or like a connoisseur who can tell a genuine painting from a copy.

There's this *je ne sais quoi*, an ineffable authority that distinguishes the master from the helpers.

That's a very old, deeply embedded distinction, and one reason I really admire Michelle Caswell's work. She passionately rejects that hierarchy, and I think she's absolutely right to do so.

NATIONAL ARCHIVES, DC (1937)

In David Lowenthal's *The Past Is a Foreign Country*, he suggests that the discipline of archival studies really emerged in the nineteenth century, especially in the wake of the French Revolution. That moment, along with the rise of Romanticism and a renewed interest in the past, seems to have institutionalized antiquarianism in a new way.

But for those who work on the sixteenth century, antiquarianism and philology feel much older, we associate them fundamentally with humanism and the Renaissance. How does nineteenth-century antiquarianism relate to its sixteenth-century counterpart? Are they connected, or are they fundamentally different intellectual projects?

Yes, this touches on the role of humanism.

It's actually a conversation I've had (quietly, over time) with Tony Grafton. He sees humanism, including antiquarianism, as a critical, transitional, and transformative force. And he's absolutely right that humanist-trained actors played key roles in shaping emerging disciplines.

Now, I used to say (though I've backed off this a bit) that the disciplinary technologies used by archivists in the early modern period were still largely scholastic, that is, they inherited the methods of the older dispensation. Things like pagination, indexing, *mise en page*, this whole apparatus of how information was laid out, was still operating within that scholastic

framework well into the seventeenth century.

But yes, in the seventeenth century, things clearly begin to change.

There's a lovely episode that leads to Jean Mabillon's *De re diplomatica*, which I mentioned earlier. It arises from a debate about whether very old records can actually be authenticated. That debate was triggered by the Bollandists, Jesuit scholars based in Antwerp (in what is now Belgium). They were working to collect and publish authentic records of the saints, the *Acta Sanctorum*, organized by month and saints' feast days.

What really sparks the controversy is an essay by Daniel Papenbroeck, written in 1675. I wrote an article that begins with

that moment. Papenbroeck says, in effect, many of these documents are inauthentic, and it's very hard to tell, because they're very old.

And he makes a striking claim: what might stand up in a court of law can collapse under scholarly scrutiny. In other words, the standards of truth in legal or administrative contexts don't necessarily hold up under critical scholarship.

That sense of incommensurable truths, that something could be valid in law but invalid in scholarship, struck me as remarkably modern. In some ways, even post-humanist. It signals a moment where the hegemony of any single interpretive tradition, whether scholastic or humanist, is breaking apart.

Returning to your point about how placing one document next to another can reshape its meaning, could we say that's analogous to, say, placing a corpus of Indian law next to a corpus of American law and treating the same terms, like 'privacy', as if they meant the same thing, even though they differ significantly in context?

In other words, when we place the bylaws of one Swiss village alongside another's, does that have the same kind of jarring semantic effect? Is that what you mean when you say archives modify the meaning of texts?

Yes, I think that's a good example. Contexts do change, and they do so at multiple levels.

Peter Burke, who seems to be everywhere in intellectual history, has a great article called 'Context in Context', where he explores how difficult and problematic the notion of context is, and yet how central it remains to certain ways of thinking.

As for the idea that discretion is being replaced by technique, yes and no. The example that immediately comes to mind from archival theory is appraisal:

the decision about what not to keep.

And the first thing to note is that there was no concept of archival appraisal in Western Europe before 1800. There was no intellectual framework for throwing things away. The prevailing model was simply: keep everything.

But by 1800, that had become untenable. The story of how the great national archives were formed is deeply bound up with this shift.

One example is the Haus-, Hof- und Staatsarchiv in Vienna, founded in 1749. The man in charge had an imperial licence to create a unified archive of the Habsburgs' most important documents. But as soon as the archive was created, all the Austrian ministries began dumping their records on it. He hadn't anticipated that, and had no theory for how to deal with it.

Appraisal, as it developed later, especially in the twentieth century, is now one of the trickiest areas of archival practice. There's no way to do it without discretion. You have to decide what has archival value.

And here, I think the post-modern argument is very compelling: there is no neutral ground from which to declare value. What you consider important depends on who you are,

where you're situated, and what your institutional role is.

In fact, in most developed countries today, public archives keep, at most, 1 per cent of what they receive, and even that is overwhelming. And that was before the digital deluge.

As late as the 1920s, Hilary Jenkinson, who wrote the classic work on European archival theory, still argued that archivists should not make such choices. His only model was: keep everything.

Then, when the Americans entered the field, Theodore Schellenberg proposed a model in which appraisal becomes a core responsibility of the archivist: distinguishing what has enduring evidential and informational value from what does not.

Australian theorists later pushed back on that. Scholars like Sue McKemmish, Frank



Upward, and others pointed out that if archives are to serve the ethics of transparency and justice, it's not acceptable to let the originating agency make those decisions.

So who decides?

This is still unresolved. Archival science has tried to develop rigid procedures to replace discretionary judgement, but I don't think that has worked. There's no consensus today on a positive theory of appraisal.

Instead, there's a growing awareness that discretion is inevitable. What archivists can do is be transparent about their criteria and document their decisions. You may not be able to record every part of the process, but you

can explain why certain materials were kept and others discarded.

And it's astonishing to reflect on this in historical terms. One thought experiment I use with graduate students is to ask: why did late medieval and early modern states keep what they did?

Take Florence, for example. In 1427, the city conducted an extensive tax survey, the *catasto*. Why are those records still in the archive? They serve no conceivable use today other than for historians. And why did they survive from 1427 to 1800, when historians weren't yet considered a relevant audience?

That's the mystery. Somehow, they made it through centuries when their survival served no obvious purpose.



Speaking of ‘throwing things out’ of the archive: digitization can inadvertently erase crucial features of a document, especially non-visual ones. For example, during an eighteenth-century cholera outbreak, correspondence from affected regions was sometimes disinfected with vinegar by the postal service. That means an original reader might have inferred something important from the smell, something that disappears entirely in a scan.

Do you think this is the kind of context-loss you’re describing in archival practice? And does archival theory really grapple with how even basic acts of copying or preserving can transform meaning?

Yes, smell the letter! Absolutely.

There's a parallel story, I believe, in Arlette Farge's *Le goût de l'archive* (*The Allure of the Archives*). She describes working in the Paris police archives and coming across a document with a faint trace of vinegar, again, because letters were disinfected when there was fear of contagion.

Any transmediation, digitizing, copying, recataloguing, is a change in context. You always lose something. That's entropy.

But archival theory tends to focus less on information loss per se and more on contextual transformation. And that's a critical point.

One underexplored area, at least among early modernists, is how much time people spent copying things. Secretaries, scribes, and officials copied texts constantly.

In Swiss archives, Zurich is a good example, but even smaller cities, you find countless manuscripts labelled *copiarium*, *miscellanea*, or simply 'copybooks'. These are notebooks filled with records transcribed and reshuffled into new forms.

I once gave a paper on this using an inventory from the 1950s of the State Archive in Chur, in eastern Switzerland. The catalogue devotes three hundred pages to describing charters, some of which are just slips of parchment, a couple of lines long, with meticulous summaries, dates, issuers, and serial numbers. Then you turn the page: the remaining 2,700 volumes, copybooks, ledgers, books of accounts, get five pages.

Copying was constant, and each copy reshaped its context. Later, these copies were collected in libraries, reorganized again, meaning that the archive's context is never fixed.

One model that archival theorists have developed to try to capture this is the records continuum. It's quite abstract, and as a historian I still struggle with it, but its core insight is that context

is always in flux. The model is meant to help us observe and reflect on how those contexts change over time.

There's really no way to limit what counts as context. It could be material: studies of the physical document, the ink, the parchment, the folding marks. Text is only one layer of meaning.

Brigitte Bedos-Rezak, a medievalist at NYU, has been central in advancing this line of thinking. She argues that, often, the text itself is not the most important thing.

Let me give you a case I meant to mention earlier. In 1415, the Swiss Confederation, with imperial backing, seized a number of towns and territories from the Habsburgs. To legitimize this, Emperor Sigismund, who, importantly, was not a Habsburg, issued charters to these towns.



DETROIT PUBLIC LIBRARY (1921)

The charters are fascinating because they have no specific content. They just say: ‘Whatever privileges you have, I confirm.’ They don’t list the privileges. They don’t grant new ones. They just acknowledge the towns already have them.

Peter Brun wrote his dissertation about this, noting that the point wasn’t legal but symbolic. The emperor was performing the role of someone who could issue privileges. If you read these documents only as legal texts, you miss that.

Meanwhile, when Sigismund wanted to ensure the loyalty of local petty nobles, he didn’t give

them charters. He ate with them. He hosted feasts. Because those nobles didn’t keep records. The ritual mattered more than a document.

So you get a kind of continuum. Texts are just one tool for creating relationships. A hundred years later, those same charters, still in town archives, are used as legal weapons. The towns dig up older privileges and say, ‘Look, Sigismund confirmed them.’ And they were used very effectively against the Habsburgs. That kind of action is one of the reasons Switzerland exists.

So yes, it’s deeply consequential, even when it has no content.



Oh, I think that's absolutely central, and none of them solved it.

Any archive held by a seven-teenth- or eighteenth-century political body was incoherent, incomplete, and disorganized. You constantly find laments: important things couldn't be found.

There's a great example from the 1590s, when there was a succession crisis in the House of Brandenburg. They couldn't locate the marriage contracts from the previous generation, vital papers, the kinds of things you or I would keep in a labelled folder or box. But apparently they did what we do too: tossed them in a closet and then forgot. The media system of the time just wasn't equipped to manage dynamic records.

Printed forms were one workaround. So was greater centralization, setting up bureaucracies where everything flowed through a single channel, whether at the county or principality level.

Let me give a personal example. I live in California, where we still use an old-fashioned land-title system based on chain of title. Much land can be traced back to Spanish and Mexican grants.

You were talking earlier about how copying a record can change how it's perceived over time. It also made me think about how record systems struggle when the world changes faster than the paperwork can. In large-scale efforts to document something like property ownership, the challenge often isn't collecting information once, it's keeping it current and getting updates to circulate reliably across different offices. When every transaction triggers a change, 'up-to-date' becomes an ongoing, resource-intensive project.

So how aware were eighteenth-century bureaucracies of those limits? And what pragmatic choices did they make about what to update, how often, and what to let lag behind?

About ten years ago, I sold a house in a small town outside Los Angeles. It was part of a five-acre parcel my grandparents had bought in 1960. I sold my portion. Six months later, I got an inquiry via my realtor: the new owners couldn't find a record of the water easement, there was no official document stating that water lines could run across neighbouring properties.

And I had no idea, I hadn't lived there since I was twelve.

It turns out there was no easement. In 1960, you just did it. Nobody wrote anything down. But here's the twist: the new owners wanted to build a swimming pool. The county pulled up the digitized records, and they discovered that the property had never been legally subdivided.

In other words, the five acres were still a single parcel, because

when everything was just paper books and deeds, no one ever saw the bigger picture. But once it's all online, someone can run a simple query and see it. That's the general point: the media system determines what can be known, not just in terms of access, but in terms of visibility.

So it's not just about historical 'context'; it's about technology. And early modern and even nineteenth-century states struggled mightily with this.

Prussia is often held up as a model of successful state-building and bureaucracy. But when I began research at the Geheimes Staatsarchiv (the Secret State Archive) in Berlin, the archivist told me, 'We couldn't find anything.'

Even in the eighteenth century, Prussia only managed to meet its immediate archival

needs. Keeping a corpus coherent and continuously updated? That remained beyond them.

And frankly, that's what Kafka was writing about.

Now, since we're on this topic, and since I was asked to say a bit about my current research, this connects quite directly.

I've been working on the record-keeping of the Swiss Confederation before 1848. Switzerland was a confederation, not a state, so there was no national archive; most records are in the cantonal archives.

But there was a shared decision-making body. The thirteen sovereign cantons, along with some associate members, met regularly, three or four times a year formally, and more informally in between.

Every formal meeting ended with the production of a recess (German: *Abschied*). These records go back at least to the early fourteenth century; they reach a stable format by the 1550s and don't change much until 1798.

Here's the interesting bit: each canton would receive its own copy of the *Abschied*. The scribes would write up one version for every participant, and each version was different, containing only the articles relevant to that canton.

So every archive today has a slightly different collection of the same meetings. It's a distributed archive.

They're not 'proofs'; they're documentation of ongoing conversations.

Much of what appears in the *Abschiede* is actually what couldn't be resolved during the



NATIONAL ARCHIVES, DC (1974)

meetings. Issues were kicked back to the cantons, which would then deliberate and send more material. That creates a whole cloud of related documents: evidentiary materials, instructions, responses.

And the heart of that cloud is the *Abschiede*.

In the eighteenth century, they began cataloguing and indexing them to make them more accessible. In the nineteenth century, they attempted to produce a printed edition, a single, centralized archive where none had existed. But that effort was a complete disaster. They had to omit documents, they drew very unequally on surviving

material, and the organization was inconsistent and often bizarre.

What I'm doing now is studying this distributed, heterogeneous archive as an archival corpus. What kind of object is it?

Each canton had slightly different versions, some more complete, some less. They're usually filed chronologically by meeting date. But despite the differences, by the eighteenth century people had come to imagine the collection as a single archive.

It's a bit like the cloud, where everyone sees their own version of Facebook, and yet we all think of it as one thing.



You mentioned the cloud and Facebook earlier, and the concept of the information state, how a state begins to mobilize large bodies of documents to strategize for the future, kept coming up while reading the article. It also calls to mind Shoshana Zuboff's idea of surveillance capitalism: big tech firms collecting and monetizing massive amounts of behavioural data, the commodification of personal experience. Given the shift from states trying to 'keep it all' to the digital age making that possible again (in a very different way), does the idea of the information state anticipate surveillance capitalism in any way? Or is there a meaningful intersection between them?

Yes, I think there is. The way I've tried to frame it, and this may well overlap with what Zuboff is getting at, is that what we're seeing now is not just a digital

revolution, but a transformation in political economy. The systematic commodification of information is central to this.

We're in the middle of a major media revolution, and a structural change in how information operates in power relations.

Late medieval and early modern Europe operated under what I'd call a political economy of privilege. Power, access, control, these were managed through privileges: rights granted to individuals, towns, guilds, institutions.

That system doesn't vanish overnight, however, it does break down. Whether you date it to the sixteenth or the nineteenth century, it is gradually replaced by the political economy of modern national capitalism.

Now the nation-state becomes the guarantor of property, now redefined as capital.

Property is no longer a privilege; it's a universalized, abstract right.

Why that shift happened is a vast question with too many competing theories to unpack here. But what's clear is that each of these dispensations has its own logic of record-keeping.

In the privilege economy, archives are about tracking the particular, who holds what right, where, and under what conditions. That's why medieval archives are full of charters, lists, seals: highly specific, situational records.

With modern state capitalism, the unit of governance becomes the agency. Archives start to reflect that: the record group, the *fonds*, becomes centred on institutional function.

Now we've entered something new, whether you call it post-capitalist, financial capitalism, or surveillance capitalism.

There's clearly a transformation under way.

Information itself is now commodified. We're seeing what you might call a 'user-centred' paradigm of archiving. Each of us is, in a way, our own archivist. I have something like 50,000 emails from the last fifteen years. That's an archive that would've been unimaginable on paper. But it's also unwieldy. It just sits there in my inbox. I don't have the tools, or even the language, to manage it properly. And Silicon Valley is working very hard to develop those tools.

So yes, I think we are in a new dispensation. And it's not surprising that archival theory and archival practice are being transformed in response. We're not just talking about methods any more, we're witnessing a shift in how information relates to power.

If we start thinking of things like the city or the body as archives, how do we meaningfully distinguish the archive from other kinds of textual or interpretive material? What makes an archive distinct, if its boundaries are so porous? Does the conceptual fuzziness of the archive risk making archival studies fuzzy as well?

Texts are transmitted in different ways, and non-texts are transmitted in different ways. Not all of that is archival. Some of it is, but what makes it archival?

One dimension is its association with authority. Not necessarily the authority of the individual as creator of the archive, but historically it's typically been associated with the authority of some kind of dominion: a state, a city, a church, a business, some kind of institutionalized authority, or at least identity, that makes

it capable of producing an archive.

The second is investment of effort. It takes effort to distinguish, to create an archive, archives, out of a pile of paper.

One of my early thought experiments in this project, trying to define this, was: if I take my twenty-year subscription to the *American Historical Review* and throw it into a dumpster, is that an archive? Well, it's ordered, I've thrown it in chronological order, old issues first, then newer

THE TIMES OF INDIA, BOMBAY (1898)



ones. But I put it in a dumpster, so it's probably not an archive.

Another wonderful thought experiment is the genizah. The most famous is the Cairo Geniza. Jewish ceremonial law requires that any inscribed material that might contain the name of God must be preserved or given a formal burial. So Jewish temples and synagogues, even today, have what's called a genizah: any unused paper goes in.

In one Cairo synagogue, they designated a storeroom next to the building as their genizah and threw stuff in there over many centuries. Much of it was ceremonial texts, prayer books, but they threw in almost anything, because it might have the name of God in it. Some things were eventually removed and buried in a graveyard; there are cemetery genizot as well. But most of it just sat there, forgotten,

until the nineteenth century, when antiquities dealers began buying it up. Much of it ended up in Cambridge, where people realized this was a priceless set of historical sources, not just about Jewish society but about Egyptian society.

So at no time is everything archival. But at any given time, in relation to the institutions and the practices of making, keeping, and using documents, what can be included in a given time and place, or for a given context, is quite fluid.

You can have an archive of dance. I had friends who were dancers in college, and dance notation is notoriously difficult. Now you have videos, you can create a dance archive in ways that were difficult before. Of course, there are older forms of notation and older dance archives, but the creation of a


dance archive depends exactly on appraisal, on description, on organization. You can produce an archive of movement.

Same with images. The emergence of photography allowed the archiving of images. Terry Cook has another article in which he struggles with this: the way archives dealt with images, and the very strong tendency, until quite recently, to extract all the photographs and put them in a separate photo archive. Which is, of course,

terrible archival practice, you're destroying context when you do that. But it took a long time to think, well, maybe the way to archive the visual isn't to just pull it all out.

So, at any given time, in any given institutional context, of authority, of practice, of intention, there is certainly a fuzzy boundary, but there are also ways to distinguish what, at the time or potentially later, operates as archives.





RADIO OPERATORS
ON AN EXPEDITION IN
CENTRAL NEW GUINEA (1920)

**Paul
Duguid**

**On the Social Life
of Information**

I want to revisit a term that now sounds almost passé, but in the 1990s was very much alive: the ‘information age’. Information was hot. Society, it was said, was being driven by information. We were told that an information revolution had transformed everything.

Marshall McLuhan, in his influential 1964 book *Understanding Media*, had already spoken of an ‘age of information’. He argued that it was driven initially by electricity and, increasingly, by television, which, for the 1960s, was an apt observation.

A striking illustration of the rhetoric of the time comes from a two-page IBM advertisement in *Fortune* magazine from July 1977. The headline announced: ‘Information: There is growing agreement that it’s the name of the age we live in.’ By then, it had become familiar to speak of the ‘information age’, and IBM was actively helping to define it. The message was that we were drowning in information; it was shaping our lives; and if you wanted to manage it, you needed IBM. The company cast itself as both the source of the problem and the solution.

This fed into a broader understanding that we were not only living in an ‘information age’, but that a thing called ‘information technology’ was driving it: the computer, then the internet, then the World Wide Web. Today we might add large language models and artificial intelligence to that list, all framed as technologies fuelling a continual information revolution.

One of the ways ‘information’ became such a critical term goes back to the Second World War and the years just before it, particularly through the work of Claude Shannon. Shannon developed what he called a ‘mathematical theory of communication’, which came to be known simply as information theory.

He worked at Bell Labs, a research laboratory created by AT&T, where he focused on improving message transmission over telephone lines. His model was simple but powerful. There is a source of information, from which the message to be transmitted is extracted. It becomes a signal, sent across a wire. Noise is stripped away. At the receiving end, if the message that arrives matches the one that was sent, the

transmission has been successful. This was the basis for what became information theory.

Shannon developed most of this during the Second World War, when he was involved in secret signalling, working on how to disguise messages and make them secure. In that context he formulated what is often described as an entropy measure, dealing with the probability of a message being predictable or not. It gave us a calculable, formal model for talking about information, something we could measure, extrapolate, and project into the future.

Other models followed. One that most people have at least heard of is Moore's Law, the idea that information technologies become increasingly efficient over time. From this, it follows that the transmission of information and its effects will continue to grow dramatically. On this view, the changes we are seeing now are merely the prelude to even more dramatic transformations still to come.

This idea persists today, particularly in conversations about AI. The prevailing notion is that even if current tools seem somewhat inept, faster machines will eventually arrive, machines of the sort imagined by Moore and pursued by IBM, and they will solve the problems we currently cannot. That is the promise. It is also deeply problematic, a point I will return to, but there is no denying how influential this way of thinking has been.

One of its more surprising but important areas of influence is economics, where 'information economics' became a critical part of the field, leading to several Nobel Prizes.

One of the first was Friedrich Hayek. He promoted the idea that economics was the foundation of society and social systems. Markets, he argued, did not need government interference; people simply needed the right information. If they had it, they would make the right decisions, and society would arrive, more or less automatically, at an Adam Smith-style harmony.

A later Nobel winner, George Akerlof, at Berkeley, approached the matter differently. He pointed out that markets, and society as a whole, become problematic when information is asymmetrical. If one party knows more than another, markets fail. His idea was that if we could develop technologies to equalize information, to make it symmetrical, people would make the right decisions. Technology, then, could satisfy not only our economic but also our social needs, promoting harmony and coherence.

This thinking was deeply influential and extended far beyond economics. Cognitive science grew up around some of these ideas. Social scientists leaned heavily on them. And we ended up in a world where, as the old English saying has it, 'to a man with a hammer, everything looks like a nail'.

Once we framed everything in terms of information, then, this approach presumed, we had the hammer to fix it. If the world could be defined as an information problem, then technology could provide the solution. John Seely Brown and I wrote *The Social Life of Information* (2000) while working at Xerox PARC, the Palo Alto Research Center, at the heart of Silicon Valley. It was a fascinating place to be.

MANILA, US ARMY SIGNAL CORPS
EXTENDING TELEGRAPH LINES
FROM THE TRENCHES (1898)



George Pake, whom Xerox had asked to set up the centre, was instructed to build a research hub that would explore what they called the ‘architecture of information’. As he once told me, he had no idea what that phrase meant when they asked him, and, fortunately for him, no one at Xerox seemed to know either. That uncertainty gave him considerable freedom.

He used it remarkably well. He recruited many of the leading computer scientists from across the United States and beyond. At the time, ‘leading computer scientists’ described a relatively small group, unlike today, where the phrase might denote thousands. Together, the group at PARC accomplished extraordinary things. One of their most significant achievements was the Alto, the first desktop computer. It featured a graphical user interface (GUI). It did not invent the mouse, but it integrated it, allowing users to move around the screen visually. In effect, it provided the basic interaction model that most of us still use.

There is a famous story, often told rather too simply, about what happened next. Steve Jobs was brought to Xerox PARC, shown the Alto, and realized that Xerox did not know how to market it, but he did. He took the underlying ideas and used them to build the Apple computer. That was followed by the IBM PC, and then by Microsoft and others.

What interested us was not just the process by which the technology developed, but the way the company that had created it failed to bring it effectively to market. A well-known book of the time, *Fumbling the Future: How Xerox Invented Then Ignored the First Personal*

Computer, captured this irony. You could achieve extraordinary technological breakthroughs and still fail to communicate or distribute them. When you possess an information technology yet seem inept at spreading information about it, something strange is clearly going on.

‘Tunnel vision’ here might be better described as ‘linear vision.’ It was symbolized by *The Road Ahead*, Bill Gates’s book about the future of technology. The road on its cover is strikingly straight. You simply follow it, and it carries you steadily into the future: no forks, no choices, no oncoming traffic, just forward motion.

Brown and I challenged this linear vision by examining what we called the ‘six Ds,’ recurring predictions about what technology would do to the world. According to this view, technology would bring about the demassification of objects, the decentralization of society, the denationalization of states, the despatialization of activity, the disintermediation of markets, and the disappearance of traditional institutions. All of this assumed technology would unproblematically transform everything.

There were confident claims that the conventional workplace would vanish, that we would all work from home, or from anywhere we chose. Centralization would be consigned to the past. In academia, the argument ran, we would no longer need universities. Everything could be put online. Education would be ‘delivered’ rather than experienced. As long as learning was understood as the absorption of information, technology could handle it.

A French illustration from around 1900 captures this imaginary future with uncanny precision. It is a prediction for the year 2000: a

teacher feeds books into a machine that grinds up their contents, while students sit with wires running into their heads, passively absorbing the processed information. No more learning, just automatic intake of facts. It is a wonderfully vivid illustration of what happens when education, society, and indeed the whole of life are reduced to information problems waiting to be technically solved.

That fantasy of frictionless information has never quite gone away. Nicholas Negroponte at MIT promised ‘one laptop per child’ and claimed that this would bring every child into a global educational system. Later, in a more unsettling variation, he suggested, and Elon Musk has echoed the idea, that we might eventually inject information directly into people’s brains. No earphones required. People would simply know.

There is an economic version of this dream as well. Paul Romer, another Nobel laureate, argued that the root of inequality lay in physical distance. We cannot ship material goods everywhere, but we can ship ideas. Once information replaces objects, inequality disappears, or so the claim goes. These were grand, sweeping visions founded almost entirely on the power of information. They prompted us to ask: how might we think differently about all this? The answer took us, interestingly enough, back to Xerox itself.

Xerox had invented the photocopier. With strong patents in hand, the company sued competitors, including IBM, to defend its position. As rivals emerged, Xerox realized it had to stay ahead by producing more sophisticated machines. The problem it then discovered, rather reluctantly, was that although a new machine might offer 150 features,

most users only ever used about six. All those added capabilities were, in practice, irrelevant.

Xerox's solution was to include a set of instruction cards to help people use the machine. But users did not understand the cards. Someone then proposed creating a second set of cards to explain the first. So began a kind of infinite regression: cards explaining cards explaining cards, all attempting to explain how to operate the machine.


The underlying assumption persisted: if we just delivered the right information, the user would be able to solve the problem.

The anthropologist Lucy Suchman, a colleague at Xerox, captured this dilemma in a wonderful video titled *When User Hits Machine*. One lesson that emerged from such work was that whenever designers told management the machines were hard to use, the response was dismissive: 'That's nonsense. Only idiots can't use it. I've used the machine many times, it's fine.'

But closer observation revealed that when a manager was escorted to a machine, innumerable subtle nudges, social cues, contextual hints, other people's gestures and comments, guided them towards the right buttons or actions. Without those cues, the machine became far less intuitive.

Suchman created a brilliant video to make this point. In it, two people are given a task: take a document, extract certain pages, reorder them, and bind them into a new document. Chaos quickly ensues. They tear pages from the manual, shuffle piles of paper, and gradually approach despair.

EMPLOYEES OF THE POST, TELEGRAPH AND TELEPHONE SERVICE IN BANDUNG (c.1930)



When the video was shown to managers at Xerox, the first reaction was predictable: if you get two janitors, or ‘someone like that,’ to do the task, of course they will struggle. Then it was revealed that the two people in the video were Allen Newell, one of the founders of artificial intelligence, and Ron Kaplan, a PARC colleague who had helped design earlier systems. In other words, world-class computer scientists still could not operate the machines smoothly.

This revelation helped trigger a shift in thinking. PARC funded a new initiative, the Institute for Research on Learning, which is where I joined Xerox. The Institute brought together cognitive scientists and anthropologists, united by the conviction that we needed a new way of thinking about learning and about how information fits into it.

Most conventional theories of education treat learning as a straightforward transfer: someone who knows passes information to someone who does not, and if the recipient ‘gets it,’ learning has occurred. But this model takes a great deal for granted.

Information, in this view, comes from the outside world, flows, like a telegraph signal, through our nervous system, and ends up in our heads, where it becomes knowledge. In that simple story, we already have three quite different understandings of ‘information,’ yet they are often treated as if they were the same thing. More importantly, the learner is reduced to a passive machine. There is no room for interpretation, understanding, or reflection.

Within Xerox, this view was challenged by the work of another anthropologist, Julian Orr, whom we discuss in the book. Orr studied Xerox repair technicians and followed them in the field. Most of the

day, these technicians worked alone. But they had breakfast together, lunch together, and often dinner as well. These gatherings were not merely social; they were essential.

Over meals, they exchanged stories and techniques. The official repair manuals often did not work, but through shared experience the technicians developed their own methods for solving problems. Crucially, they could never admit this to the company. Officially, they were ‘following the manual’. In reality, they were constantly improvising and innovating because the manuals were inadequate.

This work revealed two key ideas that became central to the book: ‘situated learning’, learning that happens in specific material and social contexts, rather than in abstraction, and ‘communities of practice’, groups that build and share knowledge collectively through their ongoing activity. These concepts offered a very different understanding of information: not as isolated packets passed from sender to receiver, but as something deeply embedded in human activity and social interaction.

The implications were significant. At the time, as technology companies fixated on the various ‘D’ predictions, one of the most fashionable was the ‘death of distance’, the notion that place no longer mattered. Yet the actual geography of technology told a more complicated story. A well-known *Datamation* cover story in 1981 depicted Boston’s Route 128 as a hotbed of computing. Companies like DEC, the leading computer manufacturer of the period, clustered there. But over time they were displaced by a very different place: Silicon Valley.

That shift is revealing. Consider Facebook. It began as a project at Harvard, with the rhetoric of a global platform on which place did not matter. Yet one of Mark Zuckerberg's first major decisions was to get on a plane and move to Silicon Valley. 'This is where I've got to build my company.' The claim that distance and place no longer matter turns out to be far more complicated than the rhetoric of the information age suggests.

AnnaLee Saxenian's *Regional Advantage* explores this geography of technology with great clarity. In Boston, firms operated in highly isolated silos. Conversations across corporate boundaries were treated as potential leaks of trade secrets, discouraged and, in some cases, punished.

Silicon Valley looked very different. People moved easily between companies. Someone might leave Xerox PARC to start a company, return for a period of research, and move on again. Ideas flowed back and forth in a much more communal, social way, both within and across corporations. The region's 'advantage' lay not just in its technologies, but in the social and organizational patterns that supported the circulation of knowledge.

To begin drawing these threads together, it is useful to return to some of the classic ideas around information. Claude Shannon's work gave us the concept of the 'bit' and a model built on binary oppositions: signal versus noise, information versus distortion. Information is what you want in the channel; noise is what you exclude. But if you read Shannon carefully, you realize that we define information by identifying the noise. It is about what is excluded as much as what is included.

TELEGRAPH TRAINING,
CAMP FUNSTON, KANSAS (1920)



A similar logic appears in economics. Economists construct models and designate anything that does not fit as an ‘externality’, a factor to be set aside. But who decides what counts as an externality? And is it always irrelevant?

Similar questions can be asked of information asymmetry. The standard claim is that markets fail when people do not have equal access to information. Get the information right and symmetrical, and everything will function. But even if we have the same information, will we necessarily do the same thing? That is far from obvious.

We see comparable binaries in our accounts of learning. We often rely on the simple opposition between instructor and learner: the teacher imparts information, the learner passively receives it. Those of us who teach know that this is only a partial truth. We learn a great deal from our students, even if we do not always rush to acknowledge it. In practice, learning is reciprocal.

Workplace organization has been framed in similar terms. Scientific management, the old Taylorist model, was built on a principal–agent binary: the boss tells the employee what to do, and the employee does it. Companies such as Amazon still rely on versions of this idea. Yet, like the photocopier technicians at Xerox, workers are often doing all sorts of things that are not officially visible, what is now called invisible labour. Organizations benefit from this unacknowledged work while pretending it does not exist.

All of this brings us back to the social. Much of the dominant economic and cognitive science analysis focused on individuals: the lone learner, the rational market actor. We never denied the importance

of the individual, but we argued that you cannot understand information, learning, or technology if you ignore the social context in which they are embedded.

One way we tried to make this point was by distinguishing between invention and innovation. Xerox effectively invented the personal computer but failed to innovate, failed to bring it successfully into the world. That final step is a major leap. Invention alone does not guarantee adoption. To innovate, you must understand the communities you are addressing: what they need, how they might want to use the technology, what they value and resist. You have to adapt to them, not simply expect them to adapt to you.

The question of knowledge itself illustrates a similar tension. In the mid-1990s, when ‘knowledge management’ was fashionable, I attended a conference that brought together ‘knowledge managers’ from across organizations. They were sharply divided.

Half insisted that their central problem was that knowledge was sticky. Their companies had great ideas, but they struggled to move them to the right place at the right time. The other half insisted that the problem was precisely the opposite: knowledge was leaking out everywhere. Everyone was stealing it and taking it away. In one view, knowledge is impossibly hard to move; in the other, it is impossible to contain. How do we make sense of that? Not by treating ‘knowledge’ as an isolated entity, clearly, but by focusing on social context, on engagement, on the practices of the people who use and produce it.

Another term that became important, especially when thinking about how people actually engage with technologies, was the tendency

to divide everything into resources and constraints. We say: here is the resource, and here is what holds it back. Remove the constraint and the world will work better.

Consider the document. The standard view is that a document simply contains information. If we could strip out the information, get rid of the paper, that clumsy, analogue nuisance, then the information could flow more easily. It could be stored, processed, transmitted around the world. Over time, we began to realize that documents are more complex than this.

While working in an archive in Portugal, a researcher was seated beside me with a set of nineteenth-century letters. I noticed that he was occasionally sniffing the letters and examining them very closely. What he was doing was fascinating. During cholera outbreaks, the Portuguese postal service disinfected letters with vinegar. This allowed scientists and historians to trace how the disease spread, not by reading the letters, which often insisted that ‘everything is fine, business as usual’, but by attending to other signs. The people writing those letters did not want to admit that cholera had broken out. They feared that trade would suffer, that goods and money would stop arriving. But a recipient who noticed the smell of vinegar or the stains on the paper might read the message quite differently.

If we were to digitize those documents, we would certainly gain some benefits: wider access, searchability, preservation. But we would also lose something crucial. The very cues that tell you how to read the document, the smell, the stains, the material traces of disinfection, would disappear. What looks like a ‘constraint’ in a narrow view of

information turns out to be a rich and essential resource in human exchanges.

This brings us back to Marshall McLuhan's famous claim that 'the medium is the message'. You cannot, as Shannon's model might tempt us to think, simply extract a neutral content from its vehicle and send it across a channel unchanged. Form matters. The way information is constructed, embodied, and conveyed is critical to how we understand it.

After Xerox 'fumbled the future', ironically, just as *The Social Life of Information* appeared around 2000, the company began to shrink. Despite earning vast sums from inventions like the laser printer and Ethernet, it never really embedded those innovations into its core business. There is a growing sense that history does not apply to information; that information is new, fast-moving, and inherently present tense, whereas history is old, outdated, and irrelevant. John Perry Barlow, a prominent figure in the Electronic Frontier Foundation and in the radicalization of the early web, once declared: 'On behalf of the future, I ask you of the past to leave us alone. You are not welcome among us. You have no sovereignty where we gather. . . . You are terrified of your own children, since they are natives in a world where you will always be immigrants.' The past, in this view, has nothing to contribute.

Anthony Levandowski, another emblematic figure in Silicon Valley, now perhaps better known for his conviction for trade-secret theft and subsequent presidential pardon, once remarked: 'The only thing that matters is the future. I don't even know why we study

TELEPHONE SWITCHBOARD ROOM
IN HOTEL CRILLON, HQ OF THE
AMERICAN DELEGATION AT THE
PARIS PEACE CONFERENCE (1919)



history. . . . In technology, all that matters is tomorrow.' As it happens, he was also a Berkeley student; I sometimes wonder, half-jokingly, whether he was the one who wrote that single-word complaint on my course survey.

Behind such statements lies a strong conviction that the past is irrelevant rather than instructive. This is precisely what we need to resist. History does have something to say about how we understand information, facts, values, indeed, many of the questions that preoccupy us today.

Consider Vicesimus Knox, an English clergyman writing in 1778. In an essay on books and technology, he observed that books had become widely available 'even in this age of information.' Already in the late eighteenth century, someone was characterizing his own time as an 'age of information.' What Knox went on to say is equally striking. Although books were plentiful, he argued, what people really needed were communities, teachers, and those around them. Learning, he warned, could not take place in 'solitary retirement.' Isolation would lead to ignorance and inequality, because some people would have access to resources while others would not. His concerns about access, community, and inequality sound remarkably contemporary.

Knox was also an optimist. Intrigued by the French Revolution, he wrote, 'Give them fair and full information and they will do the thing that is right, in consequence of it.' He believed in democracy; he believed that if people had access to good information, they would make the right decisions.

This kind of optimism runs deep. There is a persistent faith that if we simply get information ‘out there’, people will find what is true and act accordingly; that democracy, social order, and justice will naturally follow. The lineage of this belief can be traced much further back. Shakespeare, though he did not use the word ‘information’, gave us the phrase ‘truth will out’, the idea that truth will eventually surface, so long as it is allowed to circulate.

During the English Revolution in the seventeenth century, John Milton made a similar argument. In his famous defence of unlicensed printing, *Areopagitica*, he insisted: ‘Let her and Falsehood grapple; who ever knew Truth put to the worse, in a free and open encounter?’ Even if lies and nonsense abound, if truth is present too, truth will win. These are powerful, idealistic notions, and they continue to shape how many people think about information and technology today: just let the truth be out there, and everything else will take care of itself.

What the history of information suggests, however, is that things are rarely so simple. Who has access, how information is framed, the communities in which it circulates, the media and institutions that sustain or distort it, all of these matter. The social life of information is not a recent invention of the internet age; it has been with us for as long as people have tried to learn from, act on, and live together through shared representations of the world.

Plato takes us even further back into the history of these questions, and his work remains uncannily relevant. In the *Charmides*, Socrates poses a dilemma that still resonates today, especially in the wake of Covid and contemporary debates about who should lead institutions

such as the National Institutes of Health. He asks: if a wise person, or any person, wants to distinguish the true physician from the false, how should they proceed?

It is a profound problem. If you are ill and consult a doctor, how do you know whether the doctor's advice is sound? Socrates points out that if you can reliably judge the advice, then you do not need the doctor; if you need the doctor, it is precisely because you cannot judge the advice. We remain caught in this tension. It is about trust, expertise, interpretation, and the limits of information, issues that are still very much with us.

Across Plato's dialogues, a partial answer emerges. We must ask not only what the advice is, but who gives it and in what context. Who recommends the doctor? What is their standing, their quality, their reliability? How can we know? Centuries later, George Akerlof essentially reformulates the same insight. Doctors, he observes, are authorized by licensing. So are lawyers. So, in a more modest way, are barbers. A high school diploma, a BA, a PhD, a Nobel Prize, these all operate like brand names.

We do not rely only on information; we rely on the social and institutional structures that surround it and validate it. Academic titles and institutional affiliations, for example, function as brands. When a speaker is introduced as a professor at a major university, that introduction lends their words a particular kind of weight. Students listen in part because they want to carry that brand out into the world with them. When they apply for jobs, employers do not look only at the individual; they also look past the individual to the university that stands behind them.

GERMAN TRENCH TELEPHONE STATION (c.1915)



...ation.
...ion.

The 'Cal' brand used by Berkeley is itself historically rooted. In Gold Rush California, state assayers stamped gold ingots with 'CAL 1850', certifying that the gold was genuine and validated by the state. That branding has changed over time, but the underlying logic persists. It remains a mechanism for certifying value and validating trust.

This led me to an interest in trademarks and branding more broadly, and eventually to the history of trademarks in wine. A bottle of wine is a complex substance. Faced with it, we want to know: what is in this? Is it safe? Is it worth the price? We turn to a whole chain of associations: the seller's name, the producer's brand, the grape variety, the region of origin. Each layer provides a form of endorsement, a promise, a signal of quality.

This offers a useful way to think about information. We rarely take 'raw' information on trust. Instead, we attend, consciously or not, to layers of endorsement: who produced it, through which institutions, under what constraints, with what forms of authorization. When people predict that information technologies will make institutions obsolete, that we will no longer need governments, universities, or regulation, we should be wary. These institutions play a vital role in assembling, validating, and circulating information.

All of this feeds back into the larger argument. We often treat information as something autonomous: a substance that can be taken from one place, dropped in another, and expected to work unchanged. We imagine it as unidirectional: flowing from teacher to learner, from expert to novice. We treat it as causal: if people just receive the right information, they will change their behaviour. And we consider it

objective: something we can quantify like Shannon's bits, counting how much information a book or a webpage contains, as though we could simply tally it up.

Gregory Bateson famously defined information as 'a difference that makes a difference.' This implies causality: information changes how its recipient sees the world. But it also implies subjectivity. What counts as information for me may not count as information for you. That point, too, is implicit in Shannon's work.

If we ask how much information is in a book, we cannot answer simply by saying 'this many bits.' We must also ask: who read it? Did they understand it? Did they read all of it, or only parts? For someone who has already read the book, only some of it may still be informative. The quantity of information cannot be separated from the history, knowledge, and situation of the reader.

This yields a very different, far less purely objective view of information. And yet, time and again, we return to information as a universal solution. We encounter a problem and immediately type a query into Google, assuming that whatever answer comes back will fix it. We know, from experience, how limited that assumption can be.

Information, then, is not a simple, autonomous object. It is produced, interpreted, validated, and used within social worlds, communities of practice, institutions, professions, publics.





TELEGRAPH OFFICE, SURABAYA (1919)



Could you say more about how the classic distinction between information and knowledge collapses once we look at the material and international division of labour behind information flows? Along the lines of Donna Haraway's 'informatics of domination' and 'women in the integrated circuit'. How does that history complicate the illusion of information's immateriality and the binaries you work with?

We've inherited this neat hierarchy, data, information, knowledge, nicely captured in that much-quoted, little-read line from T. S. Eliot: 'Where is the wisdom we have lost in knowledge? Where is the knowledge we have lost in information?' We talk as if these are cleanly separable categories, but

once you look closely, especially in context, they start to blur.

'Information' caught on in part because it seemed to let us sidestep all the mess of epistemology, centuries of argument about how we know what we know, and replace it with something apparently precise and mathematical. But epistemology doesn't go away just because we rename it. The history still weighs on what we do.

You can see the same move with each new technology. Every era discovers a device that will finally annihilate time and distance and produce informational harmony, canals, railways, the telegraph, and so on. Zuckerberg, as I like to say, runs the story from Gutenberg to Zuckerberg, as if print revolutionized the world and now Facebook will do it again. The rhetoric is remarkably repetitive.

Haraway's work is crucial here because she insists that these supposedly immaterial 'flows of information' are grounded in very material labour, often performed by women under exploitative conditions. Her phrase 'women in the integrated circuit' captures how gendered labour is literally built into the hardware and yet rendered invisible.

You see similar patterns in something as banal as Xerox marketing: the pitch that ‘these machines are so easy, even women can use them’, and the infamous advert in which a chimpanzee operates the copier. Women and chimpanzees collapse into the same figure: the user who doesn’t really count. That attitude carries through to things like the so-called Google memo, with its suggestion that women lack the ‘right’ temperament for engineering and that empathy is a problem, not a resource.

These binaries, rational/emotional, technical/social, immaterial/material, do a lot of ideological work. They make certain kinds of labour disappear. In our Xerox study, the technicians kept the system running by constantly improvising, breaking the rules, and drawing on local

knowledge. Management insisted they were just ‘following the manual’. Their skill had to be invisible for the ideology to hold.

The same is true of today’s information economy, Amazon Mechanical Turk workers, content moderators, and the global supply chains that produce devices and run data centres. A great deal of poorly paid, precarious, often gendered and racialized labour lies behind the shiny story of frictionless information.

So, to your point, once you bring those material and international divisions of labour into view, the immateriality of ‘information’ looks much more like an illusion, and the tidy binaries we rely on, information/knowledge, technical/social, start to crack. That’s precisely why the social context matters so much.

Bacon, in ‘Of Studies’, criticizes treating reading as an end in itself and insists that wisdom comes from combining study with experience. That seems close to your critique of an info-centric view, that simply acquiring information is enough. How do you see Bacon’s empiricism and his suspicion of institutions fitting into your argument about information, knowledge, and the social contexts that validate them?

Yes, Bacon is an excellent figure to bring in. He’s usually cast as a precursor of empiricism: go out, look for yourself; don’t let Church or state do your thinking for you. That runs through Locke, Hume, and on into the early Royal Society: an ideal of open contribution, anyone sending in observations, no names, no institutional titles.

But as *Philosophical Transactions* evolves, you start to get Hooke, Newton, Boyle, then the letters after their names: FRS, Oxon, and so on. Authority and institutional weight creep back in. Bacon might be tearing his hair out, yet he might also concede that, once we move from

knowing for ourselves to sharing knowledge, we do need some means of validation.

So there’s a tension: we want Bacon’s independent observer, but we also don’t want every opinion to carry equal weight. That’s where institutions come in, and where they can go wrong. The comforting Shakespearean idea that ‘truth will out’ regardless of who speaks is attractive, but in practice the interplay of individual judgement and institutional authority is much messier. That’s exactly where the info-centric view, with its promise that ‘more information’ will solve everything, tends to obscure what’s really at stake.

In many of your examples, companies act as if they've got enough information to control everything, then overlook or ignore crucial factors. Is that just a matter of hidden variables they fail to consider, or are there things an info-centric worldview simply can't grasp?

We don't want to suggest, 'They are fools and we are wise.' Xerox, for instance, did extraordinary things. The question is: how can you do something so impressive and still miss the mark so badly?

What we found wasn't simple stupidity but a sorting process: deciding what counts as 'relevant information' and what can be ignored. You can't take everything into account, of course, but how you draw that line, what goes in, what gets left out, shapes everything.

That's why I like the image from *The Road Ahead*: a perfectly straight line into the future. Organizations almost need that. In capitalist economies, you have

to have a goal, a story of progress, to hold the structure together.

But research, particularly at places like Xerox PARC, doesn't move in straight lines. If you propose something and can't say how it will hit the bottom line, it gets shut down. That's increasingly true at places like Google Research. People doing good, exploratory work are pushed out because they can't immediately translate it into profit.

So, yes, you need goals. But you also need to see how those goals blind you. Organizations find that very hard to acknowledge; it threatens their own structure.

This is where AnnaLee Saxenian's work on Silicon Valley is so helpful. What mattered wasn't just innovation, but mobility. If you hit a wall at one firm, you could move to another. Ideas moved with people. That



AN AMERICAN SOLDIER USES A
FRENCH RADIO AT WAR (1918)

circulation loosened the grip of any one company's teleology and created something like a collective horizon. I may romanticize it a bit, but I think there was something real there.

Most organizational theory, and especially economics, struggles to recognize that. It works with very narrow notions of what information is and what it's for. Once that frame is set, everything outside it disappears. 'Tunnel vision' was our image, not only for what's seen, but for what must be excluded in order to keep that straight road in view.

Has the rise of behavioural economics helped bring social context back into how corporations and data analysts think, or is it just a new instrument built on the same assumptions, another way to control rather than understand?

If you go back to Hayek and the early stages of information economics, he suddenly realizes how central information is to his work. But his argument is essentially: governments can't really know people's interests, so they shouldn't be in the business of shaping or directing them. People should be left to make their own decisions, very much the Adam Smith model. That was the ideal economy.

Economics followed that line for a long time. What's remarkable is how behavioural economics marks a huge shift from that position with almost no vis-

ible theoretical debate about the turn. How did we get from 'don't shape people's preferences' to 'let's design ways to nudge them'? It's a transformation that seems to have slid in under the radar.

On one level, behavioural economics looks like a recognition of the social, of context, bias, interaction. But at the same time, it does precisely what Hayek feared: it offers tools to manipulate people, often without their awareness, and to justify that manipulation as 'scientifically grounded'.

That ties into the picture I mentioned of forms of learning where you don't get to discriminate or reflect, you just receive input. Musk's fantasy of injecting ideas directly into your nervous system captures this starkly: no participation, no reciprocity, just implantation.

We've seen versions of this for a long time in advertising and public relations. The market is no longer simply about open choice; it's about how people can be nudged, shaped, guided, while still being persuaded that they are making free, rational decisions.

That contradiction, that we are both carefully steered and constantly told we are freely choosing, seems to me one of the most troubling aspects of the current information economy. And I don't think economics has yet done the theoretical work to confront how radical that shift really is.

In the seventeenth and eighteenth centuries, the anxiety wasn't so much about information circulating too freely as about how to collect and record observations in the first place. Do you see that empiricist culture as a kind of prehistory of

the information age? And, as in Steven Shapin's *A Social History of Truth*, was empiricism already offering a supposedly neutral method while masking questions of credibility and authority, much as today's technologies do?

Yes, I think there's a deep continuity there. One of the things I've tried to work on is this idea of the 'information revolution', remembering that 'revolution' is a cyclical term. If you look back, you see recurrent anxiety: too little information, too much, the wrong kind.

When Gutenberg appears, we now like to say, 'He gave us democracy.' At the time, people felt overwhelmed: suddenly there were too many books. If before you had a handful of manuscripts, now you had a flood. So people had to work out how to sort, interpret, and validate what they were reading. That's a very familiar problem.



RADIO ON THE HEADPHONES, NEW YORK (c.1925)

571485

20

By the eighteenth century, you start to see something like our modern notion of information, circulating, shaping minds, tied to democracy. There's enormous optimism at first. In 1778, someone proclaims an 'age of information,' convinced that if you give people the right information, they'll do the right thing, very Enlightenment. But almost immediately a counter-current appears: suspicion and anxiety. Who shapes this information? Who's behind it? Some fall back on 'good information comes from God,' which is a rather desperate move, but it shows how troubling the issue is.

Jefferson captures the swing perfectly. Early on, he argues that a free press is essential, that he would rather have newspapers without a government than a government without newspapers. A few years later, he complains that the man who never looks into a newspaper is better informed than he who

reads them. That rapid flip, from faith in circulation to fear of manipulation, is crucial.

So yes, empiricism and modern tech culture share a pattern. They begin with promises of openness, neutrality, access; but they quickly confront questions about who can speak, who is believed, and what counts as credible. That was very much on my mind writing about the early web, utopian talk of democratization, then, by around 2012, a crash into disillusion. It felt very eighteenth century to me.

Part of the problem is that we keep telling linear stories, this leads to that, progress all the way down. The cycles of enthusiasm, abundance, anxiety, and disillusion get written out. If we thought more cyclically, we might be less prone to the naïve belief that this time it's all new, all progress, and all we have to do is lie back and let technology carry us forward. That's a trap we've fallen into more than once.

You contrasted Boston’s closed, hierarchical model of innovation with Silicon Valley’s more fluid one. That sounds very much like Ron Westrum’s ‘pathological’ versus ‘generative’ organizational cultures. Boston hoards and suppresses information; the Valley (at least historically) lets it flow across roles and institutions. Does that comparison fit your experience?

Broadly, yes, though I’d add that Silicon Valley is, sadly, drifting towards the more pathological model. Still, there’s some residue of that earlier, more generative culture.

When organizations treat information as something to be tightly quantified and controlled, they tend to build rigid hierarchies. You see it in principal–agent models, in scientific management, and in the economic thinking that underpinned the Boston firms AnnaLee Saxenian describes. Those structures helped them become powerful, but they also constrained their ability to adapt.

In the Valley, by contrast, mobility and networks mattered. People moved between companies; ideas moved with them. Informal channels, corridor conversations, meet-ups, spin-offs, were crucial. That loosened the grip of any one hierarchy and allowed for more bottom-up innovation.

Now NDAs, intellectual property battles, anti-poaching agreements, all the mechanisms of control, are eroding that fluidity. Saxenian herself worried she was documenting something already disappearing, and in some ways she was right. Yet some of that generative culture has proved more resilient than expected.

My concern is that as we shift back towards hierarchy and as information gets narrowed to what fits an economic template, what you can model, own, measure, we undermine the collaborative, cross-disciplinary experimentation that made the Valley distinctive in the first place.



Richard Firth Green

**A Crisis of Truth,
Revisited**

When I wrote *A Crisis of Truth* (1999), I was well aware that my fascination with the cultural changes brought about by a late medieval surge in literacy was driven in part by my own experience with a new electronic medium, albeit just the newfangled word processor on which the book was taking shape.

It was a time when anyone familiar with the work of my fellow Canadian, Marshall McLuhan, who, incidentally, began his career as a professor of English, could hardly ignore the importance of the medium to the message. So, as I wrote about *truth* in the fourteenth century, I felt I was witnessing a comparable shift in attitudes to *truth* in my own day. And I believed Ricardian truth had something fundamental to teach us about that experience.

Even though *A Crisis of Truth* had taken on something of an oracular role in my own mind, I fear that idea was, at best, only faintly present in the book itself. Still, whenever I come across a word like *truthiness*, or read that we're now living in a *post-truth* era, I think: anyone who'd read my book carefully could have seen this coming. Thankfully, my failure to spell it out as a clear message for our times hasn't stopped perceptive readers from detecting it, including, I'm sure, members of this very group.

A quick trawl through the web turns up citations like: 'Truth and Truth-Telling in the Age of Trump,' 'Truth and Truth,' 'Truthiness and Women's Bodies,' and 'Millennial Bashing in Medieval Times.'

The earliest sense of the word *truth* in English was a bargain, a contract. That meaning survives in the archaic form *troth*, as in 'to plight one's troth', usually in the romantic context of promising marriage. And since the ability to depend on the given word of one's fellow citizens can be seen as fundamental to social cohesion, the so-called social contract theory of Hobbes, Montesquieu, Rousseau, Rawls, and others, any disruption in a society's mechanics of trust is likely to lead to anomie and discord.

One such disruption occurred in the Middle Ages and was brilliantly analysed by the historian Michael Clanchy. His book *From Memory to Written Record* was a major influence on mine.

Now, Wikipedia describes *From Memory to Written Record* as a study of the triumph of literacy in medieval England. That's certainly not how I read it. One of its great strengths, in my view, is that it eschews any UNESCO-like triumphalism and instead soberly documents the social upheaval caused by the incursion of the written word into daily life.

In Clanchy's thirteenth century, the time of Magna Carta and the Crown's inquiries into landowners' legal rights, this upheaval primarily affected the landed classes. But by the fourteenth century, as I tried to show, it had spread across society. A late echo of it can be heard in the Peasants' Revolt of 1381, in Margery Starre's cry: 'Away with the learning of clerks, away with it!'

This general upheaval, I argued, drove the poets of Richard II's reign to focus on the polysemous word *truth*, and the new meanings it was beginning to take on. Chaucer's Franklin, in *The Canterbury Tales*, says, 'Truth is the highest thing that man may keep.' *Piers Plowman's* pilgrimage is a pilgrimage to Saint Truth. And Gawain's pentangle, the symbol on his shield in *Sir Gawain and the Green Knight*, is 'in tokening of truth.'

All signs that *truth* had become a fourteenth-century keyword: a term under social pressure, undergoing consequential semantic development. Because the success of oral truths depended on the reliability, the oathworthiness, of the parties involved, *truth* early on developed

a sense of trustworthiness, dependability. We still hear it in expressions like ‘as true as steel.’

‘He has plighted me his truth’ could mean either: he’s made a contract with me, or he’s assured me of his trustworthiness. But as written contracts gradually displaced oral pledges, this earlier meaning began to give way to the extra-human quality of objectively verifiable truth, something literally (I’m tempted to say literally literally) embodied in a sheet of paper or a piece of vellum, rather than in a handshake.

Truth’s dominant modern sense, conformity to fact, first emerged in English in the fourteenth century. I’m curious how this developed in other languages, perhaps in India. You might be able to fill that in for me.

A similar technological shift has happened in our own times. We live in the age of *truthiness*, as some have called it. I’m regularly asked by the checkout clerk at my grocery store, ‘Do you want a receipt?’

I just as regularly answer no. The first time it happened, I remember thinking: but what if someone accuses me of shoplifting as I leave the store? Then, after a moment’s thought, I realized that the entire inventory of my shopping cart was already in the store’s computer. Evidence that I’d paid was preserved by my credit card provider.

The paper receipt, once the replacement for a handshake, had, in its turn, been made redundant by a string of electronic code. It’s a trivial example, of course, but it’s not hard to find more serious ones.

THIS CHAPTER: IBN BULTAN (11th CENTURY),
TACUINUM SANITATIS (MS c.1380, NORTHERN ITALY)



Not long ago, the Canadian Broadcasting Corporation (CBC) reported on the case of a Langley, B.C., man named Brian Pike. In 1991, Mr Pike bought two guaranteed investment certificates from Scotiabank, worth \$1,500 each. I actually sent you a picture of them, you might still have it. He was given two official-looking documents, those are the ones in the picture, with serial numbers, imprints of their value, and a teller's stamp. There used to be a photograph of them on CBC's website, and I've got one saved to my computer.

Mr Pike was, clearly, an impeccably absent-minded man. He forgot all about the certificates, but when he came across them gathering dust in a drawer 28 years later, he trotted off to the bank to redeem them, only to be told the bank had no record of them in its computers. The official certificates he was waving about were, they said, worthless.

The medievalist in me immediately thought of the story of John de Warenne, Earl of Surrey, waving a rusty sword before one of Edward I's *quo warranto* proceedings. A medieval truth plight involved more than a simple oath and handshake, it also required the *traditio* of a symbolic object, a token, however rusty, to seal the bargain in the memory of witnesses. Our exchange of wedding rings today is a survival of that very archaic ceremony.

When a royal commission asked Warenne for written proof of his right to certain titles or land, he must have been as astonished as Mr Pike was. That sword was a charter sword, it had been in his family time out of mind. To him, it was all the evidence he needed.


But of course, what Warenne felt was being challenged wasn't just a piece of metal. It was his word. It was the word of his ancestors. It was his family's honour. His truth against theirs.

At around the same time Mr Pike was arguing with Scotiabank, another Brian, Brian Chesky, CEO of Airbnb, was dealing with his own crisis. Someone had rented a house in Orinda, California, a quiet neighbourhood described by *The New York Times* as a white-collar bedroom community, home to tech executives, bankers, and Berkeley professors. The renter held a Halloween party with over a hundred guests. Things got seriously out of hand. Gunfire broke out. Five people were killed.

Here are a few lines from the letter Mr Chesky sent to Airbnb employees. I sent you the full email beforehand, but these are selected lines:

Airbnb is a business fuelled by trust. Our real innovation is not allowing people to book a home. It's designing a framework to allow millions of people to trust one another. Trust is the real energy source that drives Airbnb. We believe that trust on the Internet begins with verifying the accuracy of information on Internet platforms. The world moves at the speed of trust. The more trust that exists, the more access we can all have. Airbnb is founded on trust. Our vision depends on us continuing to increase this in our community.

If the occasion hadn't been so tragic, one might almost relish the irony: a company complicit in an institution that's steadily undermining trust lamenting its loss.



There's a similar irony in a petition sent by the Mercers' Guild to Parliament in 1386. They complained that the last mayor of London had threatened some of their members with being drawn and hanged, and requested an inquiry so that 'the whether that trowthe may shewe, for trouthe amonges vs of fewe or elles no man many day dorst be shewed.' They claimed this was because the current mayor had thrown the city's written constitution into the fire.

In reality, it was the spread of literacy that was threatening the Mercers' trust, just as surely as internet surfing is now undermining Airbnb's trust.

Ironically, we are living through a reversal of the Ricardian crisis of truth. Yet our own crisis of trust, to borrow Mr Chesky's term, is no less critical. We're entering an era of secondary orality, in which the semi-literate are being encouraged to put more faith in facile internet posts than in the carefully reasoned written depositions of legal authorities.

As we reflect on the disrepute into which the U.S. Supreme Court has fallen, we might recall that Sir John Cavendish, Chief Justice of King's Bench, was killed by a mob during the Peasants' Revolt in 1381.

As Mr Chesky says, we have to find a way to trust the internet, or, better yet, to make the internet trustworthy, just as we once had to learn to trust written documents. Unless we do, we won't just see more Airbnb shootings. We'll see more Margery Starres on the steps of the U.S. Capitol, or, dare I say it, more Marjorie Taylor Greenes within the Capitol itself.

One scholar has argued that, during the transition from the medieval to the Renaissance, or early modern, period, truth shifted from being associated with personal standing, where it had multivalent meanings, to a more Enlightenment idea of truth as objective fact. That earlier sense persisted, though, for example, in Montaigne, 'I may indeed contradict myself now and then; but truth, I do not contradict'. And today, we hear phrases like 'I'm speaking my truth' in internet and popular-culture contexts. Do you think, in a sense, we've come full circle?

Yes, I actually do think that. And, as I said earlier, I'd attribute it to a profound shift in the medium of communication. It's a very current issue, of course. You've got Trump supporters saying things like, 'I don't care about the facts. If my leader says this, then that's my truth.'

That way of thinking is, in some ways, quite medieval. If Thomas Aquinas said it, it must

be true. There was a great cult of authority in the Middle Ages, and it seems to be making a return.

We think of the Enlightenment as having undermined that kind of uncritical reliance on received authority, that we should instead establish truth through reason. But that seems to be disappearing now, at least in the public sphere.

It also seems to depend on what kind of authorities we choose to trust. Michael Gove, during the Brexit debates, said, ‘The people of this country have had enough of experts.’ So, on one hand, there’s a rejection of authority, but on the other, we still look to some kind of authority, just different ones. It’s a bit contradictory.


We are, in some ways, expert-ridden. One of my great heroes, Noam Chomsky, complained long ago about the cult of expertise. What it does is fragment knowledge: this one’s an expert on this, and that one’s an expert on that.

But we rarely have public intellectuals now who try to speak to society as a whole, who attempt to draw connections, to synthesize. Instead, we get isolated domains of expertise, bombarding us from all directions.

I’m absolutely with Chomsky on the dangers of that. It does, as you say, replace the kind of awe for authority we saw in the Middle Ages. Then, people revered the great classical authors, Cicero, Ovid, Virgil, not just for what they said, but for who they were.

Maybe it’s a basic human tendency, we want someone to tell us the truth. But we’ve really fragmented our idea of authority in the twenty-first century.

Part of the problem lies in the dominance of the social sciences. You’ve got sociologists who study society, psychologists who study the mind . . . But the old philosopher, the thinker who reflected on everything, or the poet, the playwright, the artist: these figures are being devalued, precisely because they don’t possess narrowly defined expertise. And yet they’re the ones who think more broadly about general problems.



In the sixteenth century there's an eruption of conflict between the Court of Common Pleas and Chancery, where the latter's equitable jurisdiction can interpret not just the letter but the intent of the law. Lovedays—local, socially grounded arbitration and reconciliation meetings, typically convened within the community and overseen by neighbours, elders, parish figures, or local gentry, aimed at restoring social peace through agreed settlement (often with staged payments, apologies, handshakes, sureties, or other undertakings) rather than issuing a coercive judgment—seem to occupy a kind of informal adjudicatory space alongside (and sometimes in the shadow of) formal courts. Once lovedays start to disappear (apparently sometime after the fourteenth century), does that space get filled by the Court of Chancery and its more interpretive, equitable functions?

Lovedays must have started out as local dispute-resolution mechanisms, and I imagine they worked quite effectively in small communities. But by the late fourteenth century, there's definitely evidence that they were being abused.

I've been thinking about this recently while working on a book about *Gawain and the Green Knight*, or more specifically, the manuscript that includes *Gawain* and three other poems. It was almost certainly written in Cheshire, in the north-west of England, and I've been looking at Richard II's connections to Cheshire in the late fourteenth century.

He had what you might call a bodyguard of Cheshire thugs, local enforcers. There's a report, from someone clearly hostile to them, about how they spoke to the king like he was one of them, a mate.

After Richard married a French princess, they're said to have remarked: If you'd only married Perkyn, daughter of Lye, you could have had a loveday with anyone in Cheshire. The implication is that, if Richard had aligned himself with a local power structure, he'd have had the muscle to intimidate anyone into submission at one of these lovedays.

So by that point, lovedays weren't about justice, they were about coercion. If you had power, if you had force, you could stage a loveday and pressure your opponent to accept whatever you wanted.

As for what replaced them, yes, local jurisdictions were weakening by the fourteenth century. County courts were losing authority. Things were centralizing. You do see the rise of the Court of Chancery and equitable

jurisdiction, which eventually offered a less rigid alternative to common law. But even so, by Dickens's day, the Chancery was infamous. And access to it was limited: you still needed money, influence, and literacy to take a case to Westminster.

At the village level, it's harder to trace, because the documentation is patchy. But manorial courts continued, and for generations they seem to have been the only real option for people without resources. In towns, you had borough courts too.

But the law became increasingly formalized. Equitable jurisdiction offered some relief, especially from the fifteenth or sixteenth century onwards, but it wasn't always accessible. And there were setbacks too, Morton's Fork comes to mind.

Well, obviously, the consequences of a king breaking his word are much more serious than if a shepherd breaks his. You do find, particularly among the commons, a reliance on ambiguous or equivocal oaths and promises.

In the Middle Ages, when literacy was far less widespread, speech had a kind of power that it rarely has today, at least among literate people in the West. So you could, for example, say something ambiguous, and if someone was foolish enough to believe you, you were off the hook.

There's a scene where Tristan swears an oath to clear himself of sleeping with the king's wife. He has done that, but he manages to swear an oath that's literally true. Or rather, it's his lover, Isolde, who swears that no one has held her in his arms except the king and 'that old beggar who just carried me across the ford.' Well, the beggar is Tristan in disguise.

The oath is technically true, and God honours it, she carries a red-hot iron and isn't burned. So there's enormous respect for literal truth, even in moments of calculated equivocation.

I think, even at the commoner level, people understood that literal truth was the gold standard. But if you could phrase something ambiguously, and someone misunderstood it, that was their problem.

'When kings or rulers make rash promises, they must be held to their word, but when a miller or shepherd makes one, it needn't be taken literally.' Where does that idea come from, that rulers must be held to their word, but commoners less so?

Did the medieval mind often treat rash promises or metaphorical speech as reflections of some kind of inner truth, something separate from intention?

Intention is a very different thing in the Middle Ages. The Church, priests and schoolmen who set out rules for confession, absolution, penance, did recognize the importance of intention. Intending to murder someone was a sin, even if it wasn't as bad as actually doing it.

But among common people, intention wasn't so clearly articulated. You see this in early, or 'primitive', law: it barely registers. It comes in more fully as law becomes more literate, more intellectualized. That's one of the peculiar things about Anglo-Saxon law, for instance, how little emphasis it places on intent.

In fact, even in Victorian law, there was the idea of a *deodand*. If someone was killed by an inanimate object, or a non-human agent, like a horse, then the object itself was forfeit. The victim's relatives were entitled to the horse, say, because the horse had caused the death.

The idea was that the object had somehow committed the crime. In theory, it was forfeited to the Crown, but in practice it could end up compensating the victim's family. That practice survived into the nineteenth century and was abolished in 1846.

What finally ended it? Railway accidents. People started getting killed by trains, and the railway companies pushed back: 'We can't give away the value of an entire locomotive every time someone's run over.' At that point, the law changed.



That's marvellous. It brings us back to law. One thing I discovered while researching medieval law is how fundamental the division between tort and contract is in common law. It goes right back in English legal history.

But when you think about it, there's no strictly logical reason for that split. If you break someone's word, that's a private wrong. So why distinguish between private wrongs and breaches of contract as two separate legal areas? And yet the distinction survives all the way through.

An interesting document about inanimate objects and accountability, from 1979, is an IBM training manual, a line of which read 'A computer can never be held accountable. Therefore, a computer must never make a management decision.'

The IBM example shows people wrestling with precisely this tension, between contractual obligations and tortious harms.

What I find especially worrying is how modern American law has gone all-in on contract. On the face of it, contract seems fair, it's based on mutual agreement, after all. The idea of the 'social contract' rests on that assumption: both parties agree, so all is well.

But that only works if both parties have roughly equal power. If one party is far more powerful, say, a corporation and an employee, then the contract is

almost certainly going to be a bad one. And in the U.S. legal system today, contract law dominates in ways that are deeply troubling.

Take nondisclosure agreements, for example. A woman is sexually harassed by her boss. That's a tort, it's a civil wrong. But instead of seeking justice through tort law, she's offered money in exchange for signing an NDA promising never to speak publicly. So contract is now being used to undermine tort. And that's happening in many areas. It's a very dangerous trend, because it concentrates even more power in the hands of the wealthy and well-connected.

Your IBM example raises the related question of computer 'authority': can a computer be held accountable for a decision? Presumably the programmer or manager would bear responsibility, but the contractual authority granted to algorithms or machines is a fascinating and unresolved issue.

Ultimately, it reminds me of the deeper tension in the evolution of law, especially in the way contract emerged, historically, from actions that look much closer to tort in early common law.



CADAVER LAB, HALIFAX MEDICAL COLLEGE, NOVA SCOTIA (c.1890)

**Daryn
Lehoux
& Sergio
Sismondo**

**Epistemic
Corruption**

A Dialogue

DL. Sergio and I aren't, on the face of it, the most obvious people to collaborate on a project like epistemic corruption—we work in very different periods. But we kept finding ourselves talking about shared interests, and at some point, we decided to work on something together. It turned out there was quite a bit of overlap in our work, even though Sergio focuses on modern pharmaceutical marketing and science, and I work on much older material, the ancient world, far removed from questions of marketing or certain sociological concerns.

I don't have the evidence to ask the kinds of questions Sergio can ask in his period. But we're both really interested in how people perceive facts—how they're manufactured, how they come into being, and how they move through the world. In my case, I'm also fascinated by how they disappear.

When I teach my second-year course on ancient science, I always begin by telling students that everything they're going to learn is wrong. Every fact, every argument, every cosmological model—they're all things we no longer believe. That, to me, is a really interesting phenomenon. I want to know where those facts went, why they looked like facts at the time and how they stopped being facts.

SS. And I take a step back from individual facts and look more at knowledge systems—how knowledge is produced and distributed. What I'm interested in is how knowledge systems come into being and decay—how they fail to do what we expect them to do. In this case, the systems for producing medical knowledge, especially about pharmaceuticals, get co-opted and 'corrupted'. That's the term Daryn and I use. They're redirected by specific interests—usually drug companies—so the knowledge produced ends up serving those interests. It's not just information released into the world; it's knowledge shaped for particular purposes.

Daryn's work, I think, also speaks to the decay of knowledge systems—or at least constellations of knowledge. Take something like the sympathy–antipathy framework. That's not a single fact—

it's a whole system that structures what counts as a fact, and what kinds of questions people think to ask. That's quite analogous to what I study in pharmaceutical knowledge production.

- DL.** I don't have access in my sources to the sort of systemic forces that you can trace in yours. But you're right—the questions I ask often deal with how facts get grouped into ways of understanding the world. Theories, in a broad sense—frameworks that cluster facts together. Why do certain facts seem to belong together? Why do people link them the way they do?

Could you walk us through the sympathy–antipathy framework in the ancient world?

- DL.** The central paradigm for me is the curious phenomenon of garlic–magnet antipathy.

It first caught my attention over twenty years ago when I was teaching ancient science. I came across this claim in a text—Ptolemy's *Tetrabiblos*, I think—and it struck me as both amusing and odd. He's trying to explain how medicines work, particularly how wounds heal, and he says something like: Well, we all understand this mechanism—it's just like garlic and magnets. If you get garlic on a magnet, the magnet stops working.

He brings in that claim not to argue for it or introduce it, but as if it were a piece of shared factual knowledge, just something everyone accepts. It's like me saying: We all know water is H₂O. It's treated as obvious, a matter of fact.

So I started digging around, trying to see where else this idea shows up. Turns out it's fairly common in ancient sources. An early reference is in Plutarch. In a banquet scene, someone tells a story about a fish—a remora—that sticks to the bottom of ships and makes them stop. Now, remoras do stick to sharks and other things, and the Latin word *remora* even means delay or hindrance. But in the dialogue, some characters laugh it off—Oh, that's just a silly story. But someone responds: No, no, I saw it happen once. And now we have a claim to experience.

That's what makes it interesting. Is Plutarch just having fun, or are we supposed to believe this character genuinely thinks he experienced this? To support the story, another person chimes in: Well, that's not so surprising—we all know garlic ruins magnets. Again, it's brought in as a self-evident fact to support a more controversial one.

That's what got me hooked. I read 'garlic ruins magnets' and immediately thought, 'That's ridiculous.' But then I asked myself—why? Why does that feel so obviously false to me? How do I *know* it isn't true?

My first answer had to do with experience—just like Plutarch's character claims. I've handled garlic, I've played with magnets. I've never seen them affect each other. But as I dug deeper, I realized it had more to do with classification—how we mentally sort and relate different things. How we think about what garlic *is*, what magnets *are* and what kinds of interactions we imagine are possible between them.

That classification work tends to hide behind a veneer of experience. I *think* I know garlic because I've cooked with it. I *think* I know magnets because I've fiddled with them. But actually, what's doing the work in how I judge that garlic and magnets don't interact is the classificatory framework I'm relying on—what kind of thing I think each one is.

That's what I think is happening in Plutarch, too. Garlic and magnets are treated as kinds of things that belong to a broader sympathetic or antipathetic system. They're imagined to operate in the same mode, so the interaction seems plausible within that worldview.

This, in turn, shapes how people see the world—not just in antiquity, but also how I see it. It's something I try to reflect on in my work: where my own assumptions come in, where I might have blind spots, and what I might be missing in how I approach epistemology.

Ancient sources are great for this because they're often radically different from how we think today. The contrast throws things into sharp relief. If I'm talking to someone in the office next door, we probably share enough of a worldview that small disagreements don't really illuminate much. But when you go back in time or across cultures, the differences become stark, and certain phenomena suddenly stand out in really productive ways.

If I've understood you correctly, your central claim is that our sense of what's plausible is shaped more by our classificatory systems than by raw experience. It's not direct observation that tells us garlic and magnets don't interact—it's how we've already sorted those things conceptually.

DL. Exactly. And I think a key part of that is how unaware we are of it. We tend to think we're drawing on specific experiences, when really it's the classificatory system doing the work. But in our minds, we give credit to experience—as if we've seen it with our own eyes—even if we haven't.

Right, it just feels obviously true—it fits so well with what we think we know that it must be true.

But here's what I'm wondering: some examples of sympathy and antipathy could, in principle, be tested. Why don't those more verifiable examples become the common tropes? Why does garlic and magnetism—a claim most people likely hadn't witnessed—become the one that circulates?

SS. That's a good question. I'd be cautious, though, about saying that no one had seen it happen. Someone might well have *thought* they saw it, the same way someone in Plutarch's dialogue says they saw a remora stop a ship. So I wouldn't dismiss that entirely.

But yes, your broader point stands. Why don't the more obviously demonstrable examples become the standard ones?

DL. It's possible they did. Garlic and magnets just happens to be the one I focused on because it has such a long and visible history—you can really trace its trajectory. But there are others. There's also the ancient claim, in Pliny, that goat's blood can break up even diamond.

If I'd come across that first, I might've built my whole story around it. But garlic and magnets was the one I saw initially, and it has this lovely, winding path in the historical record. You see it repeated over and over as a fact, then you start to see doubts emerge. One person says, 'I've seen it ruin a magnet'; another says, 'I've tried it—it doesn't.' You get this back-and-forth.

That's where Sergio's earlier point really hits: what exactly are they seeing?

One of the amusing things for me is that I've never tested it myself. I still haven't. Which is funny, considering how often I've talked about it over the years.

Although, once, some students ambushed me. They wanted to demonstrate it for real. They brought in fridge magnets and said, 'Look, it still works with garlic.' I had to tell them, 'That's not a fair test—those are ferrite (ceramic) fridge magnets. We're talking about lodestones. Natural magnets.'

So they went away. But one of the students had keys to the geology department. He borrowed a lodestone from their display cabinet and brought it into class, reeking of garlic. He showed me it still picked up paperclips.

And, well, what could I say? I didn't watch him rub the garlic on it, but I could smell it, and I trusted him. Still, the magnet was very weak. A natural lodestone is nothing like the powerful neodymium ones I have in my office. It barely did anything.

So yes, it's possible—as Sergio said—that someone, at some point, had a real (or at least plausible) experience that convinced them garlic weakens magnets. Maybe they thought, 'Oh, this used to work better.' The difference might have been marginal. I don't think every experimental claim from antiquity is disingenuous.

And there's also the fact that natural magnets were rare. People didn't have magnets all over their homes like we do now. I've got them in phone chargers, cases, latches—you name it. But in the first century BCE, a lodestone would've been something special. So perhaps people avoided exposing them to garlic just to be safe—'Let's not take any chances with this expensive, fascinating rock I like showing off.'

- SS.** That's what's interesting with the diamond and goat's blood example too. These are valuable things—you don't want to risk damaging them. I think there's a similar claim that dog saliva will dissolve gold. I once brought one of my dogs to class and joked, 'It's not rabid, but we could test whether dog saliva really dissolves gold.' And, unsurprisingly, no one volunteered their jewellery to find out.

DL. Another place I've seen this kind of phenomenon is in the context of spontaneous generation. There's a recurring set of examples that people working on biological generation used to explain why certain things seem to appear regularly.

In a world where spontaneous generation is possible—where animals emerge from rotting fruit, for instance—there's a theoretical problem for those working within an Aristotelian framework. According to Aristotle, if there's no seed or organizing principle in the material, then the animal should arise randomly. So why do specific creatures emerge from specific substances?

To solve this, ancient thinkers developed and circulated a set of regular examples. They say: bees come from cattle (*bugonia*), wasps from horses (*hippogonia*), and sometimes beetles from donkeys. If a cow dies and rots, it produces bees—not wasps. If a horse dies and rots, you get wasps—not bees. These examples function almost like tropes—metonyms for natural regularity. They let people argue that spontaneous generation isn't random, but ordered. 'As we all know,' they'd say, 'bees come from cattle.' You see those claims repeated again and again over the centuries. It's a similar kind of move, though in a different context.

Would it be fair to say that one feature of a trope like this is that people tend to explain away contradictory observations rather than question the trope itself? Like, if a magnet still works after exposure to garlic, you might say, 'Well, that must not be a real magnet.'

DL. And we're very good at doing this. That's one of the things I find fascinating—how readily people explain away contradictions when they're invested in a belief. They assume something must have gone wrong in the experiment, not in the idea.

SS. And we do that all the time. Every day, we encounter things that don't quite fit the taxonomies we rely on—not garlic and magnets, maybe, but broader conceptual categories. And we explain them away. We're remarkably efficient at that.

Can we extend this classificatory theory to include beliefs in things like gnomes and fairies, or perhaps Pliny's Monocoli, the one-legged, umbrella-footed people of India? Is the garlic–magnet story part of the same spectrum?

SS. Well, I'd say gnomes and fairies—or things like them—still have a place in the world. For example, a small number of North Americans genuinely believe in lizard people—shape-shifting elites who control the world. For them, those figures play a similar role. They help explain power dynamics, hidden forces, things happening behind closed doors.

So yes, elves and fairies might not be gone—they've just changed form. Today, they might be spies, deep-state actors or alien overlords. They still help people make sense of unseen influences in their lives. And there are still people who say they believe in various kinds of gnomes and fairies.

But that's where I'd want to interrogate what we mean by *belief* in those cases. We often use the word casually, as if people have worked it through and decided, 'Yes, I believe in elves and fairies,' with a clear, systematic rationale. But I think these objects—these beings—function at a different level than the standard philosophical examples of belief, the reasoned ones.

They're often more playful, maybe more subterranean. They operate through other explanatory modes—symbolic, narrative, emotional.

DL. And you mentioned earlier that good historical practice involves taking people seriously when they say they believe something. We shouldn't just dismiss those beliefs as silly. We need to engage with the world that allowed, enabled or even encouraged people to hold them.

I didn't do that with the garlic and magnets story because I was asking a different kind of question. But if a student asked, 'Was Plutarch just gullible?'—I'd defend him to the end. There were good reasons for him to believe what he believed, and we should explore what those were.

I do the same thing with astrology. In my period, and well beyond it—even into the early modern period—many scientists believed in astrology. Up until Newton's time, it was widely accepted. People thought there were sound physical reasons for it. Today, of course, it's not part of mainstream science. You're not going to land a tenure-track job or win a research grant for inves-

tigating the effects of star positions on personality traits. That's completely outside institutional science now.

But it wasn't always. And I'll argue passionately that there were good reasons to believe in it when people *did* believe. I think that's in line with what you were saying—that we have to take historical beliefs seriously, even when they now seem implausible.

And yet, even when we try not to dismiss something like belief in elves or witchcraft, our sense that it isn't true inevitably shapes how we interpret it. If I'm truly open-minded about whether something exists, I'll ask whether it's true or false. But if I start from a position of scepticism, I'm not asking whether elves exist—I'm asking what kind of worldview could make room for them.

So even if I'm not dismissing the belief as silly, I'm still reading it through the assumption that it's false. That must shape the kind of understanding I come to.

DL. Yes, I think that's about right. It's something I've been grappling with myself as a methodological problem. How far can I actually go in suspending disbelief when I study ancient science? Can I fully immerse myself in these older ways of thinking—or not?

There are places where I *don't* want to do that. And this ties into the epistemic corruption project that Sergio and I are working on. For example, in Book One of Aristotle's *Physics*, when he explains that the Earth must sit still at the centre of the universe—because heavy bodies fall and light ones rise, and planets orbit around it. I think he had plausible reasons to accommodate that belief in his worldview.

But then we get to other parts of his work like the *Politics*—for instance, where Aristotle claims that some people are *naturally* slaves. Should we similarly say, ‘Let’s think our way into this and see how it makes sense’? I don’t want to do that. There’s a moral hazard in it, and I avoid it.

That’s clearly a methodological inconsistency on my part. I’m stepping away from the model we just described for elves and fairies. But it’s hard to fully enter a worldview when doing so feels morally or politically dangerous.

Some of these beliefs still have real-world implications today. We’ve spent centuries trying to move beyond the notion that some people are inherently inferior or born to serve others. That’s not ancient history—it’s a live ethical issue with real consequences. So I’m not willing to go back there, even for the sake of empathetic reconstruction. It feels too raw. That’s not quite the right word, but it’s still morally charged. And that makes it different from something like geocentrism, which no longer carries that kind of ethical weight.

I think another issue arises when we’re studying what we’d now call pseudosciences—phrenology, for instance. Whether or not the historian believes the object of study is ‘true’ changes the starting point of the inquiry.

If I’m studying phrenology, I’m likely to focus on its ideological underpinnings. But if I’m studying something like Darwinism in the nineteenth century, I might still acknowledge ideological factors, but I’m also going to say, ‘They believed in this because it made sense—because it worked.’

Whereas with something like homeopathy or eugenics—even if the belief is ideologically neutral—it’s treated with a degree of suspicion that affects how we frame our analysis.

DL. This underlies the Strong programme in the sociology of science, which calls for methodological symmetry: when we study historical sciences that we *now* think are true, and those we *now* think are false, we’re not allowed to use different types of explanations for each.

We shouldn’t say, ‘People believed in the true stuff because it worked,’ and ‘People believed in the false stuff because they were misled or ideological.’ Methodologically, we’re supposed to treat all beliefs in the same way—to explain *why* people believed them using consistent categories, regardless of whether we now judge them true or false.

So if people historically believed two different things, and I now think only one is correct, I still need to account for both beliefs in symmetrical ways. Symmetry means we explain both what we now consider rational and what we now consider irrational using the *same kinds* of explanatory tools.

So, is it fair to say that we can’t explain people’s belief in a scientific system by saying, ‘They believed in it because it worked’? Because we can’t apply that same reasoning to beliefs we now consider false?

DL. I think that's right. But we *can* say they believed in it because it *appeared* to work, or because it *seemed* to them to be effective. We just have to soften the claim a bit. That way, you can treat phrenology and Darwinian evolution using the same kind of explanatory framework.

But if you say people believed in Darwinism *because it's true*, and you can't say the same for phrenology—since we now think it's false—that would violate the symmetry principle.

SS. And 'working' is a tricky, complicated category. When does something count as working? We've all heard the old programmer's line, "That's not a bug, it's a feature."

Things can appear to work or not work depending on how we define success. Take the pharmaceutical example: there are some who genuinely believe opioids 'work' very well. And that's because their definition of 'working' differs from those who believe opioids are ineffective in the long run or harmful.

Could we fit the COVID vaccine debate into this framework? With vaccines, the category of 'working' tends to get even murkier, because it is difficult to ascertain whether it is immunity, individual or herd, or the vaccine that renders you safe from the disease.

DL. What's interesting about the vaccine debate is that the arguments for efficacy are statistical. They're about large populations, number-crunching. And with any vaccine, there's always a risk: some people will get sick from it, some may even die. The question becomes, on balance—do the benefits outweigh the harms?

I had an interesting encounter just last week. I overheard someone loudly insisting they wouldn't get another vaccination. She said, 'Last time, I got sick for days. And look around—no one's getting COVID anymore. It's just like a cold.'

So she was arguing from personal experience. And I believe her—she probably *did* get sick from the vaccine. Maybe it was a coincidence. Maybe it was unrelated. But for her, that experience was real.

And honestly, I see what she sees. Most cases I've heard about recently have been pretty mild. Some are worse than others, sure—but we're not seeing people being intubated, we're not seeing refrigerator trucks in Manhattan. It's not like the early pandemic.

So she's offering an experience-based claim that I still want to say is wrong. But what's my counterargument? It's a statistical one I don't have to hand. I just trust—on faith, really—that when public health authorities say, 'Yes, you should get vaccinated,' it's because someone has done the analysis.

And that means I'm placing my trust in a system, in institutional science, over my own experience—or hers.

It's also strange, isn't it? People take medicines all the time, and if one doesn't work, they don't say, 'Well, I'll never take medicine again.' But with vaccines—if one doesn't seem to work, or makes someone feel ill—you often hear, 'That's it, I'm done with vaccines. I'm not vaccinating my kids either.'

Anti-vaxxers rarely differentiate between types of vaccines unless it's something mandated by the government. Vaccines seem to be treated as a coherent category of intervention. But you don't see anyone saying, 'That antifungal didn't work—Big Fungus is out to scam me.'

But this raises a broader question: how we think about background beliefs that shape our scientific understanding. Take the two examples we've been discussing: rubbing garlic on a lodestone, and vaccines.

With the garlic example, most people would scoff immediately—'That can't possibly be true.' But with vaccines, we engage the statistics, weigh the science, consider the mechanisms. So what determines whether a belief gets taken seriously or dismissed out of hand?

You might think it's about how long the belief has been around or how embedded it is, but even recent examples complicate that. Think of thalidomide or the demonization of fats in nutrition science. People accepted those narratives quite uncritically for a long time.

So do you think some beliefs—especially scientific convictions—become consciously argued for, and gain that sort of 'moderate' standing in our minds, while others just sit quietly in the background?

SS. I think there are a lot of complex dynamics at play. Some beliefs become tightly linked to political issues, for example.

Take the COVID vaccine. I think the vaccination effort during the pandemic—even in a small way—contributed to the re-election

of Donald Trump in 2024. In the US, populist sentiment has been strongly aligned with COVID vaccine resistance and with a preference for homegrown or ‘homespun’ solutions.

This ties into a deep distrust of elites. A significant chunk of the American population simply doesn’t trust elite institutions.

There is a fascinating study out of the US: Researchers tested numeracy—math skills—by giving two groups of participants a data chart. One group got a graph about the effectiveness of a face cream. The other group got the same data, but it was about gun laws and crime prevention.

They had already controlled for how good the participants were at understanding numbers. And what they found was striking: regardless of mathematical ability, people interpreted the face cream data correctly. But for the gun laws, people’s interpretations aligned with their political beliefs—not their numeracy skills. Political conviction was the stronger predictor.

DL. What a great study. And I think it also relates back to what I was saying about garlic and magnets—about classification. In politically charged environments, it’s not exactly about classification, but something analogous is going on.

If you’ve got strong political commitments—and in the US, almost everyone does on these issues—you start to *see* facts differently. You don’t just interpret them differently, you actually perceive them through that framework.

With face cream, nobody has a strong commitment one way or the other, so you can just process the data. It's a great neutral case.

- SS.** Especially in highly polarized environments, almost anything can become politicized. The boundaries of the political can expand into domains you'd never expect—things that once seemed entirely neutral.

But then, aren't we left with two different systems for explaining belief? For one kind of belief, we invoke political convictions. For another, we say there's no political variance. That seems like a split.

- SS.** Yes, but we still want to compare apples with apples and oranges with oranges. That face cream vs gun laws study is a good example.

We're not trying to explain *beliefs about face cream* and *beliefs about gun laws* in the same terms. The face cream is a control. What we really want to explain is the *variation* in how people interpret the *gun law* data—why some people read the graph one way, and others read it differently. That variation correlates with political convictions, and that's where the interesting explanatory work lies.

- DL.** I wonder: in the gun law study, are people misunderstanding the data, or are they simply *explaining it away* to preserve a worldview? Those are two very different processes. One is a cognitive failure, and the other is a defensive move.

The study showed that people confronted with the gun law data weren't confused—they wanted more context. Even though the graph was simple and clear, the participants asked: 'When was this implemented? What kinds of crimes are we talking about?' They weren't answering the yes/no question directly; they were trying to reframe it.

It also connects with what you were saying earlier—how we choose which numbers or studies to trust is shaped by our ideological commitments.

Sergio, do you think the shift in attitudes towards cannabis has affected the opioid industry?

For years, marijuana was demonized—associated with addiction, paranoia, criminality. Now suddenly it's seen as therapeutic, even benign. Meanwhile, people who take opioids are viewed more negatively. So are we seeing a political distinction between 'good' and 'bad' drugs?

SS. If we're talking about North America—especially the United States—then yes, definitely.

The opioid crisis began around the year 2000. Purdue Pharma had introduced OxyContin a few years earlier, but it was around 2000 that prescriptions really started to rise. Purdue's not the only player, but they've received the most attention. They did spectacularly well—sold vast quantities of opioids.

And yes, they saw cannabis as a threat. They understood that marijuana legalisation would cut into their market. Cannabis was direct competition.

At the same time, legalisation also undercut the Mexican cartels, who had long dominated illegal drug distribution. So the cartels diversified. They started importing heroin from Afghanistan, and later began growing opium poppies in Mexico. Heroin began flowing north.

And this is important: heroin and prescription opioids occupied different geographies. Heroin gained traction in cities; prescription opioids were more common in rural areas and small towns. That had a lot to do with distribution networks and marketing.

Cannabis was in the background of all this. But eventually, the patent on the original OxyContin formulation expired on 16 April 2013 and public awareness of the crisis exploded. By the early 2010s, people were horrified by the numbers of overdoses and addictions.

As profits declined, pharmaceutical companies scaled back. They effectively ceded the territory to the cartels. And the cartels didn't hesitate—they expanded. They started producing both heroin *and* synthetic opioids.

And then, finally, *they* got pushed out—by Chinese fentanyl producers. So it's been a story of overlapping markets, shifting players and commercial competition at every level.

Would you mind summarizing your work on this—particularly on the business of pharmaceutical knowledge?

SS. My paper, ‘Promoting Opioids’, co-authored by Maud Bernisson and published in *Frontiers in Medicine*, an open-access journal, focuses on the company Mallinckrodt, an opioid producer that was sued, and as a result, a trove of internal documents became publicly available. These are now housed in a fantastic archive, mostly at the UCSF–Johns Hopkins Opioid Industry Documents Archive.

This article explores how Mallinckrodt tried to shape medical knowledge in ways that would lead to increased opioid prescriptions. My co-author and I trace the whole infrastructure behind this knowledge-shaping effort—contracted research firms, so-called ‘key opinion leaders’ (who give talks or publish journal articles), the companies that ghostwrite and shepherd those articles through publication, continuing medical education programmes, speaker circuits—It’s a whole ecosystem.

This paper follows a single, concrete case of what I wrote about in my book *Ghost-Managed Medicine*—which lays out the same kinds of structures, but at a more general level and not focused on one drug. The core argument in both is that pharmaceutical companies invest heavily in shaping what people—especially doctors—believe and know about their products.

This isn’t about whether what they say is *false*. The point is that companies operate behind the scenes, controlling how medical

knowledge is produced and distributed. And this knowledge—these narratives—help make doctors feel more comfortable prescribing opioids. Doctors know the risks, but they're presented with a framework that helps them feel they're prescribing responsibly.

So the real issue here is not falsity, but control. It's about how commercial interests steer scientific and clinical knowledge in subtle ways, without necessarily falsifying it. That's where the epistemic corruption lies—not in outright lies, but in structural co-optation.

Our paper in the *British Medical Journal*, 'How an opioid giant deployed a playbook for moulding doctors' minds', focuses more on problematic conclusions and specific misleading claims.

For example, one of the concepts we highlight is 'pseudoaddiction,' which was promoted by several companies, including Pfizer and Mallinckrodt. Pseudoaddiction refers to cases where a patient *appears* to be addicted to opioids—but the claim is that they're not actually addicted, they're just under-medicated. The proposed solution? Increase the dosage.

The logic here is based on the medical understanding that over time, treating chronic pain with opioids requires escalating doses. But this leads to a fork in the road: if it's *addiction*, you try to *reduce* the dose. If it's *pseudoaddiction*, you increase it. That divergence is crucial—and potentially dangerous.

So the two articles, while overlapping in content, differ in emphasis. The *Frontiers* piece focuses on the structural corruption of medicine—the systems and institutions. The *BMJ* article focuses on the corruption of particular facts or claims.

And you can see the disciplinary difference. *BMJ* didn't want a structural analysis—that's for humanists and social scientists. What they wanted was a clearer medical framing: someone did something wrong, here's the evidence, here's why it matters for clinical practice. That's the mode of argumentation that works in medical journals.

When you talk about 'pseudoaddiction' and 'chronification', is it fair to say that the company was deliberately creating a category—a taxon—under which to package and legitimise a claim that would support their product? Is that a rhetorical move to create a concept that accommodates a profitable fact?

SS. Yes, I think that's exactly right.

Take 'chronification,' for example. The idea is that a patient comes in with what looks like *acute* pain—pain from a recent injury—but there's a concern it might turn into *chronic* pain. The term 'chronification' shapes how the physician perceives that situation. It frames the experience in a way that opens the door to a particular medical intervention—usually more opioids.

From the *BMJ*'s point of view—and I'll put on my clinician hat for a moment—I'd say there's no real evidence for chronification. But as a concept, it's powerful. It fits a certain narrative. We

do encounter patients whose pain starts out acute and then continues long-term. Whether that counts as a disease process or a classification trick is a matter of interpretation.

Something similar happens with the claim that ‘pain is a disease.’ Some pain specialists really push for this framing. Others prefer to see pain as a symptom or consequence of disease. These different framings change how people respond, and how medical decisions are made. So yes—this is very much in line with Daryn’s work on how taxonomies shape perception.

It’s one of those concepts that doesn’t emerge from observation but tells you how to observe something in the first place.

About the Contributors

STEVEN SHAPIN is Professor Emeritus of the History of Science at Harvard University. His work examines early modern science, scientific authority and the social conditions of knowledge-making. His books include *Leviathan and the Air-Pump* (with Simon Schaffer), *A Social History of Truth* and, most recently, *Eating and Being*.

DANIEL ROSENBERG is Professor of History at the University of Oregon. He studies the history of information and temporal/visual forms such as timelines and data, with strong interests in early modern science. He co-authored *Cartographies of Time* (with Anthony Grafton) and directs the National Endowment for the Humanities-supported *Time Online* project.

LISA GITELMAN is Professor of Media, Culture and Communication and Professor of English at New York University. A media historian, she researches American book and document history, techniques of inscription and new media in historical perspective. Her books include *Paper Knowledge* and *Always Already New*; she also edited the volume *'Raw Data' Is an Oxymoron*.

MATTHEW STANLEY is Professor at the Gallatin School of Individualized Study at New York University. He works on the history and philosophy of science, especially modern physics and astronomy and debates about science and religion. He is the author of *Einstein's War* and *Huxley's Church and Maxwell's Demon*.

RANDOLPH C. HEAD is Professor Emeritus of History at the University of California, Riverside. He studies early modern Europe (especially Switzerland), political culture, and the history of archives and record-keeping. His books include *Early Modern Democracy in the Grisons* and *Making Archives in Early Modern Europe*.

PAUL DUGUID is Adjunct Professor Emeritus at the UC Berkeley School of Information. His research examines the concept and history of information, and issues of authenticity and trust in digital and organizational settings. He co-authored *The Social Life of Information* with John Seely Brown and co-edited *Information: A Historical Companion*.

RICHARD FIRTH GREEN is Professor Emeritus in the Department of English at the Ohio State University. A scholar of late medieval English literature and law, he has written on Ricardian culture, court writing and popular culture. His books include *A Crisis of Truth* and *Elf Queens and Holy Friars*.

DARYN LEHOUX is Professor of Classics & Archaeology at Queen's University. He works on ancient science and epistemology, including Greco-Roman philosophy and ancient astronomy, astrology, medicine, and mathematics. His books include *What Did the Romans Know?* and *Creatures Born of Mud and Slime*.

SERGIO SISMONDO is Professor in the Department of Philosophy at Queen's University. He works in science and technology studies at the intersection of philosophy and sociology of science, with recent research on pharmaceutical knowledge and 'epistemic corruption'. He is editor of *Social Studies of Science* and author of *Ghost-Managed Medicine* and *An Introduction to Science and Technology Studies*.



Acknowledgements

We'll keep this short.

We want to thank our contributors for writing their parts; Sejuti Malakar for her help with the design; Lorraine Daston for her insights on objectivity; Otto von Bismarck for sheer administrative poise under absurd pressure; Gregor Mendel for counting things until they became science; Agatha Christie for plot discipline; and the late Franz Reichelt for inventing the parachute coat.

We'd also like to thank Sunandini Banerjee, Naveen Kishore, the Viennese Waltz Committee, the age of sail, the left-handed wicket-keeper of Namibia, the porter from Macbeth, the Society for the Prevention of Unnecessary Alarm, the planet Jupiter and, of course, some of its moons, the International League of People Who Say 'As It Were', the last dodo, the brass section of a slightly overconfident wedding band, Søren Kierkegaard's less cheerful afternoons, and, of course, John Cleese.

FACT OR VALUE is an online forum based in Calcutta, with a focus on politics, literature and intellectual history. Our effort is to trace the genealogies of ideas which inform the inheritance of the human condition.

For queries, please write to:

Doyeeta Majumder: doyeeta@gmail.com

Subha Prasad Sanyal: subhaprasadsanyal@gmail.com

Diven Nagpal: divennagpal1@gmail.com

DESIGN DIVEN NAGPAL

