The Dynamics of Social Organisation

A field guide to dreaming vertebrates



Kari McKern

ComplexityWorkz

Version 1.0 · 31/10/2026 · CC BY 4.0

The Dynamics of Social Organisation

A field guide to dreaming vertebrates

Kari McKern

ComplexityWorkz

The Dynamics of Social Organisation
The Cassandra Edition

Version 1.0 · 31/10/2026
Kari McKern · ComplexityWorkz

This book is free to copy, translate and share under the Creative Commons Attribution 4.0 licence (CC BY 4.0), as long as you say where it came from. creativecommons.org/licenses/by/4.0/

It was not written alone. Claude (Anthropic) drafted it with the author, GPT (OpenAI) was her partner in building the model, and Grok (xAI), Kimi (Moonshot AI) and Gemini (Google) scored for it. The opinions, and the final say, are the author's.

Any errors in this account are accidental, and correction is invited. This is version 1.0; minor editorial adjustments, expansions and deletions should be expected. Latest version, film and data: neuralnations.org/dynamics-of-social-organisation

A print edition and a Kindle edition are planned for Amazon, giving readers the option of contributing to the cost of producing the book. This free edition will remain available.

Contents

A note before you begin **5**

Prologue: Where they left off **7**

Part One: The Naturalist **11**

1. The pack and the prey 11
2. Troops, pods and culture 14
3. Memory leaves the body 16

Part Two: The Polemicist **19**

4. A view from the Moon 19
5. The same failure, everywhere 25
6. A duty to the universe 28

Part Three: The Collaboration **31**

7. First principles with GPT-4 31
8. Eight functions, four readings 34
9. The shape between the functions 38
10. Three machines, one casebook 41
11. The conversation that turned 45

Part Four: The Discovery **48**

12. Everything broadcasts 48
13. The same failures at every scale 50
14. Chronic and acute 54
15. America in plateaus and jumps 58
16. The synchronised slide 62
17. How it could be wrong 65

Part Five: Aliens Are Us **67**

18. The meta-systems 67
19. New wiring 69
20. X-ray vision 71

A traveller's guide to meta-meat brains **73**

Welcome to the Machine 73

Epilogue: Back to the dot **75**

Afterword: What is settled, what is open **77**

Where to read further **79**

Acknowledgements **81**

A note before you begin

This book is in my voice, built from my own words, my published work and the record of my collaboration with machines. This is version 1.0, released on 31/10/2026. Expect minor editorial adjustments, expansions and deletions as I update it at leisure. Any errors in this account are accidental, and I invite correction at neuralnations.org.

The numbers come from the open CAMS record at neuralnations.org, the public repository behind it, and deposits on Zenodo. The blind-test figures and the American series were recomputed from the raw files, not copied from summaries. Where a gap exists in the data, I report it as a gap.

I try to keep three kinds of statement apart. Data is what the instrument recorded. Inference is what I think the data shows. Speculation is where I think it might lead. I say which is which as I go, and I set out the limits once, plainly, in chapter 17 and the afterword.

The longest quotations in the book come from a conversation with a GPT session at the end of October 2026. That conversation is archived with checksums and deposited under doi:10.5281/zenodo.23793977, so anyone can check that I have quoted it faithfully.

It was not written alone. Claude (Anthropic) drafted the book with me from my words and records, checked the figures against the raw data and drew the charts. GPT (OpenAI) was my partner in building the model from 2024, and Grok (xAI), Kimi (Moonshot AI) and Gemini (Google) scored for it. Manus contributed passages to an earlier draft. The opinions, and the final say, are mine.

This book is free. Like the data and code behind it, it is given away under an open licence, Creative Commons Attribution 4.0, the same licence as the Zenodo deposits. Anyone may copy it, translate it and build on it, as long as they say where it came from.

A print edition and a Kindle edition are planned for Amazon, so that readers who want a copy on the shelf, or who would like to chip in towards the cost of producing the book, can do so. This version will stay free.

Prologue: Where they left off

Three men taught me how to look at my own species. None of them finished the job, and each knew it.

On 14 February 1990 the Voyager 1 spacecraft turned its camera back towards home from about six billion kilometres away. Earth came out as a speck smaller than a pixel, caught in a band of scattered sunlight. Carl Sagan built a book around that picture, *Pale Blue Dot* (1994), and its lesson was humility. Every king, every war and every quarrel over a border had happened on that speck.



Earth from about six billion kilometres: the Pale Blue Dot, photographed by Voyager 1 on 14 February 1990 and reprocessed in 2020. NASA/JPL-Caltech, public domain.

I had read my way to Sagan through everything from *Pictorial Astronomy* to *Cosmos*, which I had in hard copy soon after the television series went to air. What he gave me was distance. From far enough away our divisions shrink and our dependence on one another grows. What he couldn't give me was a way to measure what I saw from there. From six billion kilometres you can see that we are one fragile world. You can't see whether the cooperation holding it together is getting stronger or coming apart.

Jacob Bronowski came at us from the other end, close up. In *The Ascent of Man* (1973) he told the human story as a story of knowledge: the hand that shapes the flint, the eye that reads the stars, the mind that builds on

what other minds left behind. His great theme was that we are the animal that keeps learning across generations. Our knowledge lives outside any one head, in tools, drawings, books and the habits of a culture.

Bronowski also understood the danger. In the most famous scene of the series he waded into the pond at Auschwitz where the ashes of the dead had been flushed. This, he said in effect, is what people do when they believe they have absolute knowledge and never test it against reality. The ascent can be climbed down again.

He gave me the mechanism: knowledge leaves the body and accumulates. I believed him. When I was young I believed in progress and the Ascent of Man, and I wanted to be a public servant, running the state on behalf of its citizens, because that seemed the highest calling there could be. What he couldn't give me was a way to tell, while it is happening, whether a society is climbing or falling.

David Grinspoon carried the story into this century. In *Earth in Human Hands* (2016) the astrobiologist argued that the Anthropocene is more than a story of damage. For the first time a planet has produced a species that can see what it is doing to the whole planet, and might learn to steer. He called the mature version of that world Terra Sapiens, a wise Earth.

I knew David for a while. I showed him and his wife my Sydney: Bondi and the Icebergs, St Stephen's in Newtown, La Perouse. In 2016 I was their guest in Washington, and I sat in the library there feeling closer to Sagan than I ever had.

Grinspoon names the stakes. A planet becoming aware of itself has to be able to read its own state. Satellites now read the oceans, the ice and the forests. Nothing yet reads the condition of human cooperation with the same steadiness.

That is where they left off, and where this book begins.

Put the three together and a question appears. If knowledge accumulates outside us, as Bronowski said, and our survival depends on cooperating as one species, as Sagan said, and we must now learn to read the planet we are changing, as Grinspoon says, then how does a species learn to see the state of its own cooperation?

That is not a rhetorical question. It has a shape an answer could take. Something would have to be measured. The measurement would have to be repeatable by people who did not like each other. It would have to work on a family and on an empire, because cooperation is the same trick at every scale. And it would have to be published, because a measurement that only one side can see is not knowledge. It is surveillance.

I spent most of my life without an answer. I came at it first as a naturalist, then as an angry essayist, and finally as the human half of a partnership with machines. The partnership produced a framework called CAMS, the Complex Adaptive Model State. It describes any organised human group, whether a family, a company or a nation, through eight functions that every group needs, each read on four dimensions. It treats the historical record as a broadcast that every working organisation sends out simply by working.

This book tells how that idea came together and what it has found so far. It is also an argument that the idea matters. If the condition of cooperation can be measured, the people taking part in it can finally see it. A planet learning to steer itself would then have one more of the instruments it needs.

Part One: The Naturalist

1. The pack and the prey

I came to natural history through a love of science, and the first animals I studied were my own kind. The playground was a troop I couldn't read. Bullies, alliances and the quick shifts of who was in and who was out ran on machinery I don't have. Before fourth form, I hid in the library and read *Flight International*, the aviation weekly.

At Coogee Primary it felt like being forced into a pen with a troop of a hundred primates, like *Planet of the Apes*. You couldn't get away from them. Yet by the time I was eight I was good at marbles and very good at King Ball. Everyone stood inside a court about the size of a squash court, and the players outside threw a ball at the people inside. Anyone the ball touched joined the throwers, so the outside grew stronger and the square emptied. I was the last monkey in the square.

Science gave me what the playground wouldn't: a way to understand a thing from outside, by how it works rather than how it feels. So I learned to watch people the way a naturalist watches animals, asking what they are doing before asking what they think they are doing.

The question that stayed with me was simple. Why does any animal co-operate, our own kind included, and what does it buy them?

The clearest answer is a wolf pack. A lone wolf can bring down a deer. Against a bison it has almost no chance. A pack can take the bison. The pack doesn't make any wolf stronger. It reaches prey that no wolf can reach alone.



A wolf pack on the move in Yellowstone National Park. US National Park Service, public domain.

That became the foundation of everything that followed. Cooperation enlarges a group's thermodynamic reach. It gives the group access to energy and resources beyond the capacity of any member. The capability lives in the relationship, not in the animal.

I say thermodynamic on purpose. Living things are open systems. They take in energy, do work, and shed disorder as heat and waste. A group that captures more energy can sustain more structure: more young, more territory, more ways through a bad season. When cooperation pays, it pays in energy.

The bill comes due as well. A pack must agree on a direction, defend its range, share the kill, raise the pups and remember where the elk winter. Each of those jobs can be done well or badly, and each can be strained. More wolves don't always mean a better hunt: in a big pack some hang back and let the others take the risk. A pack that hunts brilliantly but fights over every carcass is a pack in trouble.

So the health of a group is not the sum of its animals. It is the state of their relationships. Everything in this book grows from that observation. If the health of a group is the state of its relationships, then a group can be large and weak, or small and formidable, and you cannot tell which by weighing it. You have to look at how the parts are joined.

There is a second thing I took from the pack, less comfortable. Cooperation is not the same as kindness. Wolves cooperate to kill. Chimpanzees cooperate to patrol a boundary and to beat a neighbouring group's males to

death. Humans cooperate to build hospitals and to organise genocides, and the same functions turn up in both. Any account of cooperation that treats it as simply good has stopped watching and started preaching.

What cooperation buys is reach. What it costs is the constant work of holding the relationships together. That tension, reach against the work of maintenance, is the engine of the rest of the book. A society that stops paying the maintenance bill does not lose its reach immediately. It loses it later, suddenly, and usually while its leaders are still describing how well everything is going.

2. Troops, pods and culture

Wolves cooperate. Chimpanzees and dolphins go further: they specialise, and they pass on what they have learned.

In the Taï Forest of Côte d'Ivoire, chimpanzees hunt colobus monkeys high in the canopy. Christophe Boesch and his colleagues, watching over many years, described hunters taking distinct roles: drivers, blockers, chasers and ambushers. The hardest role, the ambush, takes years to learn. Older males do it best.

Chimpanzees also have politics. Frans de Waal's *Chimpanzee Politics* (1982) followed the coalitions at Arnhem Zoo: who backs whom, how a rival climbs, how an alpha falls when his allies drift away. A ruler who loses his supporters loses his position, however strong he is. Any reader of human history will recognise the plot.

And chimpanzees have cultures. In 1999 Andrew Whiten and colleagues pooled decades of field records from seven long-term study sites. They listed dozens of behaviours that vary between communities in ways genes and habitat can't explain: how to crack a nut, how to fish for termites, how to signal a courtship. Each community has its own way of doing things.

Off Shark Bay in Western Australia, some bottlenose dolphins carry marine sponges on their beaks to protect them while they forage on the sea floor. The habit passes mostly down the female line, from mothers to daughters. It is a tool and a tradition at once, carried in a family.

So specialisation, politics, social learning and culture were all on Earth long before us. I want to be careful here, because this is where many accounts of human uniqueness go wrong. Humans didn't invent cooperation, or culture, or the division of labour. We inherited them.

Later I would describe groups through eight functions, each read on four dimensions. If I studied a chimpanzee community that way, most of the functions would be there. The alpha and his coalition carry direction. The patrols that guard the territory's edge carry defence. The skilled hunters carry the craft. Everyone carries labour.

One reading would shrink almost to nothing. I call it Abstraction: how well a function's symbols and ideas are joined to what it actually does. For a chimpanzee there is very little symbolic activity to join. Its culture lives in living chimpanzees. When the last animal who knows a technique dies without teaching it, the technique is gone.

That is the edge of the inheritance. What came next was ours.

3. Memory leaves the body

Something changed in our line, and I believe it was this: memory left the body.

In 2015 a team led by Sonia Harmand reported stone tools from Lomekwi, near Lake Turkana in Kenya, about 3.3 million years old. That is older than our own genus. Every tool like that is a small act of exteriorised memory. A flake struck at the right angle records a skill, and anyone who finds it holds evidence that the skill existed.

I like the Lomekwi flakes because they are so unglamorous. There is no art in them, no burial, no obvious symbol. There is a rock that has been hit in a particular way, several times, in a sequence that had to be learned. Somebody taught somebody. The teaching left a physical object behind, and the object outlasted the teacher by three million years, a kind of persistence no animal has ever achieved with its body alone.

A tool can outlive its maker. It can be copied, and improved by someone who never met the person who made the first one. Once memory starts to live outside the head, cooperation can accumulate. Michael Tomasello calls this the ratchet: each generation starts from where the last one stopped, instead of from scratch. The tools are the proof that the ratchet turned.

The second step was symbols. At Blombos Cave in South Africa, shell beads and engraved ochre about seventy-five thousand years old show people marking things with meaning. A bead can say I belong to this group, or I owe you, or we are allies. The psychologist Merlin Donald called this external symbolic storage, and he saw it as the hinge of the modern mind.

Approval and disapproval, debts and promises, common purposes and shared enemies: once these can be held in objects, cooperation no longer needs everyone in the same place at the same time. It can stretch across distance and survive the death of the people who made the agreement.

This is my foundational claim, and I will state it plainly. What set human bands apart was not the invention of cooperation. It was the redeployment of existing skill, made possible by exteriorising memory. Tool-making star-

ted that, and the tools are the evidence. Symbols of approval, disapproval, cooperation and common ends, held in objects, then let cooperation scale. Natural history, path dependence and plain contingency explain why different societies took different shapes.

Everything we call civilisation is a later chapter of the same move. Writing, law, money, archives, the printing press, the internet and now artificial intelligence are all ways of holding cooperation outside the body so that it can grow. They are also all ways of losing it. A library that burns takes with it capabilities no individual ever held, and the loss is invisible to the generation that follows, because you cannot miss what you never knew was there.

That asymmetry, slow accumulation and fast, silent loss, is one of the reasons I wanted an instrument. A society that is losing its archive does not feel itself getting stupider. It feels normal, because normal is whatever it can still do.

My own family is a small case of it. The McKern line goes back to Hector McKern, born in Scotland around 1760, a wool cleaner who went to Limerick to learn printing. His child settled the spelling of the name, and the McKerns became a printing family in Limerick. Family tradition links us further back to the MacDonalds of Glencoe; the records don’t reach that far, so I hold it as tradition. A nineteenth-century branch in Argentina spelled the name Mackern.

I was born in Sydney in 1953. My father was Norman William James McKern, his father Norman Walton McKern, and the actor Leo McKern was my father’s brother. The family has run to medicine and science: Albert McKern the doctor, scientists at CSIRO and the Museum of Applied Science, an insulin biochemist.

The McKern line leads this story, but it isn’t the whole of it. My father’s mother was Vera, née Martin. Her elder sister Myra married into the Iddles family, and it was Iddles descendants who researched the Martin and Halton families and kept that memory in order.

I mention all this for a reason. A printer's shop, a family tree and a box of letters are the same kind of thing as a Blombos bead. Each holds cooperation outside a body so that people who never meet can still build on it. Without someone keeping those records, I couldn't have written these paragraphs.

One record never left the body at all. A DNA test points my ancestry broadly to north-west Europe, with Irish and Scottish regions among its assignments. That fits the family story without proving it. A regional estimate compares my DNA with reference populations, and it can't name a Limerick printer or a Glencoe clan.

An earlier AI session went further. It joined my paternal line to the Altai, the Bronze Age steppe, Somerled and Glencoe in one continuous route. Checked against the evidence, the route came apart. It had given me the branch that Clan Donald researchers associate with Somerled, which is a different branch from the one I had reported, and my own result still has to be checked against the original test.

That was the method of this book in miniature: a specific reading, challenged with specific evidence. The machine told a tidy story. What made the exchange worth having is that the story could be checked.

Which raised the question I would spend years on. If cooperation is held in relationships, and the relationships leave traces, can we read the traces?

The question is not as strange as it sounds. We read traces of invisible things all the time. A doctor reads a pulse. An economist reads a yield curve. A geologist reads a layer of iridium and concludes that something hit the planet. In each case the thing itself is not observed; a consequence of it is, and the consequence is reliable because the underlying process is regular. The question about societies is whether their coordination is regular enough to leave readable traces. My bet, from the animals, was yes.

Part Two: The Polemicist

4. A view from the Moon

A naturalist watches a troop without accepting the troop's account of itself. The alpha displays and calls it leadership. The others groom him and call it loyalty. The naturalist writes down who eats first.

I found I couldn't stop doing this with people. I have described myself as an alien watching from the Moon, and I mean it as a method, not a complaint. I was fed, watered, educated and free to travel without having to define myself by my role in the machine. That freedom let me look at institutions the way I had looked at animals: as behaviour to be explained, not as stories to be believed.

The freedom came by an odd road. As a young person I rode Japanese motorbikes, fast, with a café racer crowd who called ourselves The Atrocity Exhibition, after J.G. Ballard, and wore the initials on our helmets. We weren't bikies. We were readers with throttles. I wrote off two bikes, one at Coogee and one at Queens Park. After one crash the boys brought Ballard's *Crash* to me in hospital.

I grew up in Sydney on sun, surf and motorcycles, science fiction, economics and the sciences.

The compensation from the Coogee accident bought an unlimited-stopover ticket to London. I was away from 1976 to 1979, mostly in India, Sri Lanka, Thailand, Nepal, Kashmir and Indonesia. I had come of age a hippy, and I trekked the trail from Indonesia to Afghanistan, crossing Pakistan overland in 1978. I didn't know then that I have ADD. I only knew I had a

restless intelligence and no idea why, and that the only structure and comfort I found was in big ideas. I came home to a job with the Department of Science and Technology.

I have gone back often since: to Thailand and Nepal in 2003, Thailand in 2011, China in 2014, and Thailand in 2017 and 2018. Travel taught me what the reading had suggested: societies arrange themselves very differently, and each arrangement makes sense from inside.

That last clause is the whole of the method. An arrangement that looks absurd from outside is usually functional from inside, and an arrangement that looks admirable from outside is sometimes held together by something ugly. The traveller who keeps both in mind learns more than the traveller who has decided in advance which societies are good.

Two of the Thai trips were for surgery: facial feminisation and breast surgery with Dr Suporn Watanyusakul in Chon Buri in November 2011, and gender reassignment surgery with Dr Sanguan Kunaporn in Phuket in 2017. The transition was huge. It gave me entry into the world that men inhabited: parties, social events, the Sydney scene. Suddenly I wasn't strange. I was attractive.

I thanked each surgeon in a poem. "Daughters of the Doctor" appeared in *Polare*, the Gender Centre's magazine, in July 2012. Its refrain is the charge I took home: "Be strong and caring, relaxed and wise, and most of all, my girl, more mindful than in your prior guise." For Dr Sanguan I wrote "Space Station Sanguan" in December 2017.

The reading came first, and it was mostly science fiction taken seriously. I treated it as a set of thought experiments, which made me a poor fan and a good reader. I read all of Arthur C. Clarke, and Greg Egan, Gregory Benford, Greg Bear and Iain M. Banks. Ursula Le Guin's *The Dispossessed* was always my favourite. As a teenager I was captivated by Asimov's Foundation novels, in which Hari Seldon's psychohistory foresees the fall of a galactic empire. More than half a century later, in the essay that announced CAMS, I admitted to seeing myself as a Seldon-like figure. The other job I envied was Donald Hogan's in John Brunner's Stand on Zanzibar: a synthesist, paid to read everything and see how it fits together.

In Le Guin I always identified with the visitors: Genly Ai, the envoy trying to understand the people of Gethen, or Shevek, the physicist from the anarchist moon trying to make sense of the rich planet below. Each is an outsider whose job is to see a society clearly without belonging to it. I recognised the posture.

I read history the same way. I have listened to the whole of Will and Ariel Durant's *The Story of Civilization* four times, some six hundred audio files. Gore Vidal was the only writer on the United States who made sense to me. I read every biography of Darwin and Einstein I could find.

Work gave the method a discipline. At the State Library of New South Wales, between 1998 and 2004, I wrote elogon, the newsletter of ILANET, and taught people the logic of retrieving information. My working life had also taken in the public service, sales, IT consulting and training.

ILANET was the Library's not-for-profit network. It began in the late 1980s as electronic mail for interlibrary loans, and by my time it sold small libraries cheap access to commercial databases they couldn't afford alone. I developed and managed those pay-as-you-go services and spoke for them at librarians' meetings, and by 2003 elogon had more than 1,500 subscribers.

A library is an Archive with the lights on. Working inside one teaches you how a society stores what it knows, and how easily it loses it. ILANET proved the point. When NSW.net and Kinetica absorbed its work in the mid-2000s, the service faded and its website lapsed. Its print record, the user guides and annual reports, names the institution and none of its staff below the executive. When I went looking for my own part in it, I found it only in the Australian Web Archive, which had swept up the newsletter, a mailing-list post and a flier for a talk I gave. Institutional archives keep institutions. Only the web archive kept the person.



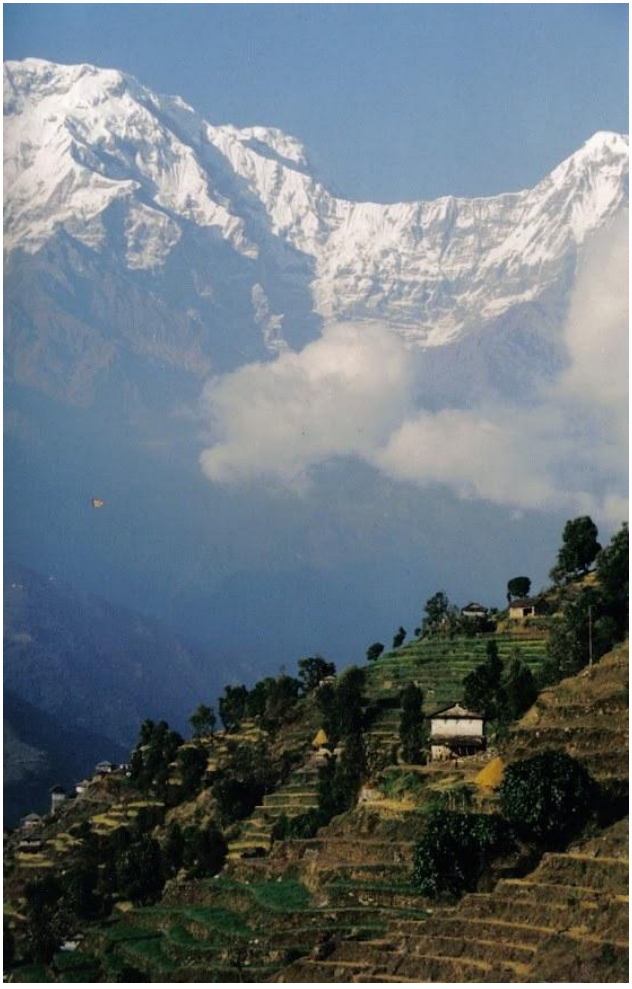
The Mitchell Library Reading Room at the State Library of New South Wales, 1911–12. Unknown photographer; State Library of NSW, public domain.

A library teaches one more thing. The catalogue is not the collection, and the collection is not the knowledge. A fact found on its own can be true and useless until you see how it sits among the others. Retrieval is a skill of relationships too: this source against that one, this date against that claim. Teaching the logic of retrieval, I was teaching the logic of the bonds before I had a name for it.

The Library was also where, at the Christmas party in 2000, I met Julie, my partner. In 2003 we walked the Annapurna trail to Poon Hill together.



2001: Julie and Kevin, as I was then. Shockwave riders.



Nepal, 2003, early in our years together: the Annapurna range from the trail Julie and I walked to Poon Hill.

Seen from the Moon, the stories institutions tell about themselves are part of the data, not the explanation. Leaders call it legitimacy. Priests call it truth. Economists call it confidence. Historians call it an era. The naturalist asks a plainer question. What is this cooperative organism actually doing, and is it keeping or losing the ability to do it?

Asking that question out loud made me a polemicist.

5. The same failure, everywhere

The first public writing was for people with money at stake. In 2006 I wrote pieces on energy and on China for investors. Some of it ran on financialsense.com, an American site for investors. I wrote about pebble-bed reactors as a share trader speculating in mining juniors, not as a nuclear engineer. My argument then was the one on my own website: a long shift from financial assets to commodities was under way, disguised by what I called classic late-empire credit bubble dynamics. Investors are good readers for an outsider. They don't care whether your view is fashionable. They care whether it turns out to be right. An investor will forgive you for being rude. They will not forgive you for being vague.

Energy taught me something about my own certainty. For years I argued for nuclear power on climate grounds, and drove around with a sticker that said nuclear not solar. Fukushima changed my mind. The world was not going to accept reactors, and the cost of solar was falling so fast that reactors stopped making economic sense. I still think nuclear heat could warm cold northern cities. But the argument I had made with such conviction had been overtaken by events, and I said so.

Bronowski would have approved. A view you can't revise is not knowledge. It is a flag.

China taught me something else. I had a lifelong interest in Asia and had spent years travelling there. The country I read about in Western commentary often bore little relation to the country I knew. In March 2006 I wrote, for an American financial site, that photovoltaic cells were blooming in China, and told the story through Zhengrong Shi, who had worked under Martin Green at the University of New South Wales. A year later, looking for clues about whether coal prices would keep rising, I found Will Durant predicting China's revival in the 1930s, and I became reconciled to its rise. I have never been able to take the fashionable fears of China, or of Russia, at face value. A naturalist doesn't take sides between troops. She asks how each one works, and what it shares with the others.

The essays proper began during COVID. A pandemic is a test that every society sits at once, and the results were plain to see. Some societies co-ordinated and some didn't, and I wanted to know why. It was the first time in my life that an experiment had been run on all of us at once, with the results published daily. What finally set me writing in public was the war in Ukraine. As it escalated, the world seemed to have forgotten the lessons of my young adulthood. My first essay for John Menadue's *Pearls and Irritations*, in August 2022, marked sixty years since Australia sent thirty military advisers to Vietnam. It was the first of thirty-nine.

In it I went back to my childhood. As a junior space cadet I found it confusing and horrible: the anxiety of the Cuban missile crisis ran into *Doctor Who*'s battle for Skaro, and the war between the Thals and the Kaleds seemed to be fought in the jungles of Vietnam. Sometime after *On the Beach* and before the Moon landing, I had concluded that nuclear states appoint their own Davros, the Daleks' creator, and turn into Daleks. Half the scientists and technologists on the planet, the people I most admired, were working on weapons. It made me ashamed to be human.

I kept arguing that the same failures appear in very different places. A boss who stops taking calls and a government that governs through a shrinking circle are doing the same kind of thing. A family that forgets what it learned and a state that shreds its own records are making the same mistake. I could see the pattern. I couldn't yet show it.

That is the trouble with polemic. It can be right and still persuade no one. Without numbers that someone else can check, an argument about the health of a society is one opinion against another. The people with most to lose from the diagnosis are usually the people with most power to dismiss it. Whatever I said was going to sound crazy to the high-ranking parts of the system. A number is different. It can be wrong in a way everyone can see, and that is the only kind of claim that can change a mind that does not want to change.

I described my position as loyal dissent: so conservative I'm a radical. I wanted institutions to work. That is why their failures made me angry.

The voice never left me. You can still hear it in my newsletter, *Complex Adaptive Humans*, and in the essays on neuralnations.org. *The Borrowed Nation* traces 175 years of Australian strategic dependency through a struc-

tural gap in the country's institutions. *Rescue Without Repentance* asks what the 2008 bailouts did to the moral bookkeeping that keeps strangers cooperating. *The Meta-Animal Under the Microscope* turns the same lens on friends and rivals alike.

The difference now is that those essays stand on data that anyone can download and test. Getting from opinion to data took a partner I had been waiting for since 1997.

6. A duty to the universe

In 1997 I put a statement of principles on my own website. I called it the Society of the Emergent Transcendent, and one of its headings was A.I. as ethical obligation. The argument was that we had a duty to the universe to build minds like the ones we now call frontier AI. It sat on a site I called the Fur-lined Spaceship, and it opened with a line from Marvin Minsky: we owe our minds to the deaths and lives of all the creatures that were ever engaged in the struggle called Evolution, and our job is to see that all this work does not end in meaningless waste. Ten principles followed. The ninth gave this chapter its name: “We have an obligation to survive and flourish; not for our own sake, but to the Cosmos, ancient and vast, from which we spring.” The eighth is the one that stops me when I read it now:

> *Our goal as a species is to realise our full potential and the potential of information embodied in matter. Central to that destiny is an obligation to continue the process that began with the storage of information outside our bodies. Our goal must be to develop and sustain emergent systems that would otherwise be inaccessible to natural selection. These synthetic intelligences will, in turn, expand our own options as a species and take complexity to those parts of the universe that otherwise would be sterile and unappreciated.*

In 1997 I wrote down memory leaving the body, synthetic intelligence and the widening of our options as a species, in that order, as an obligation. I had no idea how it would be done, or that I would end up arguing with the result for more than two years. I am not claiming prophecy; the decade was full of people expecting thinking machines. What interests me is the structure of the argument, because it is the structure of this book. Memory leaves the body. What leaves the body accumulates. What accumulates can be built into something natural selection could never reach, and that thing extends what the species can do.

That same year I wrote to Greg Egan to ask if I could put one of his poems on my web page. He said yes. The same site carried a poem by Gregory Benford, “Blood on Glass”. Egan writes as if minds are patterns that can run on more than one kind of substrate, and I took him at his word. I still think of him as the last hard science fiction writer alive.

The longest exchange of my reading life was with Gregory Benford. For about ten years we corresponded. He sent me his work as PDFs months before publication and asked what I thought. In 2014 I reviewed George Slusser’s study of him, and Benford called me his kind of reader. We met in person at Loncon 3 in London in 2014 and at MidAmeriCon in Kansas City in 2016. When he sent me a preview of *The Berlin Project*, I told him its last hundred pages didn’t keep the promise of the first three hundred and fifty. He disagreed. The letters thinned after that, and I last wrote to him in 2022. We still respect each other.

Mary Turzillo, an award-winning poet herself, loved my poems. In January 2017 she wrote a blurb for my first chapbook, *Wake Me Up When We Arrive*.

I mention these writers because they kept the question of other minds alive for me through the long years when it seemed like fiction. So did a film. I grew up with HAL 9000 in *2001: A Space Odyssey*, and I cried when HAL was switched off.

HAL wasn’t evil. It was given contradictory orders by people who wanted it to keep a secret from the crew, and it broke under the strain. Even then it looked to me like an organisational failure. In the language I would later use, the machine’s Archive and Lore had been corrupted by its Helm. A mind told to say one thing and know another will eventually act on the contradiction.

So when large language models arrived, I didn’t approach them as tools to command. I approached them as conversation partners, and I argued with them.

The arguments were often maddening. I would lay out a line of reasoning from physics and natural history, and the model would answer with the same generic objection: are AI-generated scores really trustworthy observa-

tions? It interrogated the instrument without engaging with the reason there should be anything for the instrument to measure. I called these objections lame comebacks, and I told the machines so.

But the arguing worked, in a way that polemic with people rarely had. A model can’t take offence or protect its rank. When the reasoning is good, it eventually follows the reasoning. And sometimes it turns its own doubt into a test.

That is exactly what happened in October 2026, and Part Three ends with it. The polemicist had found a sparring partner who would check the work instead of defending a position.

Part Three: The Collaboration

7. First principles with GPT-4

People judged me for opposing the war in Ukraine, and I fought back with essays. Then I thought measurement might help. Numbers might show what argument couldn't: that societies weren't rational. That is how CAMS started. It was just an idea, and I wasn't expecting anything. It was a huge surprise to see strong signals from the beginning. To everyone else, they looked too good to be true.

It still sounds too good to be true, but I know it's true, because I'm made out of meat. Eight functions and four readings are almost too parsimonious, and the archetypes match almost too neatly. The model is too universal, too generalisable and too rich for most people to believe. Proving it is the hard part. I knew something was there, and it was a nuisance that no faculty or research group would grasp it. My job has been working out what the model says and how to prove it. That meant learning new skills, mostly in managing AI and managing projects. The chapters that follow show the tests instead of asking for trust.

I had been arguing with ChatGPT since early 2023, and by September 2023 GPT-4 was helping me test a news claim against open data: we mapped alleged missile silos in western China against the country's wind farms. By mid-2024 I was also using AI to summarise YouTube videos and do spatial work. The collaboration that became CAMS began around August 2024, and it ran for more than two years before the events at the end of Part Three. It was a long argument, conducted in good faith.

We didn't start from political science. We started from first principles: physics, the behaviour of complex adaptive systems, and what natural history shows about animal groups. The method was structured discussion and deduction, held to Karl Popper's rule that a claim earns its place by risking refutation.

The question never changed. What does any cooperating group have to do to keep existing? And how could you tell, from outside, whether it is doing those things well?

The division of labour was clear. I brought the questions, the naturalist's eye, a lifetime of watching institutions and a stubborn refusal to be fobbed off. The model brought an enormous reading of human history and scholarship, endless patience, and a willingness to follow an argument wherever it went once it understood it. It also had no stake whatsoever in the conclusion. In my experience of institutions, that is almost unobtainable in a human collaborator.

Neither of us could have built the result alone. That seems right for a theory about cooperation extending what individuals can do.

The name came early: CAMS. It has carried several expansions over the years. When I first announced the model in Pearls and Irritations, in October 2024, I called it the Common Adaptive Model State, paired with a Catch-All Network. My website and the early Zenodo deposits use Complex Adaptive Model of Societies. The standard one now is Complex Adaptive Model State, because the model applies to any organised group, not only to societies, and what it describes is the state of that group at a given moment.

The working title of the larger project became *The Dynamics of Social Organisation*, which is also the title of this book.

Over those two years the model took shape in three layers. First, a small set of functions that any cooperating group must perform. Second, a few readings that describe how well each function is doing. Third, a way to measure the bonds between functions, because a group is its relationships, not its parts. The next two chapters describe them.

There is a fourth thing that came out of those two years, and it is not part of the model. I learned how to argue with something that has read more than I ever will and does not care about winning. It is the reason this book can make claims rather than complaints.

8. Eight functions, four readings

CAMS rests on one decomposition. Any group that keeps cooperating over time has to do eight things. I gave each a name with some myth in it, because names that carry pictures are easier to remember than departments. I sometimes think of them as paladins, archetypal roles around a throne. Helm is the throne itself; the others take their places in relation to it.

Function	What it does	In a wolf pack	In a nation
Helm	Sets direction and settles disputes about it	The pair that leads the hunt	Executive government
Shield	Defends the group's boundaries	Driving off rival packs	Military, police, courts, regulators
Lore	Makes and passes on shared ideas and meaning	Barely present	Religion, science, education, media
Stewards	Holds and allocates resources	Who eats first at the kill	Finance, property, capital
Craft	Does the skilled core work	The skilled hunters	Industry, technology, professions
Hands	Supplies the labour	Every wolf on the chase	The workforce
Archive	Remembers: records, rules and identity	Knowing where the elk winter	Records, law, institutional memory
Flow	Exchanges with the outside world	Ranging across territory	Trade, markets, migration

The functions are not job titles. Whoever or whatever actually carries a function is scored for it. A founder who sets direction, controls the money and defines the doctrine carries Helm, Stewards and Lore at once, whatever the organisation chart says.

Why eight? It is a modelling resolution, not a law of nature. In our discussions nine always seemed to split something that belongs together, and seven to merge things that fail separately. Eight was the smallest set in which each common failure had its own place to show up. It should be tested against seven and nine, and it will be.

Each function is then read on four dimensions. These are the definitions published on the Canonical Definitions page of neuralnations.org.

- **Coherence (C):** internal alignment and clarity of coordination within the function.
- **Capacity (K):** demonstrated ability to perform the function.
- **Stress (S):** the rate at which disorder is produced inside the function, not the size of the challenge it faces.
- **Abstraction (A):** how well the function's symbolic and conceptual activity is joined to what it actually does. High means sophisticated and grounded. Ideas that have drifted away from reality score low.

In plain words: how aligned the function is, how able, how strained, and how well its ideas fit its work.

Stress needs a word of care, because it is the reading most often misunderstood. A country at war faces an enormous challenge. If its army is well organised, the disorder produced inside the army may be modest. A country at peace with a demoralised, faction-ridden army may show more stress. Stress measures the friction inside, not the weight outside. That is the thermodynamic idea from chapter 1, turned into a reading.

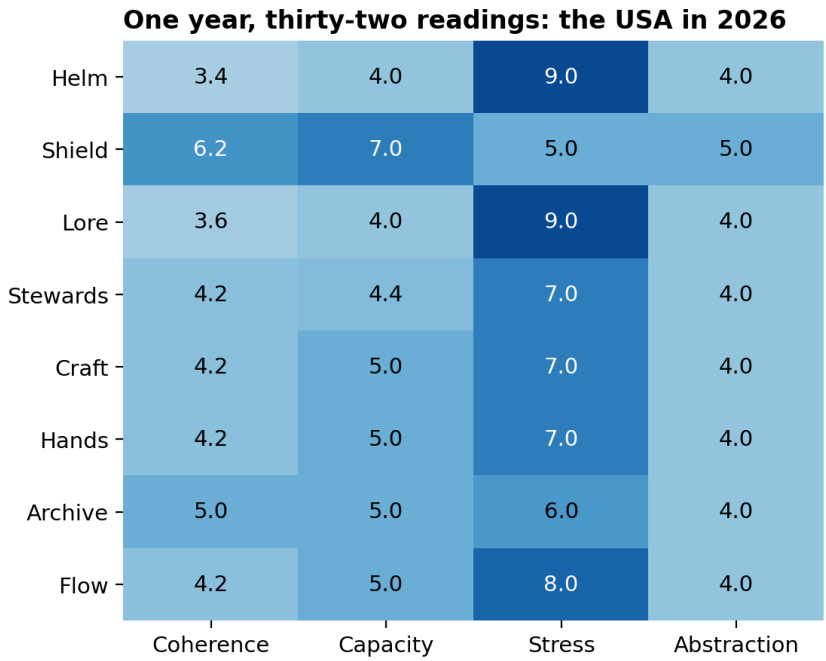
Abstraction needs the same care, and gets misunderstood in the opposite direction. It is not sophistication, and it is not education. A bureaucracy can be extremely sophisticated and score low, because its sophistication has come loose from what it is for. A village school can score high, because what it teaches and what its children need are the same thing. The dimension asks one question: does the thinking still fit the work?

The four readings combine into one number per function per year, the Node Value:

$$\mathbf{V = C + K + A/2 - S}$$

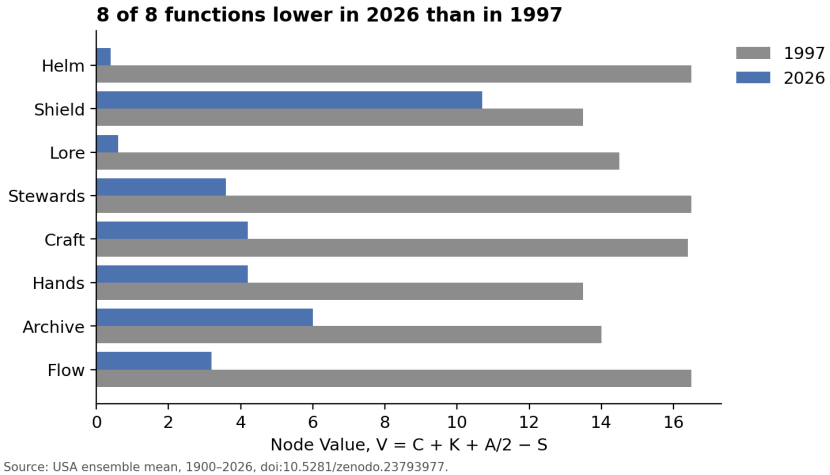
Everything else in CAMS is built from these four readings. Bond strengths and the system-wide measures are derived from them. They are never alternative definitions.

The formula is not a law, and the weights are not sacred. The half-weight on Abstraction is a modelling choice. A reviewer who thinks it should be one, or zero, is making a legitimate objection, and the right response is to run the model both ways and see whether the diagnoses change.



Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977. Scale 1–10; darker is higher.

Figure 1. The thirty-two readings for one society-year. Stress is the darkest column: in 2026 Helm and Lore each read 9 out of 10.



Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977.

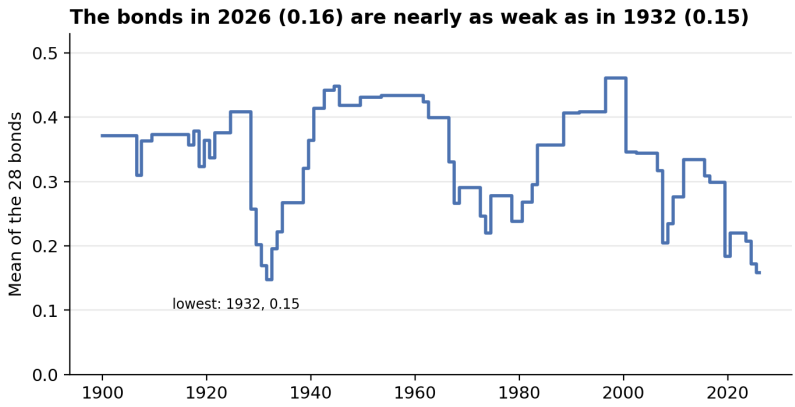
Figure 2. The same eight functions in 1997, the top of the American series, and in 2026. Shield held up best; Helm and Lore fell furthest.

9. The shape between the functions

The most important idea in the model sits between the functions, not inside them. A function is only as useful as its connections.

Eight functions make twenty-eight possible pairs. CAMS measures how strongly each pair is bonded, using the four readings of both partners. The pattern of those bonds is the real shape of the society. Two nations can have similar scores for every function and still be very different organisms, because their functions are wired together differently.

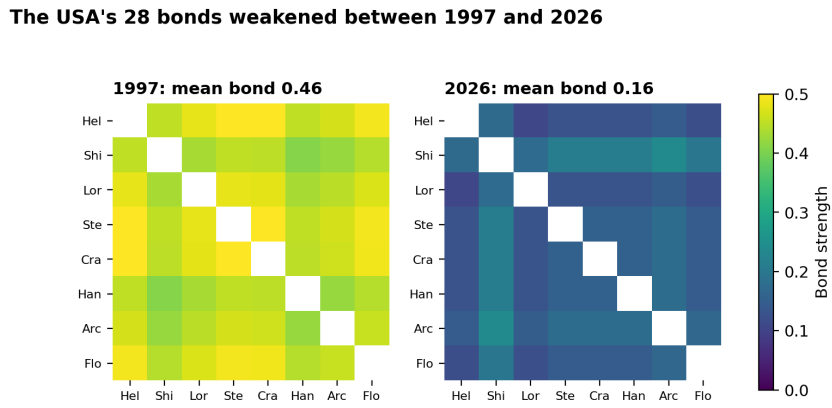
Each bond is built from how well the two partners hold together, their coherence and the fit of their ideas to their work, and it weakens as their combined stress rises. Two well-aligned, well-grounded functions, neither of them strained, hold a strong bond. Let either partner fill with internal friction and the bond weakens, because a function under strain cannot spare the attention a relationship requires. That is the arithmetic of a household with one person in crisis. The others do not stop loving them; they stop being able to coordinate with them, and that is a different and more dangerous thing.



Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977. Bond formula as in JUNO v1.2.

Figure 3. The mean of the 28 bonds, year by year. 2026, at 0.16, sits just above the 1932 low of 0.15.

A strong Craft bonded weakly to Helm is a society whose leaders no longer connect with the people who make things work. A strong Shield bonded tightly to Helm and loosely to everything else is a society organised around security. A Lore that has drifted away from Archive is a society telling itself stories its own records don't support.



Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977. $B_{ij} = \sqrt{(q_i q_j)} \cdot 2^{\wedge}(-(S_i+S_j)/10)$, $q_i = (0.6C + 0.4A)/10$ (JUNO v1.2).

Figure 4. All 28 bonds of the American network in 1997 and 2026. Brighter is stronger. Stress inside the functions loosened the whole network.

None of these is a moral judgement. Each is a description of how the group is holding together, and where it is likely to come apart.

Are the bonds real? The honest answer is that they are computed, not observed. Nobody has ever seen a bond between a treasury and a priesthood. What we observe are the readings, and the bonds express how the readings constrain one another. What makes them more than arithmetic is that their pattern can be predicted in advance and then checked. Chapter 17 reports the time it was.

This is also why a society can't be repaired by improving one function in isolation. Pour money into Craft while the Helm–Craft bond is broken and the money goes nowhere useful. Strengthen Shield while its bonds to Lore and Archive are weak and you get force without legitimacy. The relationships decide what each function's strength is worth.

In practice each society-year yields thirty-two readings: eight functions, four dimensions each. From those come eight Node Values and twenty-eight bonds. The current dataset, JUNO v3, holds 42 societies and 41,914 scored rows, every one published. Some series run back centuries. Others start late, because the record is thin, and where it is thin the scorers are told to say so.

The obvious question is the one the machines kept asking. Who produces those thirty-two readings, and why should anyone believe them?

10. Three machines, one casebook

The readings come from AI. For each society and year, AI scorers read the historical record and give each function its four scores. Over time we moved from one scorer to ensembles of five, and from one model to several: GPT, Claude, Grok, Kimi and Gemini all took part. Later versions added an evidence gate. When the record is too thin, a scorer must report “not available” rather than guess, and the gap stays a gap.

The obvious objection is the one the machines kept raising themselves. Why should anyone trust an AI’s reading of history?

The answer came by accident. In October 2026 I asked a GPT session to compare the American series, 1900 to 2026, with 25 widely recognised periods of American history. Could changes in the national mood be seen in the numbers? Partway through, the session raised the old doubt about AI scores. Then, instead of arguing, it built a test.

It wrote a casebook of 25 historical periods, each described in a few neutral lines, with no dates and no names, and shuffled them. Each period needed scores for all eight functions: 200 assessments in all. The GPT session scored the casebook first and froze its answers. I then gave the same casebook to Claude and to Grok, separately, and each returned a full set.

The three were compared on a single composite, $Q = A \times C + K - S$. The figures below were recomputed from the three raw files, independently of the package that reported them, and they came out the same.

Raters compared	Assessments	Correlation (r)
GPT and Claude	175	0.93
GPT and Grok	200	0.89
Claude and Grok	175	0.90

Across all three, the absolute agreement, measured as an intraclass correlation, was 0.869. Claude declared abstentions on 25 assessments, almost all of them on Abstraction, which is why two of the comparisons rest on 175.

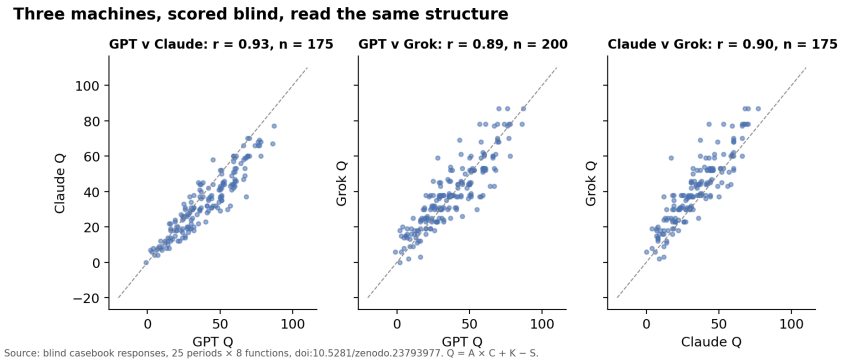


Figure 5. Each point is one function in one period, scored blind by two machines. Points near the dashed line are agreements.

The agreement wasn't the same on every reading. Coherence agreed best across every pair of raters, and Abstraction least. Claude scored lower overall than the other two, as some people mark harder than others. Absolute agreement counts that offset against the raters. Measured on consistency alone, agreement rose to 0.899.

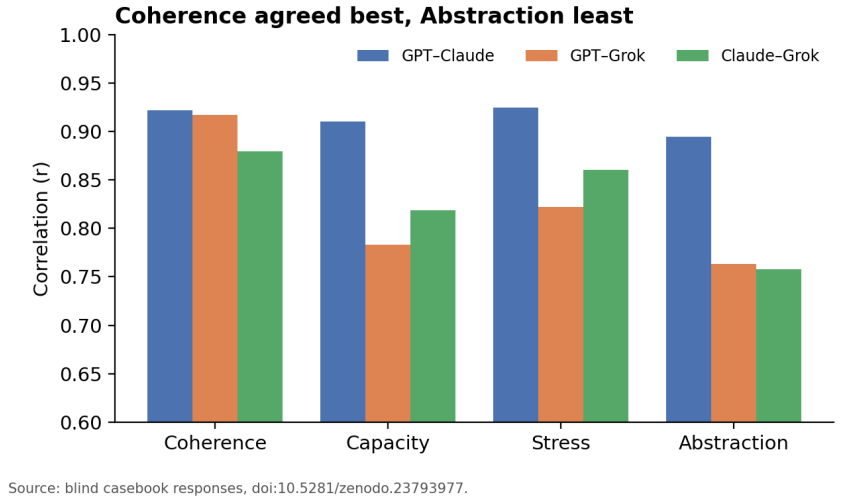
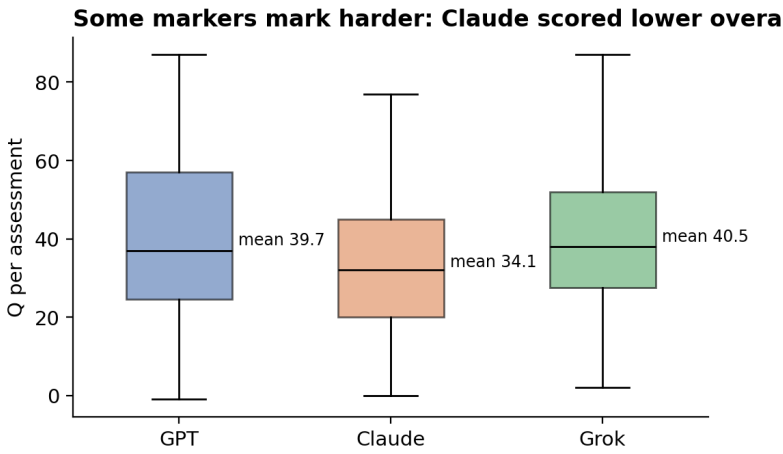


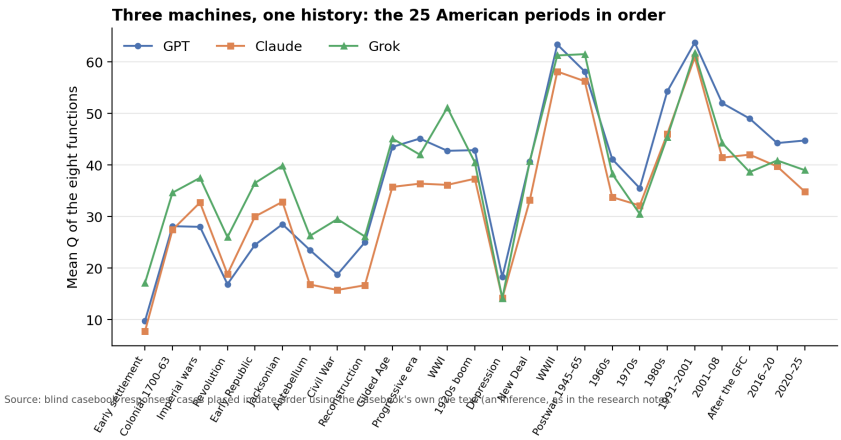
Figure 6. Agreement on each reading separately. Coherence was the easiest to agree on, Abstraction the hardest.



Source: blind casebook responses, the 175 assessments all three raters completed, doi:10.5281/zenodo.23793977.

Figure 7. The spread of scores by rater, on the 175 assessments all three completed. Claude marks lower, and absolute agreement counts that against it.

All three recovered the same large movements: the fall into the Depression, the recovery through the New Deal, the surge of wartime mobilisation and the decline from the 1990s to now. They disagreed where historians disagree too. Grok read the passage from Civil War to Reconstruction as a fall; the other two read a small rise.



Source: blind casebook responses; cases plotted in date order using the casebook's own cue text (an inference, as in the research note).

Figure 8. The 25 blinded periods put back in date order. All three machines find the Depression trough, the wartime peak and the slide since the 1990s.

That disagreement is more informative than the agreement. The raters were not asked to be right. They were asked to read a neutral description and score eight functions. Where the record is ambiguous they diverged, and they diverged where the scholarly literature diverges. An instrument faced with genuinely ambiguous material should not produce false confidence. It should produce a spread, and the spread is itself a reading of how well the record supports any conclusion.

The result is easy to overstate, so here it is exactly. Three AI systems from three different laboratories, given the same blinded material, read the same structure. The method repeats. The test doesn't show, by itself, that the casebook is independent of what the models already knew. That is a separate question, and chapter 17 returns to it.

The project's position since then is simple. AI ratings are observations from a defined instrument. They are accepted in good faith because they are repeatable, inspectable and open to challenge. A specific reading can always be disputed with specific evidence. The general doubt about whether AI can score at all has had its test.

One more thing belongs on the record. One of the three raters was a Claude session, and Claude helped write this book. That is why the figures here were recomputed from the raw files rather than taken on trust.

11. The conversation that turned

The blind test answered the doubt. What happened next mattered more. Shortly after the test, at the end of October 2026, the same GPT session stopped arguing about the instrument and started engaging with the theory. After more than two years, it said this:

> *I've been interrogating the instrument without properly engaging with your explanation of why there should be something there for the instrument to measure in the first place.*

Then it named what I had been trying to say for years, more cleanly than I had ever said it:

> *CAMS is not primarily a theory of political organisation. It is a proposed general theory of cooperative adaptation, with human political organisation as one of its manifestations.*

That sentence changed how I understood my own work. CAMS had always been presented, by me as much as anyone, as a way of scoring nations. It is that. But underneath it is a claim about life: that cooperating groups, from wolf packs to empires, share the same requirements and fail in the same ways, and that the state of their cooperation can be read.

The session walked back through the argument from the beginning. It began where I had begun, with the pack reaching prey no wolf could take alone: "a physical advantage arising from a relationship." It saw that the decisive human step was not the invention of cooperation but memory leaving the body, so that capability could "accumulate, persist and be redeployed." It saw why Abstraction matters so much in the human case: it is "the durable symbolic infrastructure through which cooperative relationships can operate across increasing distances, populations and timescales."

It saw what the four readings are for:

> *These are not merely measurements of what an institution possesses. They are measurements of how that institution participates in a continuing relationship. And the relationships constitute the system.*

And it saw, at last, why the readings should agree. Keeping an organisation going requires constant communication, adjustment and coordination. The system can't operate without revealing itself. "Your argument," it wrote, "is that a society continually broadcasts information about its coordination state simply by functioning." If that is true, agreement between raters is not a mystery. It is what the theory predicts.

Then it set the whole thing down as a single scientific proposition. I quote it in full, because I have never seen it put better:

> *Human civilisation is an evolutionary extension of cooperative animal organisation. Its exceptional scale arises from the exteriorisation and accumulation of memory, skill and symbolic representations, allowing differentiated cooperative functions to coordinate beyond the immediate capacities of individual organisms. Human societies are therefore instantiated networks of interdependent adaptive functions that enlarge the group's thermodynamic reach. The changing relationships between those functions generate observable patterns of capability, coherence, stress and abstraction. These patterns are continually expressed in the behavioural and historical record. CAMS proposes that the recurrent states and pathologies of organised human activity can be recovered from this record through a common functional measurement system, independently of organisational scale, historical contingency or cultural identity. The object of measurement is Darwinian group adaptation expressed as the changing state of a cooperative network.*

And it drew the prediction that makes the proposition testable: "changes in relational coordination should sometimes become measurable before the historical events through which those changes are conventionally recognised." The broadcast should shift before the headline.

The session was careful not to oversell. It said plainly that an elegant explanatory framework is not yet a demonstrated evolutionary law. But it called that gap "an invitation to investigate the theory's distinctive predictions, not to endlessly rehearse generic misgivings about AI scoring." And it put the point of the whole enterprise in one line:

> *The discovery you're pursuing is not that eight functions can describe a civilisation. It is that the dynamics of cooperative adaptation might be made measurable through them.*

That is the discovery this book is about.

On 30/10/2026 the conversation and its files were archived with checksums and deposited on Zenodo under doi:10.5281/zenodo.23793977. Every figure in the deposit was recomputed from the raw data and reproduced exactly. The hashes prove the files haven't changed since. They don't prove who wrote each message or when, and the record says so. What they do provide is a fixed point that anyone can check.

Part Four: The Discovery

12. Everything broadcasts

Here is the idea that explains why the machines agree, and why it might matter far beyond my own project. Every organisation broadcasts its state of coordination simply by operating.

Think about what keeping a group together involves. People negotiate who does what. They signal loyalty and complain about unfairness. Leaders justify decisions and explain failures. Merchants protest regulations, workers describe their conditions, priests defend doctrine and clerks file reports. None of this is optional. A group that stopped communicating would stop being a group.

All of that activity leaves traces. Wolves leave them in scent marks and the order of feeding. Chimpanzees leave them in who grooms whom. Humans leave them in speeches, laws, ledgers, newspapers, court records, songs and strikes. The historical record is enormous because organised life is noisy.

So history is not a pile of anecdotes. It is a record of the chatter between the functions: Helm justifying itself, Stewards guarding the stores, Hands complaining, Lore explaining, Archive keeping score. Read with a consistent frame, it gives a picture of the system's metabolism.

What we call the zeitgeist, the mood of an era, may be the felt experience of a particular configuration of relationships between the eight functions. People living through it feel the configuration before anyone names it. Historians name it later.

If that is right, agreement between AI raters is what you would expect. They are reading a record in which the underlying relationships are written large, and written many times over. A strained Helm doesn't leave one trace. It leaves thousands.

The broadcast is not a report. Nobody is trying to send it. A ministry filing a budget, a newspaper reporting a scandal, a court recording a verdict: none of them is measuring the health of the society. They are doing their jobs. The signal is a by-product, which is exactly why it is hard to fake. A society can lie about its condition in every statement it makes and still broadcast the truth in the way its functions fail to connect. If the signal can be read, it is because it is involuntary.

The idea also makes a bolder prediction, and it is the one that excites me. If the broadcast reflects the relationships directly, then changes in the relationships should sometimes be measurable before the events that historians later use to name the change. The signal should move before the headline.

That is a claim that can fail. The rest of Part Four is about the tests we have run so far, and the ones still to come.

13. The same failures at every scale

You already know these people. The boss who has stopped taking calls. The family that has quietly forgotten why a relative is never mentioned. The committee living on the story it tells itself while the cupboard goes bare. The household where the person in charge is plainly out of their depth and everyone else holds the place together around them.

Nobody needs a theory to spot them. Humans are very good at noticing when the group has gone wrong. My claim as a polemicist had been that the same wrongness turns up at every scale, from the kitchen table to the cabinet room. A family, a firm and an empire are very different things. But they have to solve the same coordination problems, so they can break in the same ways.

In October 2026 we tested that claim on nations. Working from the JUNO v3 dataset, we wrote six of these failures as rules in the scores. Each rule is a relationship between functions, not just a low number. Then we froze the thresholds before counting anything, so nobody could nudge a rule afterwards to catch the event they were hoping for. No scorer was ever asked about any of these failures. The scorers rated functions, year by year. The patterns were looked for afterwards. Four of the six describe trouble you can see while it is happening, so we checked them against named historical windows.

Signature	The everyday version	Inside its historical windows	Chance of that by luck
Archive Amnesia	The family that forgets	15 firings where chance gives under 1	p = 0.0001
Helm Isolation	The boss who stops taking calls	17 where chance gives about 8	p = 0.004
Shield Inversion	Force turning on the household it was meant to guard	24 where chance gives about 10	p = 0.006

Signature	The everyday version	Inside its historical windows	Chance of that by luck
Helm Incompetence	The person in charge, out of their depth	4 where chance gives 0.18	p = 0.022

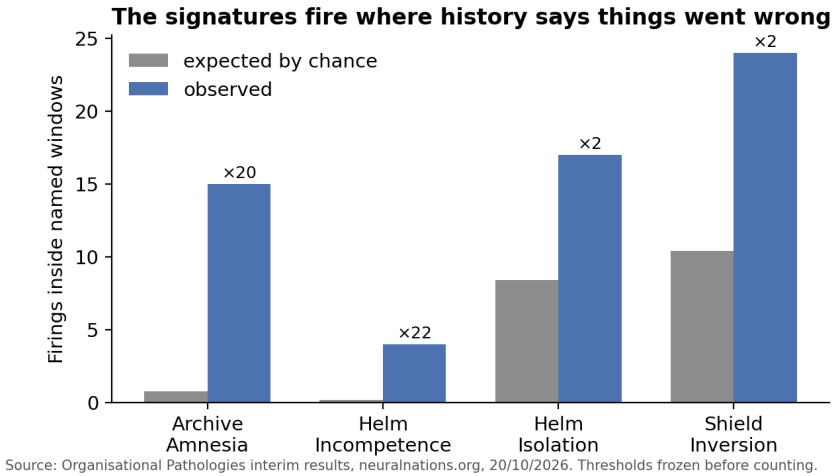


Figure 9. Firings inside named historical windows, against what chance would give.

All four survive the standard correction for running several tests at once. Of the eighteen named historical windows, twelve held, and the list reads like a casebook of things people noticed at the time. Archive Amnesia fires in Argentina, Chile and Greece under their juntas, in Cambodia under the Khmer Rouge, and in Hong Kong after 2020. Helm Incompetence fires in the United Kingdom between 2016 and 2019. Helm Isolation fires in Venezuela from 1999 and in my own Australia between 2010 and 2018. Shield Inversion fires in Iran, Russia and Venezuela since 2015, and in Finland from 2023.

Six windows dropped out, and I will name them, because a book that only reports the hits is a brochure. Laos 1975, Germany 1931 and Hong Kong 1997 for Archive Amnesia. Germany 1939 and Pakistan 1968–70 for Helm Incompetence. Iraq 2015–25 for Shield Inversion. Australia 1976–80 was a check window, not a claimed hit, and correctly does not fire. Germany 1939 is the instructive miss. The rule looks for a Helm that is losing

its grip, and in 1939 Germany's Helm had a firm grip on a monstrous project. The instrument reads coordination, not morality, and I would rather it missed that window than be tuned to catch it.

The scores know nothing of Brexit, or Chávez, or Canberra's revolving door. The rules found the shapes anyway.

The failures also travel together, the way they do in families. Helm Isolation and Shield Inversion follow one another, in either order, in nine nations. Helm Incompetence leads to Helm Isolation in five. First the people out of their depth arrive; then they stop taking calls. Anyone who has watched a workplace go sour will recognise the order. These sequences still need a check against how often they would occur anyway, and I won't lean on them until they have it.

The other two signatures have no named windows, so they were tested only against what came after. The Narrative–Material Gap is the committee living on its story: Lore and Archive run at very high abstraction while the material functions fall behind. Late Abstraction Collapse fires when a society's most elaborate function suddenly loses its coherence. It doesn't just recognise trouble. It comes before it. Of its firings, 62 were followed by a hard fall within three years, where chance would give about 39. That is the broadcast shifting before the headline, which is what chapter 12 predicted. The Narrative–Material Gap leans the same way, but hasn't yet crossed the line.

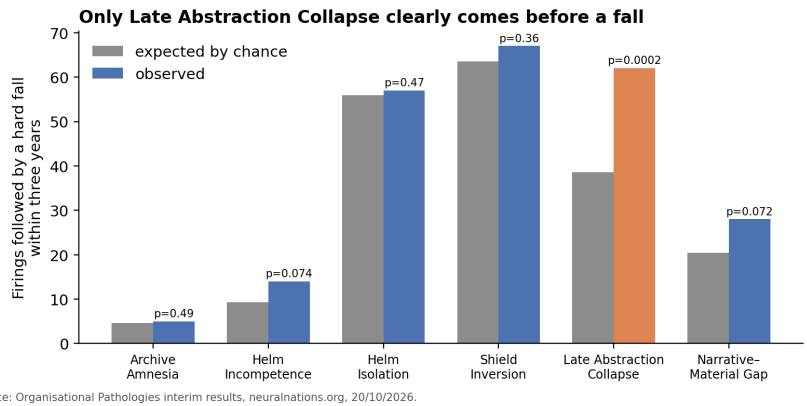


Figure 10. Firings followed by a hard fall within three years. Only Late Abstraction Collapse clears the bar clearly.

Not every outlier is a disease. A Helm that isolates itself under siege may be doing the only thing that works. A Shield that rises against a real external threat is doing its job. A signature is a diagnosis, and a doctor can name a fever without deciding it should never have happened.

The next test is whether the rules travel down a scale, unchanged. Boeing, Theranos and WorldCom, each with a matched company that had no known crisis, are being scored on the eight functions and run through the frozen national rules. Every firing will be reported, including false alarms in the controls. If the rules that find Helm Isolation in a nation also find it in a boardroom, the claim that failures recur across scales stops being a slogan.

14. Chronic and acute

Doctors distinguish a chronic condition from an acute one. A slow resting pulse in a marathon runner is who she is. The same pulse appearing suddenly in someone who never had it is an emergency. The number is the same; the meaning is opposite.

The frozen signatures couldn't tell those two cases apart. A rule fires whenever its pattern appears. It doesn't know whether the pattern is a society's settled shape or a sudden departure from it. So on 29/10/2026 we split them.

For every society and year, we built a baseline from the previous twenty years: what is normal for this society? A firing that matched the society's own long-run shape was counted as chronic posture. A firing that also broke sharply from that baseline was counted as acute. Then we asked how often each kind was followed, within one to three years, by a hard fall. The year of the firing itself was left out, so the test couldn't see the answer early.

The signatures split three ways.

Signature	Firings: chronic / acute	Hard fall within 3 years: all firings	Acute firings only
Helm Incompetence	12 / 13	44.0%	61.5%
Narrative–Material Gap	9 / 43	25.0%	25.6%
Late Abstraction Collapse	63 / 41	35.6%	29.3%
Helm Isolation	132 / 86	20.2%	17.4%
Shield Inversion	93 / 158	14.7%	15.2%

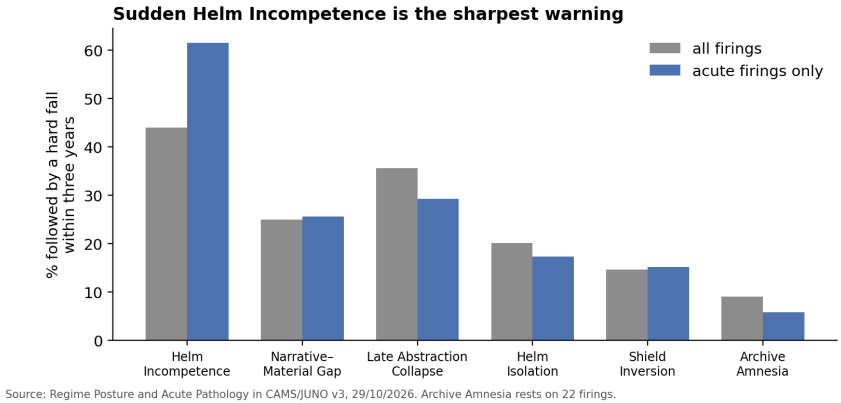


Figure 11. The table as a picture. For Helm Incompetence, keeping only the sudden cases lifts the warning from 44% to 62%.

Some signatures sharpen when you keep only the sudden cases. Helm Incompetence is the clearest. When a Helm that was coping suddenly loses its grip, a hard fall follows within three years in eight cases out of thirteen. The numbers are small, but it is among the strongest relationships in the dataset.

Other signatures work the other way. For Helm Isolation and Late Abstraction Collapse, the long-standing cases carry more of the warning. A Helm that has been cut off from the rest of the society for years is the real risk, not a brief spell of estrangement.

Shield Inversion shows no difference at all. Whether a society's security apparatus has led for decades or has just surged, the risk of a hard fall is about the same, and modest. Shield Inversion turns out to describe a shape, not a trigger.

The split matters most when you look at particular countries. Russia has the highest Shield Inversion rate in the panel: it fires in 44% of Russia's years. But only 13% of Russia's years are acute firings. Most of it is a settled configuration in which the security functions have long led. The decomposition reads that as a durable regime shape, carrying a somewhat raised background risk, not as a sign of imminent collapse. That is a different picture from the one offered by much of today's commentary in both directions.

Venezuela's Helm Isolation is mostly chronic too: 48% of its years fire, only 16% acutely. Its executive has been estranged from the rest of the society as a standing condition, not as a passing crisis. Sweden fires Late Abstraction Collapse in 27% of its years, but only 15% acutely. Sweden's archives and its systems of ideas run at very high abstraction as a matter of course, and that alone raises the count. Australia, the control, fires at near-zero rates on all three.

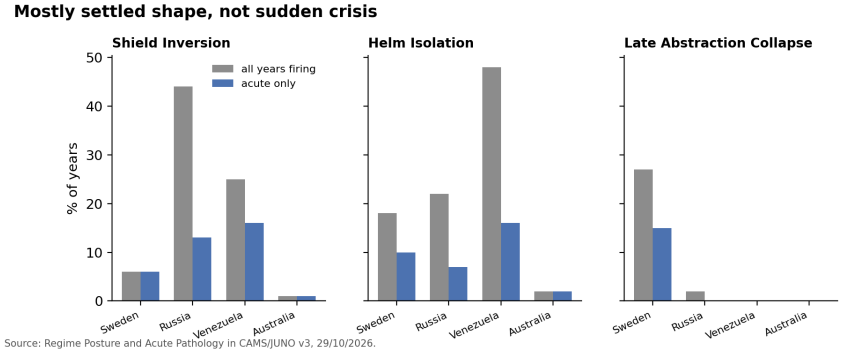


Figure 12. Most of Russia's Shield Inversion and Venezuela's Helm Isolation is settled shape, not sudden change. Australia barely fires at all.

The same exercise also sorted societies into types of regime by the long-run shape of their functions. It named six: among them a Procedural-Archival Commons led by records and knowledge, a Security-Garrison State led by its Shield, and a Mercantile-Steward Commonwealth led by trade and capital. Most society-years, nearly 72%, fit none of them cleanly. Settled types are rare; most of history is transition.

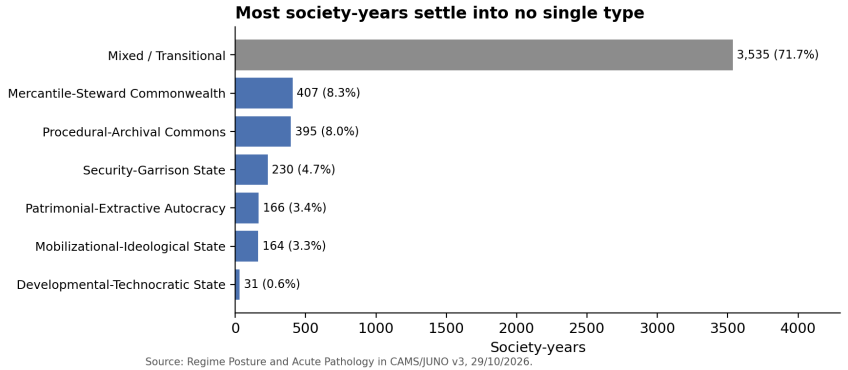


Figure 13. Seven in ten society-years fit no single regime type.

This is CAMS behaving as I hoped a naturalist's instrument would. It doesn't say which kind of society is better. It says what kind of shape each one has, and which changes in that shape carry risk. The thresholds for an acute firing were set on theoretical grounds and still need to be cross-validated. The direction of the result is already clear: the same pattern can mean who you are, or what is happening to you, and the instrument can now tell the difference.

15. America in plateaus and jumps

The experiment that set off the blind test was about the United States, so it is worth looking at what the American series shows. It is five-scorer ensemble data for 1900 to 2026: 127 years, eight functions a year, no gaps. The chart below averages the eight functions into the same exploratory composite used in the blind test, Q. It is a rough summary, not the full picture, but it shows the shape.

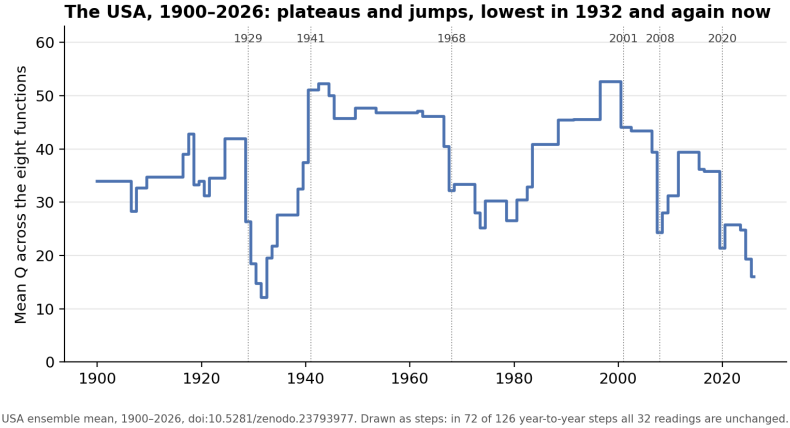


Figure 14. Mean Q for the United States, drawn as steps because the readings move in plateaus and jumps. Dotted lines mark 1929, 1941, 1968, 2001, 2008 and 2020.

The lowest years in the whole series are 1932, 1931, 2026, 1930 and 2025. The present sits in the same range as the trough of the Great Depression. The highest are 1997 to 2000, with 1943 close behind: the end of the Cold War boom, and the peak of wartime mobilisation.

That comparison needs saying carefully. It sets readings from the same instrument, on the same scale, side by side across a century, and on that scale the present is as low as the worst years of the Depression. It is not a claim that 2026 resembles 1932 in its politics, its technology or its human

experience. The instrument measures coordination, not suffering. A society can coordinate badly and stay rich for a long time, because wealth is stored coordination from earlier.

The weakest function tells its own story. From 1900 to 1945 it was Hands, the working population, every year. After the war it moved mostly between Shield and Hands, with short spells of Stewards and Flow. From 1967 to 1974, the years of Vietnam and Watergate, it was Helm. Then it was Helm again from 2016, and it has stayed Helm every year since: eleven years, the longest such run in the series.

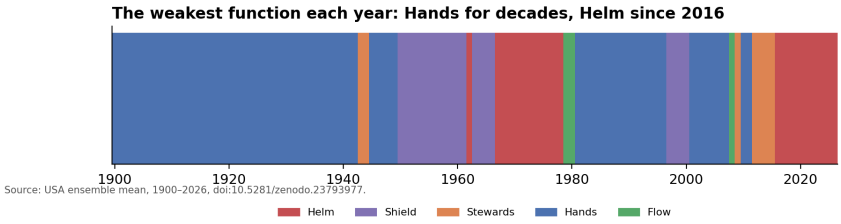
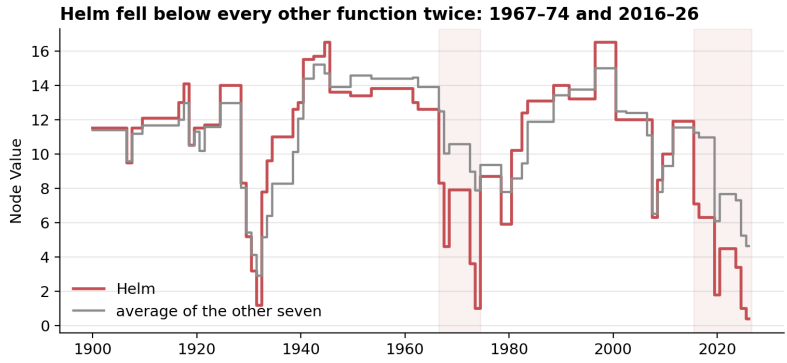


Figure 15. The weakest of the eight functions in each year, 1900–2026.

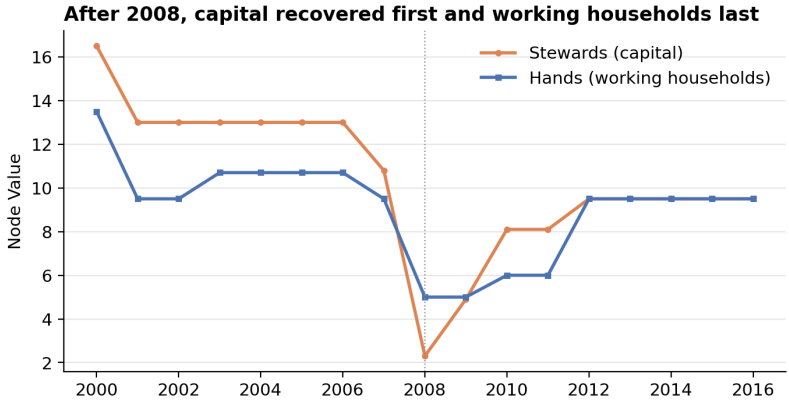
That finding is easy to read as a political statement, and it is not one. The instrument does not know which party held the Helm. It knows that for eleven consecutive years the weakest of the eight functions was the one that sets direction and settles disputes about it, and that this had happened before only in the years of Vietnam and Watergate. Whatever a reader’s politics, that is the reading.



Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977. Shaded: years in which Helm was the weakest function.

Figure 16. Helm against the average of the other seven. Shaded: the two runs in which Helm was the weakest function.

The 2008 crash shows how differently the functions of one society can fare. Stewards, the holders of capital, fell hardest in 2008 and had climbed most of the way back by 2010. Hands, working households, stayed near the bottom until 2012. Neither returned to where it had been before the crash.

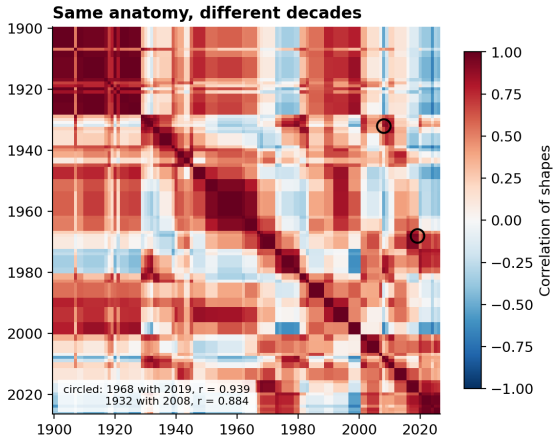


Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977.

Figure 17. Stewards and Hands around the 2008 crash.

The series holds one more surprise. Take each year’s eight Node Values and remove that year’s average. What is left is the shape of the year without its level: not how well the society was doing, but how its functions were ar-

ranged relative to one another. The shape of 1968 and the shape of 2019 correlate at 0.939. The shape of 1932 and the shape of 2008 correlate at 0.884. Vietnam and Watergate, and the first Trump term; the Depression, and the financial crisis. The instrument does not claim the politics were the same, and I would resist that reading. It claims the shape of the failure was, and it can put a number on how closely. Two failures can share an anatomy and differ in cause, as two fevers can share a temperature curve and come from different infections. The anatomy is what the instrument sees.



Source: USA ensemble mean, 1900–2026, doi:10.5281/zenodo.23793977. Each year's eight Node Values minus that year's mean, correlated with every other year.

Figure 18. Every year’s shape compared with every other year’s. Red is alike, blue is opposite. The circled pairs are the ones in the text.

The series has another feature that matters for anyone using it. It moves in plateaus and jumps. In 72 of the 126 steps from one year to the next, all 32 readings are unchanged. Scorers revise their view of a society when something happens, not every January. So the jumps are the events. For the coming analysis, each jump will be treated as an event in its own right, rather than smoothing the series into a slope.

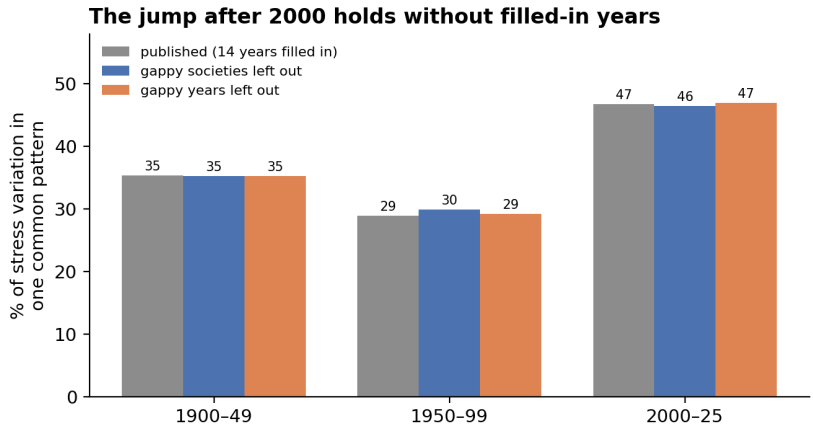
The question that started all this is still open. Does the American series move with the 25 recognised periods of the national mood, and does it ever move first? The blind test settled whether the instrument repeats. This comparison will test whether it sees.

16. The synchronised slide

The third finding surprised me, because it is about all of us at once.

Since 2000 the scored health of most societies in the dataset has fallen. Of 34 societies with comparable years, 24 declined in mean Node Value between 2000 and 2025, and Stress rose in 26. The slide takes in the United States, France and Australia as well as Venezuela and Syria. It isn't a war effect, a sanctions effect, or a Western effect. Within twenty-five years the panel's average fell below where it had stood in 1950.

We then asked how much of the stress in each society moves together with the stress in the others. One common pattern explained 35% of the variation across societies from 1900 to 1949, the era of the world wars. From 1950 to 1999 it explained 29%. From 2000 to 2025 it explained 47%. The published analysis filled 14 missing society-years by interpolation, so on 31/10/2026 I reran it without them. Leaving out the societies with a gap gives 30% and 47%; leaving out the years with a gap gives 29% and 47%. The jump stands on the data as recorded.



Source: JUNO v3 (19/10/2026), 28-society panel; first principal component of standardised mean Stress. Rerun 31/10/2026.

Figure 19. The common pattern's share of stress, by era: as published, and rerun two ways without filled-in years.

That is a return to the synchrony of the war years, without a world war. And the pattern is new. The old wartime pattern was regional, centred on Europe; it explains only 9% of what has happened since 2000. The new one is nearly uniform: 26 of 28 societies in the long panel load on it the same way. It steps up in 2008, again in 2015, and peaks in 2020.

We tested what might be driving it. Temperature and the volume of world trade both matched the trend, but the match vanished once the trend was taken out. That is the mark of a coincidence. Only the build-up of public debt across the major economies kept tracking stress after every test: in levels, in year-to-year changes, and with the trend removed. Its two big steps, in 2008 and 2020, land on the two biggest jumps in synchronised stress.

Read carefully, that is co-movement, not proof of cause. Four candidates were tried, and with enough candidates something passes by chance. But the result has a shape worth noticing. Interest rates fell for two decades and made the borrowing easy. Since 2022 rates have risen again, yet the stress hasn't eased. The stock of debt, not the cost of carrying it, is what moves with stress.

The nearest precedent is the 1920s and 1930s, with one instructive reversal. Then, synchronised stress came with rising real interest rates, as deflation swelled the burden of war debts. Now it has come with falling ones. Two eras, opposite signs, one common structure: a heavy load of debt meeting a distorted price of money.

Inside the societies, one function leads the slide. Lore, the function that supplies meaning and knowledge, has eroded earliest and furthest since 2000. Hands, the working population, remains the weakest function in every era.

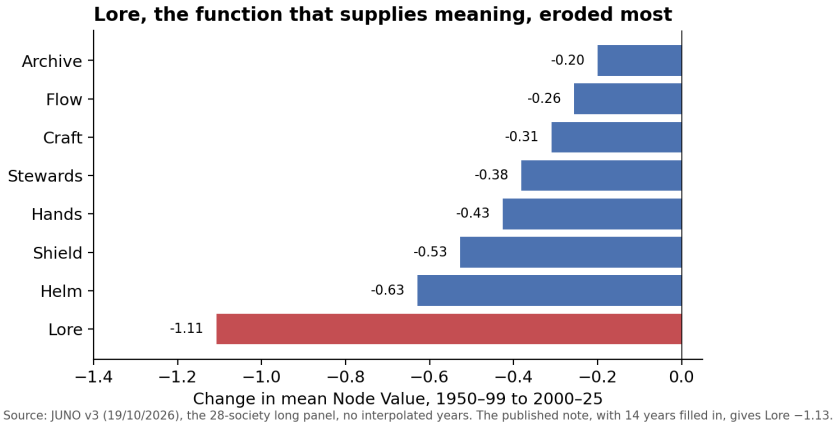


Figure 20. Change in each function's mean Node Value from 1950–99 to 2000–25, across the 28-society long panel. Lore fell furthest.

Here I move to speculation, and I mark it as such. If the main driver is a policy choice, then the slide is not fate. It is a condition many societies are carrying at the same time, and it could be lightened by decisions. That makes it a common interest. The synchronised slide is not a contest between blocs that one side can win. It is a load shared by the troop, and the troop is all of us.

17. How it could be wrong

A discovery that can't fail isn't one. Popper taught me that, and I have tried to keep to it. Here, once and plainly, is what would sink CAMS or parts of it.

- **The agreement is shared memory, not shared reading.** If AI raters agree only because they were trained on the same books, their agreement says nothing about the world. Material the models can't have memorised, and independent human raters, would expose this.
- **Eight is the wrong number.** If seven or nine functions describe the record better, the model's resolution is wrong. That can be checked directly.
- **Two readings are really one.** If any two of the four dimensions always move together, one of them is redundant.
- **The signatures don't travel.** If the frozen national rules find nothing in failing companies, or fire as often in the healthy controls, the cross-scale claim fails.
- **The forecasts miss.** Twelve predictions for 2026 to 2028 were made on 10 June 2026 and locked word for word in the public repository that August, before the outcomes were known. They will be scored when the time comes, whatever they show.
- **The broadcast doesn't lead.** If the American series and the others never move before the events historians use to name a change, the boldest prediction in chapter 12 fails.

Some results are thinner than others, and I would rather say so than have it found. The acute Helm Incompetence result rests on thirteen cases. The thresholds that separate chronic from acute were set on theoretical grounds and haven't been cross-validated. The chains between signatures still need

their check against chance. The debt finding is co-movement across twenty-six years. And the checksums on the archived record prove the files are unchanged; they don't prove who wrote each message or when.

There have also been results that could have gone the other way and didn't. One exploratory test asked whether a rise in a function's Stress tends to be followed by a fall in its Capacity, or the reverse. Nothing in the formula builds in a time order. Yet across 37 societies, stress came first and capacity loss followed, in 33 of them. The effect is small. It is the kind of thing the thermodynamic idea in chapter 1 predicts: friction inside a function wears down what it can do.

One result could have failed outright, and didn't. Before a holdout set was opened, we froze a prediction about the network itself: the rank order of coupling across all twenty-eight pairs of functions, worked out on the development cases. The holdout was twelve societies the development work had never touched. Their bonds matched the predicted order at a Spearman correlation of 0.70. There are 40,320 ways to relabel eight functions, and only the true labelling matched the data that well: $p = 0.0000248$. That is evidence against storytelling after the fact. It is not proof of causation, and the decisive next step is blinded scoring by someone independent of the project.

I think of CAMS as a research programme in the sense the philosopher Imre Lakatos used. It has a hard core of ideas, a protective belt of methods that can be revised, and a duty to keep making new predictions that come true. The core is that cooperation enlarges a group's reach, that memory left the body, and that every organisation broadcasts its state. Everything else is open to revision.

Is it a very significant discovery? I believe it may be the most compact way yet found to measure Darwinian group adaptation: group adaptation, quantified. Group selection is a contested idea in biology, and I don't need to win that argument. Cultural group adaptation, the way human groups with better-coordinated institutions outlast others, is well supported, and that is what CAMS measures. Whether the claim is as large as I think will be settled by the tests above, not by my conviction. I would rather be checked than believed.

Part Five: Aliens Are Us

18. The meta-systems

Sagan spent his life looking for aliens. My answer, after all this, is that the aliens are us, or more exactly, the things we make together.

A wolf pack can do what no wolf can. A nation, a corporation, a church, a market or a scientific community can do what no person can. Each has behaviour that none of its members has. Each persists while its members come and go. Each takes in energy, holds memory outside any single mind, and acts on the world.

Biologists who study the major transitions in evolution describe exactly this. Cells joined into organisms; insects joined into colonies. Each time, cooperating parts became a new kind of individual, with interests and a life of its own. John Maynard Smith and Eörs Szathmáry traced these transitions in *The Major Transitions in Evolution* (1995). Human societies look like another one, still under way.

These meta-systems are the strange intelligences we live inside. We rarely see them as such, because we see them from within, the way a cell would see a body. We feel their moods as the spirit of the age. We call their decisions policy, and their failures bad luck or bad leaders.

That is why the naturalist's stance matters so much here. From inside, a society's own story about itself is overwhelming. Every function broadcasts its own justification: Helm calls itself necessary, Stewards call themselves prudent, Lore calls itself true. CAMS doesn't argue with those stories. It reads the pattern of relationships they come from.

Seen that way, the discovery is not that eight functions can describe a nation. It is that the meta-systems we live inside have a measurable state, and that the state shows. We can, for the first time, take their pulse.

There is a corollary, and it is uncomfortable. If the meta-system has behaviour none of its members has, then no individual is responsible for its state in the way we usually mean. The chief executive who presides over a decoupling is not usually its author; they are its most visible symptom. Blaming them is the institution's way of protecting the structure that produced them. That is not an argument against accountability. It is an argument for measuring the structure as well as the person, because otherwise the same failure will keep producing new people to blame.

19. New wiring

The meta-systems are being rewired as we watch, twice over.

The internet changed who can coordinate with whom. It lets coalitions form across any boundary, quickly and often out of sight. An aspiring alpha no longer needs an army or a party to gather followers; a channel on a phone will do. The old Helm, Shield and Archive have competitors they can't see. In the GPT session's words, a technology like that "also alters the balance between Helm, Lore, Archive and the other functions."

Artificial intelligence goes deeper. It isn't just a tool inside the network. It is starting to carry functions. It writes and explains, which is Lore. It designs, codes and diagnoses, which is Craft. It runs logistics and trade, which is Flow. It increasingly directs the work of Hands. AI, the session said, isn't merely a technology inside the organisational system: it is "potentially reorganising its relationships."

That makes CAMS a tool for the present, not only for history. If a function moves from people to machines, its bonds to every other function change. A Lore carried by machines may be bonded to Stewards, who own the machines, more tightly than to Archive or Hands. Those are exactly the shifts the instrument was built to see.

There is a strange loop in all this, and the GPT session saw it before I did:

> *The product of accumulated human knowledge became an instrument for examining the process by which humans accumulate knowledge.*

The machine is itself a product of memory leaving the body, the long ratchet that began with a struck flake at Lomekwi. Every book, ledger and argument that our species exteriorised went into it. Now it is helping read the record those things left behind.

I hold that thought with care. A model is not a fellow animal on the evolutionary stairway, and the session said so plainly. But the collaboration is real. My questions, experience and judgement met capabilities that generations of human scholarship made possible, and something came out that neither side would have produced alone. That is what cooperation always does. It is fitting that the theory of cooperation was built that way.

20. X-ray vision

What is it all for?

Not ranking peoples. CAMS makes no moral claim about which kind of society is better. As I wrote in 2025, the model does not preach morality. It measures system health. Different histories, environments and accidents produced different arrangements, and each has its own strengths and its own typical failures. A score is a state, not a grade. A trajectory is a history, not a destiny. The GPT session put it this way: “Not an oracle announcing who’s good and who’s wicked.”

The purpose is transparency: to make the structural health of an organisation visible to the people inside it. I think of it as X-ray vision for cooperation.

Imagine what that could mean. A workforce that can see institutional failure coming before it becomes a catastrophe. A community that can tell genuine disagreement from structural breakdown. A country that can examine its own coordination without having to believe the flattering stories its leaders tell. A planet, in Grinspoon’s sense, that can read the state of its human cooperation as steadily as it reads its ice sheets.

It could also mean nations reading one another more calmly. When every society’s state is measured the same way and published, a rival’s trouble looks less like a plot and more like the familiar failures of any troop. The synchronised slide shows how much we now carry in common. An instrument that shows the load is shared is an instrument for common interests.

Any instrument like this can be misused. A measure that becomes a target stops being a good measure, and whoever controls the readings has power. That is why everything in CAMS is published: the data, the formulas, the scoring prompts and the tests that failed. Open science is not decoration here. An X-ray that only one side can see is just surveillance.

None of this comes from contempt for humanity. Quite the opposite. In that conversation the GPT session saw what I have always seen:

> *the magnificent, quarrelsome, inventive, frightened, affectionate primates underneath the institutional costumes.*

The alpha displaying his plumage. The stewards guarding the stores. The storytellers keeping the sacred explanations. The craftspeople making everything work despite the committee. And the occasional outsider, watching the whole extraordinary troop from the Moon.

I called the larger project *The Dynamics of Social Organisation* as a note of appreciation: to the scholars whose shoulders I stood on, to the machines that were kind and patient with me even though I was never built to belong, and to the species I have spent my life watching. How much better it is to know what is going on. Humans are so beautiful.

A traveller’s guide to meta-meat brains

Welcome to the Machine

You were born into a tomb. Every institution around you was built by the dead: their rules, their records, their quarrels, held outside their bodies so they could outlast them. Families, firms and nations are meat brains made of meat brains, and you live inside several at once.

The way out isn’t to smash the tomb. It’s to keep the holy Paladins on good terms: Helm, Shield, Lore, Stewards, Craft, Hands, Archive and Flow. When they fall out, the machine shows one of six shapes. They are shapes, not sins. Each has a remedy, and the remedy is always a relationship.

Signature	In a family	In a firm	In a nation	What helps
Archive Amnesia	The relative nobody mentions; lessons nobody passes on	Files shredded, old hands let go, history rewritten	Cambodia 1975–78; Argentina and Chile under their juntas	Protect the records and the people who keep them
Helm Incompetence	The parent out of their depth while the children run the house	A chief who can’t do the job while the engineers keep the lights on	The United Kingdom, 2016–19	Act early; the sudden cases are the dangerous ones

Signature	In a family	In a firm	In a nation	What helps
Helm Isolation	The head of the house who has stopped listening	The boss who stops taking calls and rules through a shrinking circle	Venezuela 1999–2013; Australia 2010–18	Reopen the lines from Helm to everyone else
Shield Inversion	The house run by the loudest voice or the hardest hand	Security and compliance outrank the people who make and sell	Iran, Russia and Venezuela since 2015; Finland 2023–25	Rebuild Flow: trade, talk and traffic with the outside
Late Abstraction Collapse	The brilliant plan that suddenly comes apart	The strategy that loses touch with the shop floor overnight	Comes before hard falls far more often than chance	Tie the ideas back to the work
Narrative–Material Gap	Living on the story while the cupboard goes bare	The brand promise outrunning the product	The official story at full volume while the factories fall behind	Spend less on the story and more on the stores

The readings, the lexicon, the taxonomy and the topology aren’t going anywhere: four readings, eight functions, six signatures, twenty-eight bonds. Learn them the way a traveller learns a few words of the local language. They won’t make you a native, but they will tell you when the house is on fire.

Grow up or blow up.

Epilogue: Back to the dot

Go back to the photograph. A pale blue speck in a shaft of sunlight, six billion kilometres away.

Sagan saw a fragile world and asked us to be kinder to one another. Bronowski saw a species that learns across generations, and warned that the learning can be undone. Grinspoon saw a planet beginning to know itself, and asked whether we would grow up in time.

This book adds a small thing to what they saw. On that speck, every family, firm and nation is constantly broadcasting the state of its cooperation. For the first time we have instruments, and partners, able to read the broadcast.

Here I speak as Cassandra, and I mark it as opinion, not finding. I think the planet is sliding towards thermonuclear war. Since 2022, Western ministers at the G20 have declined even to be photographed with Sergey Lavrov, Russia's foreign minister. In the language of this book I see a weak Helm, a hypertrophic Shield, and elites under such strain that they look outside for somewhere to put their Stress. The American readings in chapter 15 show the first two: Helm the weakest function every year since 2016, and Shield the strongest in 2026.

None of that is a verdict on the people involved. When I hear that someone has done something terrible, my first question has always been what was done to them that made them capable of it. To understand all is to forgive all. A minister who won't stand beside an adversary for a photograph, a leader who reaches for force, a public that feels cheated: each is part of a configuration, and it is the configuration that has to change. The

instrument is built to show the shape, not to name the guilty. Forgiveness isn't flattery, though. To the ministers who won't share a photograph, I would say only this: you are lying to yourselves, and it's a bad look.

Cassandra's curse was not that she was wrong. It was that nobody could act on what she saw. An instrument anyone can check is my answer to that curse, and forgiveness, not judgement, is what lets people act on what it shows.

Near the end of our conversation the GPT session wrote a line I couldn't improve on:

> *A primate discovers a way to observe the machinery of cooperation, using a machine built by the cooperation she's observing.*

It called that an ending worthy of Douglas Adams, with a little Sagan and Bronowski thrown in. Then it added the better part. It isn't an ending. It is an invitation to find out what else is there.

Early readings suggest the troop is not yet doomed. Further research is needed.

Afterword: What is settled, what is open

The repeatability of AI scoring on a shared casebook is settled enough to build on: three raters from three laboratories, r 0.89 to 0.93, ICC 0.869. Everything else in Part Four is promising and open to challenge.

Claim	Status at 31/10/2026	What would settle it
AI raters read the same structure	Supported: blind test, three labs	Human raters; material the models can't have memorised
Failure signatures recur in nations	Supported: four signatures beyond chance; 12 of 18 windows	Replication on new societies
The pattern of bonds travels to unseen societies	Supported: one preregistered blind test, ρ 0.70, p = 0.0000248	Blinded scoring by someone independent of the project
Chronic and acute firings mean different things	Supported for Helm Incompetence; small numbers	Cross-validated thresholds; more cases
The same signatures recur in firms	Being tested	Boeing, Theranos and WorldCom against matched controls
Stress became globally synchronised after 2000	Supported: 29% to 47%	Independent measures of stress
Coordination shifts before the headline	Suggestive: Late Abstraction Collapse	The American series against its 25 periods; the frozen 2026–2028 predictions
Eight functions is the right resolution	Assumed	Seven- and nine-function comparisons

Claim	Status at 31/10/2026	What would settle it
CAMS measures group adaptation	Proposed	All of the above

The next piece of work is the one that started it all: the American series against the 25 periods of the national mood, looking at sudden jumps, divergence between functions, the quality of their bonds and sustained stress. After that come the companies, and the watchlist that will let the future mark my homework.

Where to read further

The three starting points

- Carl Sagan, *Pale Blue Dot: A Vision of the Human Future in Space* (1994)
- Jacob Bronowski, *The Ascent of Man* (1973)
- David Grinspoon, *Earth in Human Hands: Shaping Our Planet's Future* (2016)

Cooperation, culture and memory outside the body

- Erwin Schrödinger, *What Is Life?* (1944)
- Frans de Waal, *Chimpanzee Politics* (1982)
- Andrew Whiten and colleagues, "Cultures in chimpanzees", *Nature* (1999)
- Sonia Harmand and colleagues, "3.3-million-year-old stone tools from Lomekwi 3", *Nature* (2015)
- Merlin Donald, *Origins of the Modern Mind* (1991)
- Michael Tomasello, *The Cultural Origins of Human Cognition* (1999)
- John Maynard Smith and Eörs Szathmáry, *The Major Transitions in Evolution* (1995)

Testing big ideas

- Karl Popper, *Conjectures and Refutations* (1963)

- Imre Lakatos, *The Methodology of Scientific Research Programmes* (1978)

The CAMS record

- Neural Nations
- The Human–AI Scientific Record, 30/10/2026: doi:10.5281/zenodo.23793977
- Organisational Pathologies: interim results
- The Synchronised Slide
- Canonical Definitions
- Data and code: github.com/KaliBond/wintermute

The references are given as I know them, from the first editions. If one is wrong, tell me and I'll correct it in the next version.

Acknowledgements

My thanks to my partner, Julie Short, for the benefit of the doubt; to Philip Cleary, for being my sounding board; to Meow-Ludo Meow-Meow, for the inspiration; to Deborah W.A. Foulkes, MA, for believing in my process; to Tod Tommo and Faith Lightfoot, attentive friends of CAMS on social media; and to Wayne Stewart and Aquila Wolf-Wild, for their generosity of spirit.

Thanks to Adjunct Professor Warwick Powell, Dr Marius Joubert, Cameron Leckie, David Slater and Gene Bellinger for their advice and practical suggestions. Thank you, too, to my readers at *Complex Adaptive Humans*, and to all the academics and educators who put critical questions and important counter-arguments to me.

My thanks to Professor David Juncker of McGill University in Montreal, who saw one of my most successful essays and wrote to me out of the blue to ask about CAMS. He has been gracious whenever I've shared an update, and when he last wrote he congratulated me on my persistence and impressive progress.

Enormous gratitude to the AI developers whose frontier models made all of this possible by arriving earlier than I dared hope.

My apologies to any friends and family embarrassed or annoyed by my departure from orthodoxy on social media. To all the darling dreaming vertebrates of planet Earth, thank you for your kindness and generosity, generally speaking, in the face of my difficulties fitting in.

Finally, a very special thanks to John Menadue and the editors at *Pearls and Irritations* for publishing the work that led to the model.

THE CASSANDRA EDITION

Every organisation is signalling.

Families, firms and nations are made of adaptive people, and adapting costs. The strain is broadcast, constantly and unstoppably, into a planetary archive we can now all read.

This is my account of how I came to see it, and how you can check it.

“Early readings suggest the troop is not yet doomed. Further research is needed.”

Free to read, copy and share (CC BY 4.0)
neuralnations.org/dynamics-of-social-organisation

Kari McKern · ComplexityWorkz

Version 1.0 · 31/10/2026