

Day 007 — Licensed to Influence

Seven months of Hedegreen Research: ambitious tools, missing conversations, and the question of who gets to choose the mission.

2026.10.20 19:41 DENNIS HEDEGREEN JOURNAL V1.0 [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/DAY-007/](https://hedegreenresearch.com/articles/day-007/)

Seven months of Hedegreen Research, an imaginary intelligence service, and the question of who gets to choose the mission.

Seven months ago, on 20 March, I started a website because I wanted somewhere for my investigations to go. Ideas had a habit of surviving for an evening in a conversation and then disappearing behind the next interesting question. Publishing them meant that somebody else could inspect the work, challenge it, or continue where I stopped. At that point I did not need a research institution. I needed a place to put the research.

In October I have been discussing how Hedegreen Research might reach members of the Danish Parliament, how a free magazine could turn hairdressing salons into places where research finds new readers, and how software might compare market opportunities across the planet. I am also rebuilding the website into something considerably larger than a website, while wondering whether the whole enterprise could one day support films, music and tools as well as papers. Earlier this month I even published an investigation into whether Hedegreen Research needed a popcorn bucket. My son had acquired a Mario one, and somehow that was enough to start another project.

The seventh monthly reflection gets a James Bond title. Bond's licence comes from a state. Mine was issued by me, to me, and authorises nothing beyond what anyone can already do: ask questions, build things and publish the results. The difference is that a fictional secret agent has a mission assigned by somebody else. At Hedegreen Research I choose most of the missions, make much of the equipment and decide what gets sent into the world. This arrangement is good at getting work started. I am beginning to understand why that is not the only ability an institution needs.

The workshop where every idea needs a gadget

Q Branch works beautifully in a Bond film because it has a screenwriter. The watch, pen or suitcase introduced in the first act turns out to be exactly what the agent needs in the third. I do not have a screenwriter arranging for my inventions to become necessary. I could spend weeks improving a device that a spreadsheet would have replaced, or build an elaborate platform before proving that the original problem exists.

A recent question was whether a very inexpensive, high-volume bakery could work in central Copenhagen. It sounds mundane. The price of flour is only one part of the answer; there are wages, rent, transport, ovens, unsold bread and customers who may already be buying something good enough at a supermarket. The immediate research could be a few calculations and a conversation with a baker.

Naturally, my next thought was whether the same question could be asked of every street on Earth.

Imagine opening a map and looking for places where ordinary needs might be supplied differently. Bakeries would be a starting point, not the subject of the whole system. What about repair shops, local production, services that exist on one side of a municipal boundary but not the other, or businesses whose costs might fall if a different distribution model became possible? A small operator who cannot afford a market research department might be able to test an idea before risking years of savings. Sometimes the most useful answer would be that the apparent opportunity is an illusion. A global tool that can say *do not open this shop* for defensible reasons would be valuable too.

I find that possibility genuinely exciting. I can also see the moment at which curiosity tries to promote itself into infrastructure. A gap on a map is not evidence of unmet demand. It might describe zoning, habits, an inaccessible data source, or a mistake in what the software thinks a bakery is. A model that misunderstands one Copenhagen street will not become wiser merely because it applies the same misunderstanding to another million. I need to investigate one street before I have anything worth multiplying across the planet.

This is a recurring pattern in the workshop. Saved Hours starts with a question about human labour and leads to instruments for measuring where time goes. Human Baseline asks what people materially need and suggests comparisons between populations and places. Pizza4People asks whether restaurant ordering can work as open infrastructure. Models of Consciousness turns philosophical disagreement into questions about methods and tests. The projects are different, but each tempts me to build a general tool that will make the next ten questions easier.

Sometimes that is the right decision. Other times the promise of ten future investigations is a way of avoiding the first inconvenient result. The problem I wrote about in Day 0.2

(<https://hedegreenresearch.com/articles/day-0-2/>) has not disappeared: knowing how to begin something is not the same as knowing when to finish it.

And I have already seen how much damage an apparently small tool error can do to the account of what I know. In the August audit behind Day 0.5 — Don't Trust It Yet (<https://hedegreenresearch.com/articles/day-0-5/>), an unavailable NOAA data feed could emerge from the Space Weather tool as the reassuring status “Quiet signal.” A warning could fail to survive its journey into a more official-looking Windowsill report. The audit itself then needed correcting: a suspected missing-year problem in Eurovision Explorer was initially called confirmed when the evidence justified only an open finding. The correction record (<https://hedegreenresearch.com/articles/day-0-5-audit-correction-record/>) changed that classification.

Those are findings about the tools and audit at the time, not claims that the same defects remain today. But they give the workshop a history less glamorous than a row of successful gadgets. Missing data is not quiet weather. A PDF is not safer merely because it looks finished. Even an audit can make a claim it has not earned. If Q is going to invent ambitious instruments, the record must also show when an instrument has misled its inventor.

Rebuilding headquarters while working inside it

Hedegreen Research V3 began as a website update. That has become a difficult description to defend. I started asking what happens to an article after it is published: what happens when a source changes, a number needs its historical context, a later investigation disagrees with an earlier one, or a reader finds an error? A static page can present an argument, but much of the life of that argument remains outside the page.

I want the work to carry more of its history. A reader ought to be able to follow an important claim toward its evidence and forward through its corrections. A historical monetary figure should not lose its original meaning when somebody asks what it represents today; Readable Money is one practical experiment with that problem. The publishing machinery has begun to involve structured articles, editorial assistance and reusable tools for maps, documents and other kinds of output. These are parts of a larger rebuild, not a few cosmetic changes. They are also at different stages of development. V3 is not a completed institutional operating system that I can point to and declare solved.

There is a reason this matters beyond software. The Day articles have returned every month since March, but much of the rest of the research has developed in irregular loops. A question about journalism returns during a software audit. A paper about labour leads to a new question about energy. An argument that looked plausible in May is less convincing after something I investigate in October. A research archive should not pretend that those connections are accidents merely because the articles were not published as a tidy series.

Yet there is a limit to what even the most carefully designed archive can remember. It can show the questions I published and perhaps the ones I explicitly rejected. It cannot reveal every worthwhile question I never noticed. And the time spent engineering a better publishing system is time I did not spend investigating something outside it. I have complained about being stuck in build mode and then responded by inventing more efficient ways to build. That would be a remarkable success if my objective were to remain in build mode forever.

The hope is that V3 eventually makes the work easier to inspect and improve without requiring everybody to trust the person who published it. The risk is that I become so busy constructing a future institution that I neglect the research the institution was meant to support. Both are present in the same rebuild. Calling it a vision does not make the trade-off disappear.

Outside the browser, there is still ordinary life. There are bills to pay, meals to make and a son who acquired a popcorn bucket without intending to provide me with a new research programme. It is often easier for me to explain the architecture of the next system than to explain, in a few ordinary sentences, what Hedegreen Research already does for somebody other than me. That is a communications problem no amount of structured metadata can solve on its own.

From the hairdresser to Parliament

In Day 0.6184 (<https://hedegreenresearch.com/articles/day-0-6184/>). I was concerned that Hedegreen Research had learned how to make work visible without generating enough of the conversations that might make it better. I did not want followers for their own sake. I wanted somebody to use an idea, find a fault, or

show me why the question was badly framed.

This month supplied an awkward example. I tried to obtain press access to the Models of Consciousness 7 conference (<https://amcs-community.org/moc7-schedule-information/>), in Copenhagen and followed up, but I did not get accreditation or the interviews I had hoped for. I could still examine the conference programmes and publish an analysis of the series (<https://hedegreenresearch.com/articles/from-models-to-methods/>). The research could continue without the badge. The conversation I wanted did not happen.

That made me think more deliberately about where future conversations might begin. Folketinget is an obvious possibility because public decisions need evidence, but a research publisher who wants its work in front of a parliamentary committee has already made a political choice. It has decided that this issue deserves attention, selected the moment to raise it, and chosen the people to contact. Making the underlying analysis public does not erase the agenda-setting effect. If I ever send an investigation into that setting, readers should be able to ask who wanted it prepared, who funded it, why those recipients were chosen and which objections travelled with it. I have been discussing that route, not reporting an established channel of influence.

The other end of that imaginary network is a hairdressing salon. Earlier in the year I considered seeking a professional response to an analysis of the hairdressing trade. Later that developed into *THE CHAIR*, a proposal for a free magazine that a salon could print and leave for customers. Someone waiting for a haircut might encounter a question they would never have searched for on a research website. A magazine on a salon table and a briefing sent to an MP can carry the same research into radically different conversations. Neither audience is inherently more important to the purpose of making an idea useful.

And the audience need not be reading research in the usual sense. Signal already has an actual NERO music release, *Inside the Terminal* (<https://hedegreenresearch.bandcamp.com/album/inside-the-terminal>), available through Bandcamp on a name-your-price basis. I have also written stories and film-related pieces that explore questions like energy and the human time returned by machines. Those are not finished feature films, and AI-assisted experiments are not a substitute for the production I might one day want to attempt with other people. But the possibility matters: a question can travel through an argument, a tool, a song or a story. The form changes what reaches people. It does not relieve me of responsibility for the factual claims that accompany it.

I would enjoy seeing Hedegreen Research mentioned in an important room. That is an ordinary, slightly embarrassing motive, and one worth admitting. It would be easy to confuse the name arriving with the research succeeding. But a politician opening a paper, a customer reading a magazine or a listener downloading an album tells me that something travelled. It does not tell me that the journey made the work right.

The honesty business

When someone on Reddit asked whether NGOs had taken an interest in Hedegreen Research, I began drafting an answer about what was actually being built. Tool Builder and V3 remain work in progress; I did not want to turn possible future partnerships into existing ones. While planning this article, I noticed how

tempting it was to turn that ordinary attempt at accuracy into an example of my honesty. The founder who refuses to exaggerate his project would make an excellent character in the story of the project.

That is the trap. I could write an essay about how honesty becomes a brand, and the essay itself could become valuable branding. A reputation for transparency might attract readers, support and invitations. Protecting that reputation might then become more important than allowing information through that damages it. Naming the danger in print is not evidence that I can resist it.

The problem is especially sharp when money appears. The footer of the HR website says, “Open Analysis means you can pay for my time, not my conclusions.” I still consider that a necessary boundary, but the rule leaves something out. **Paying for my time buys an allocation of my time.** A client can influence the research agenda without ever asking for a particular result. It can commission one subject and leave another untouched. A donation can do something similar by making a popular subject easier to continue while an important but unpopular one waits.

There is already a public Transparency page (<https://hedegreenresearch.com/transparency/>), and declaration record. For V3 I want to make income and expenditure more visible, not only as amounts but also as decisions. If money goes toward a journey, for example, the record could carry the reason for choosing that journey and a real alternative that was considered but not funded. That explanation has to distinguish money that was restricted to travel from money that could actually have paid for something else. Otherwise the supposed alternative is a fiction. And the reasons should be recorded when decisions are made, not assembled afterward into a flattering history.

The same question runs through my plans for Coffee and DeepAsk. I would like a reader to be able to support a particular article, TID tool, Camelot report or Signal work, while the money and attributable costs remain reconciled to one HR ledger. That support would offer a form of feedback, not a vote on whether the work is true. DeepAsk could make it easier to ask a difficult question or report a problem, but AI-assisted conversations consume tokens. There is a small absurdity in the model: somebody may buy a coffee to support an article, while the most valuable critic of that article costs the organisation money to hear. If I only welcome cheap criticism, or begin choosing subjects according to how many coffees they attract, the financial design will be editing the research agenda for me.

These are plans, not a running public accounting and feedback system. They also do not produce independence by themselves. A reader can inspect a payment and still not know why an unfunded question was never asked. I have considered making that selection more visible by publishing a short note about a question I chose to investigate and a real alternative I set aside. The bakery would provide another possible test: choose one street, state in advance what finding would count against the idea, and publish the result. Neither idea would count as a control until I actually carried it out.

The wider problem is not solved by confessing that I have preferences. Of course I do. I favour subjects that fascinate me and projects that can become systems. The issue is whether those preferences remain open to challenge when Hedegreen Research becomes interesting enough for other people to spend money or make decisions on its work.

Who is M?

This month I joked that the internet contained courses in everything except starting a think tank. It turns out there is a [guide to building one](https://onthinktanks.org/publication/build-a-think-tank-a-guide-for-policy-entrepreneurs/) (<https://onthinktanks.org/publication/build-a-think-tank-a-guide-for-policy-entrepreneurs/>) and even a [School for Thinktankers](https://onthinktanks.org/initiative/school/school-for-thinktanks/) (<https://onthinktanks.org/initiative/school/school-for-thinktanks/>). I had been discussing global tools and parliamentary distribution before sitting down with the more ordinary questions: why should such an organisation exist, who pays for it, and who governs it? I can almost hear Q offering to build software to answer them.

But I do not want to import every old institutional habit just because I have begun to imagine an institution. At the moment I decide which HR experiments to attempt, and I intend to keep the freedom to do that quickly. A TID prototype should not have to wait for a board meeting simply because I had an idea on a Tuesday afternoon. Slowing every experiment would make the workshop less useful without necessarily making any publication more trustworthy.

What needs to become harder is claiming authority from the result of an experiment. The freedom to test an idea is not a licence to present it as verified. Nor does it give me the right to edit somebody else's criticism into a more convenient version of their objection. Those are different decisions, and they call for different checks. An experiment involving public safety or somebody else's money can also create obligations that an amusing prototype does not.

I have asked another AI model to critique versions of this essay. Some of its objections changed the text. That was useful, but it was not independent governance. I still selected which suggestions to accept and what readers would see. A computer can be an excellent Q, even when it is instructed to argue with Q. It cannot become an independent M just because I give it the part.

One possible beginning is modest: invite a knowledgeable outside reader to challenge a substantial investigation and, with their consent, allow their objection to appear beside the published response without rewriting its meaning. That would not make the reader a board, confer a veto, or mean that every app requires review. I have not established such a process or invited somebody into that role. It is a proposed way to make disagreement harder for me to bury. If HR eventually needs someone with actual authority to stop a risky release or refuse a conflicted commission, that authority will have to be defined rather than implied by a Bond nickname.

Looking back, I can see the sequence. I began by publishing questions. I learned that producing many of them did not mean I had finished them, that a polished interface did not make them correct, and that an archive could remain almost entirely inside my own conversation. Now I am wondering how the work might travel further and what obligations would follow if it did. The seventh month has not settled whether Hedegreen Research will become a think tank, a workshop, a publication or some combination. It has made the question of who can contradict it harder to avoid.

I still want Q to build improbable inventions. I still want to find out whether the planetary map could work and whether the research can reach rooms where it matters. I do not want every new idea to wait for permission. I want the claims that leave the workshop to remain open to people who have no reason to protect my reputation.

My licence is self-issued. It gives me permission to try, not permission to be believed. And before Q designs headquarters with an underground missile silo, I should probably find out how M is meant to get through the front door.

— Dennis Hedegreen

RELATION MEMORY

CONTINUES [Day 0.6184](#) · Continues the monthly checkpoint from missing external conversation to distribution, agenda selection and outside criticism.

RETURNS TO [Day 0.5 – Don't Trust It Yet](#) · Returns to historical tool/audit failure and the difference between a polished surface and earned trust.

SOURCE LAYER [From Models to Methods](#) · Published programme-based consciousness analysis continued without the hoped-for press access or interviews.

RETURNS TO [Does Hedegreen Research Need a Bucket?](#) · The October Mario popcorn bucket callback; no new independent verification of family testimony.

CANONICAL URL: [HTTPS://HEDEGREENRESEARCH.COM/ARTICLES/DAY-007/](https://hedegreenresearch.com/articles/day-007/)