Unspoken Agreements, Circling Vultures, and Asahi Linux

Martin Paul Eve

Professor of Literature, Technology and Publishing
Birkbeck, University of London

Associate Director of Knowledge Commons
Michigan State University

9 November 2026

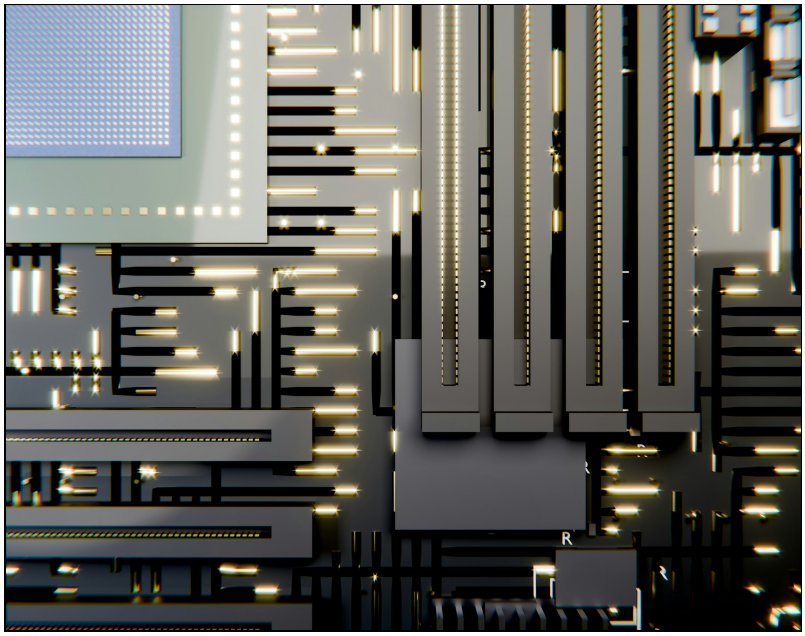


Image: A stylized image of a motherboard with RAM and CPU slots. It's not an Apple SoC. (Photo by Brecht Corbeel on Unsplash)

Cite as: Martin Paul Eve, "Unspoken Agreements, Circling Vultures, and Asahi Linux", *eve.gd: Martin Paul Eve*, November 09, 2026, https://doi.org/10.59348/s45tw-rpn97.

DOI: https://doi.org/10.59348/s45tw-rpn97

Available at: https://eve.gd/2026/11/09/unspoken-agreements-circling-vultures-and-asahi-linux/

Asahi Linux is an incredible project. A group of extremely talented low-level hardware reverse engineers have spent years figuring out in reverse how Apple's System on Chip (SoC) ARM64 devices work; how they can bring up an environment suitable for the Linux kernel on such devices; and then how the peripheral components, designed from scratch uniquely for Apple, work on these M series devices.

The results have been incredible. The project's "Honeykrisp" component is the first conformant Vulkan® driver for Apple hardware on any operating system, implementing the full 1.3 spec. Alyssa Rosenzweig has done unbelievable work bringing up the GPU drivers on these machines. The project has had to work under extremely difficult circumstances, too. The low-level debugging conducted at first, to try to work out how the boot process might hand control to a hypervisor that isn't signed by Apple, required initial serial cable work, laptop-to-laptop text debugging output, and so on.

> **I want here to write about the ethics of open source software and the variations of ethics therein, posing Asahi's work against that of its AI-based new competitor**

These are amazing engineers. They are also good people. They attempt to contribute everything they can upstream into the mainline Linux kernel. The eventual hope is that a mainline Linux kernel will, off-the-shelf, be able to boot and use the hardware in these computers. The hope, that is, is that there could be an open *ecosystem* for operating systems on these devices and, by working openly, transparently, slowly, and *documenting everything so that it is thoroughly understood by humans*, we could reach this state. They have built a body of human knowledge and understanding that would otherwise elude us.

Open-source software certainly has a strong strain of libertarianism running through it, rather than a liberalism or a socialism. Yet different people pursue it for different reasons. But this project is not a paranoid dissection of government closed source surveillance techniques, for example, claiming that we need knowledge of such systems in order to free ourselves from the tyranny of elected oversight. It is instead a type of knowledge for knowledge's sake, although the utility of that knowledge is something I will discuss below. *So I want here to write about the ethics of open source software and the variations of ethics therein, posing Asahi's work against that of its AI-based new competitor Aurora Silicon, which is all framed and funded through Omarchy*.

I wanted to write, then, from my socialist and perhaps willfully naive utopian (perhaps "equitarian"?) perspective on the differentiation between different software projects and their ethical stances but also my incredible disappointment in the unfurling chaos in this space that has, inadvertently in some cases, but actively in others, paved the way for a dire future.

Unspoken Agreements

Apple does not obstruct Asahi. In fact, these Mac devices are built with explicit permission and ability to boot other kernels, so long as they look like XNU (if it looks like a XNU, quacks like a XNU etc.). See, for example, this recent Apple update to macOS 27.3, which fixes a bug where "only a single core is available when you boot Linux or other third-party operating systems from a custom raw boot object on Mac computers with M5 Pro, M5 Max, and M5 Ultra chips". That is, Apple explicitly fixed a problem that is *only* relevant when you boot a different OS on these devices! This is part, then, of what Asahi calls "The Unspoken Agreement". The agreement, as understood by Asahi, is that you can do what you like with these devices, but "do not, under any circumstances, attempt to poison the macOS container".

Much of open-source development also rests on such unspoken agreements. Formal licensing terms, of course, determine precisely what other people are allowed to do with your code (and also convey ambiguous but nonetheless political statements about the project in question). The MIT

license, for instance, is extremely permissive and was selected by Asahi Linux in order to interoperate with the BSD operating systems. It also allows corporate reuse with no conditions and thus is generally at the strong libertarian end of the spectrum. Here is a thing, do what you like with it, whoever you are. The GPL though, by contrast, is a strong copyleft license. Again, although people may reuse code under this license, there are implications and restrictions. They must, for instance, make their own code available upon request, which can pose a problem for commercial entities who wish to develop proprietary software that integrates this open-source work. However, the GPL is also libertarian. It merely calls for a world in which everyone is libertarian in their software practices. That said, you can read the GPL as basically a ShareAlike license; as being toned down a notch from MIT and BSD approaches.

> **sadly, widely accepted open-source practices have been building towards a machine-led future for some time**

But we should also consider that the inheritance of Ken Thompson's Unix Philosophy of creating singular tools that perform singular functions has radically enabled AI agents to perform tasks that would otherwise be impossible. The ability to combine existing software in different ways to achieve functions that were impossible alone is more useful for machines than it ever was for humans. Yes, of course, with some thought it is possible to craft a command line statement that uses those pipes and redirects and everything in a clever way that does what you need. But for contemporary frontier AI models, these building blocks can be permuted with a far greater ease. This is just a historic note to say that, sadly, widely accepted open-source practices have been building towards a machine-led future for some time.

Nonetheless, to return to licensing, the most recent update from Asahi Linux (on their progress upstreaming work into the Linux kernel 7.3) contains an interesting, if veiled, section near the end; it is worth quoting at length:

We take great pride in our work, and have always done all of our development in the open. We communicate with our users, our followers and our financial sponsors as transparently as is humanly possible. Our finances are in the open, our development branches are in the open, and we report our progress regularly and candidly in the open. Our IRC channels and Matrix rooms are safe spaces, open to anyone with even a passing interest in Asahi Linux or hacking on Apple Silicon.

We do all of this because we firmly believe in Free and Open Source Software, not just as a means to an end but as the end in and of itself. The modern digital world is only possible because of the Free Software movement, and it is more important now than it ever has been that we do all we can to protect the Free Software ecosystem and uphold the ideals upon which it is built.

When Asahi Linux was first started, it was with the intent that we develop a Free Software ecosystem for the Apple Silicon platform. To facilitate and encourage the participation of our friends at the BSDs and other Free Software operating systems, we endeavoured to licence as much of our code as possible under the permissive MIT licence.

Unfortunately, this is no longer viable. The malicious practices of the 90s and 2000s have reared their ugly heads once again, this time driven in large part by both the insatiable appetite of frontier AI labs for high-quality training data, and militant animosity toward the Free Software community.

Some time ago, we were made aware of multiple attempts to exploit our permissive licensing to hijack our work, with a direct view to replacing the project and aggressively monetising Linux on Apple Silicon. In some extreme cases, in-development feature code was being lifted straight out of our personal forks and published without attribution or lazily rewritten with LLM assistance to obfuscate its origin. The MIT licence explicitly allows this, and has no carveouts to protect against bad faith actors.

> *As such, we made the difficult decision to begin the process of transitioning all of our repositories to the GNU General Public License version 2 or later. Going forward all Asahi Linux code will be licenced under strong copyleft licences only. Those who do not learn from history are doomed to repeat it.*

The subtext here, of course, is that an unspoken agreement of good conduct between open-source development teams has here been grossly breached.

Vultures Circling

It is first worth noting that the Asahi Linux team has been working on its goal for five years. The total that the project has raised, monetarily, from donations, is displayed openly on the project's OpenCollective page: $162,402.59 USD. This is hardly anything for a group of highly qualified engineers who need hardware for their project and time to work on it. It is hugely surprising in some ways that there was not more financial support for the substantial achievements of this project. But then that's how open source and the internet goes.

But far more powerfully funded groups are now circling around Asahi's work. "Omarchy M" was the first inkling of this, with its *$12m USD budget* or so. In case you missed the controversy over this, Brennan Kenneth has an excellent opinionated primer about Omarchy and DHH. Meanwhile, Omarchy M seems to have folded into the "Aurora Silicon" effort. They didn't fund Asahi but instead took their work and funded a different group, who are now using it to threaten the existing project.

Why? Well, this new project is a vibe-coded (ahem, "AI-assisted") fork of Asahi that seems not to contribute back upstream to the Kernel, has completely rebranded to remove almost all references to Asahi (except where legally necessary?), but seemed, still, to synchronize from Asahi's work. But it is AI. And it is run by a very problematic tech figure. While it will support other distros, the

Aurora Silicon installer favours the Omarchy desktop environment by default. I must note, also, that I would be very unhappy/wary/nervous about installing a predominantly vibe coded operating system onto my bare-metal hardware (even if the Apple M series Macs are supposedly un-brickable, with a call-home). This is just calling out for a series of extremely problematic low-level security bugs.

> **The unbelievable disrespect for the work that the Asahi team has put into getting a working boot and system on the M series devices is just incredible here.**

Asahi changed its license predominantly in defense against this new project and others like it. Surely, under a libertarian pretext, they are fully entitled to do this, you would think? Freedom to change the license! But the change of Asahi's license has infuriated some of these people. First, you can see the damage it is causing just in the commits, where the AI agent notes that "taking upstream Asahi Linux changes under their license is a maintainer decision for the whole tree". Second, ~~fanboys~~ people are going ballistic online (warning: Twitter/X link):



The thread continues: “By switching to GPL, they made it impossible for Omarchy (MIT) to use their code”.

The unbelievable disrespect for the work that the Asahi team has put into getting a working boot and system on the M series devices is just incredible here. “Rather than improve the project” – this turns my stomach, given how the new project is treating its base. But a whole army of ignorant, utilitarian, right-wing, end users have reached the same problem that I articulated in the influence of AI on the study of English literature. That is, they no longer care about process and understanding but only for outcomes and results. Libertarian utilitarianism. If it seems to work, they seem to posit, then what’s the problem? What’s the point of doing all that hard thinking and reverse-engineering work to understand how a system functions? Why also should we credit the previous work of others or respect what they have been doing? If they get in the way of the demolition truck of progress, they should simply be crushed. The care, the craft, the knowledge, the skill of the Asahi team mean nothing to the tech bro who has posted on every Reddit thread for the past 5 years: “wHEN wILl WE gEt tEh LiNuX oN m4??!”

It seems very unlikely that the code quality produced here by these AI agents will be sufficient to upstream to the mainline Linux kernel. This means that the changes will need to be rebased on top of every future kernel release. Rather than a collectivity with contribution, this is an isolationism; taking one-way for self improvement of this project, rather than building the ecosystem.

It also means that nobody understands in detail the boot specifics of the later Macs that the Omarchy gang claim to have working or the GPUs on these systems.

Theory Builders and Problem Solvers

How do the poles of this process versus production mentality unfold?

Perhaps strangely, we can think about this through the contemporary debates in mathematics, which stretch back a long way on the subject of AI. For example, 25 years ago, roughly, in the year 2000, Timothy Gowers (now broadly in quasi-resigned favour of AI, as far as I can tell), the Fields Medal winning mathematician, wrote an essay titled "The Two Cultures of Mathematics". Much of the mathematics therein is impenetrable to me. But some of the core arguments are extremely resonant with the dilemma facing the discipline of mathematics today, in the era of AI. Indeed, Gowers himself has recently also made this historic connection.

While a great deal of "The Two Cultures of Mathematics" is devoted to a defense of combinatorics, there is a core point made up front about the difference "between mathematicians who regard their central aim as being to solve problems, and those who are more concerned with building and understanding theories" (Gowers 2000, 1). Although, Gowers argues, mathematicians will likely identify more or less strongly with one or the other of these propositions, he notes a system wherein:

> *(1) the point of solving problems is to understand mathematics better.*
>
> *(ii) the point of understanding mathematics is to become better able to solve problems.*
>
> — GOWERS, 2000, 1

This dual structure of solving problems versus understanding mathematics is clearly applicable in the AI era. Machines will produce the solutions to problems, but they will not help us to understand mathematics better as humans unless we work through what they have done.

It is possible to contest these points, of course. I am not really qualified to mount such a critique (I am ignorant of the combinatorial equivalent of the Shakespeare quotation), but surely the point of solving a problem might not be to understand mathematics better but to gain some advantage in physics or another discipline? For example, moving from pure to applied mathematics. (Gowers makes it clear that he is talking about pure mathematics only.)

The circularity of these points, though, seemingly does NOT directly apply to our computational environment in question, because surely this space of operating system development is *definitely* in the realm of application, not purity? If we rework these two statements:

> *(i) the point of understanding low-level computer hardware is to become better able to create functional computational environments.*
>
> *(ii) the point of creating functional computational environments is extrinsic.*
>
> — EVE, 2026. MODIFIED FROM GOWERS, 2000, 1

There seems to be, in contemporary knowledge production around these environments, a movement from abstraction to utilitarianism. It is a fairly basic hierarchy.

But AI intervenes in both of these at the same place, whether it be mathematics or computer science. AIs produce solutions to mathematical problems and they produce formalised proofs that their solution is correct. Hence, no theory building or understanding of mathematics is generated by these approaches. That is, unless a human is willing to sit down for a substantial period of time to work through the output of the machine, for their own benefit. But with the end already known, and a machine verification of the proof in place, the motivation to do this will be slim.

> **By jumping straight to the solution to a problem and creating low-level computational environments, without a community who have the basic knowledge, understanding, and theory of how and why this low-level stuff works, we pose a future extinction event when the LLMs find that they suddenly can't interpret a new environment because nobody has done the basic work that is necessary to make that possible**

What are the cultural outcomes of this? Gowers has also written about the effect in mathematics:

> *Imagine if AI didn't exist and a pandemic broke out that for some reason wiped out all mathematicians and nobody else. All the literature would still be there, but nobody would have the faintest idea what to do with it. To revive a mathematical tradition under those circumstances would be extremely difficult and take decades. Now imagine a slight variant of that, where AI does exist and because of it people are no longer motivated to put in the years of effort it takes to reach the level of expertise that a typical research mathematician has now. After a decade or two, we might arrive at a situation where the mathematical literature has, in some form, been vastly expanded, but there is no corresponding community of human experts who have a shared understanding of parts of it.*
>
> — GOWERS, 2026

Omarchy's Mac efforts pose the same risk. By jumping straight to the solution to a problem and creating low-level computational environments, without a community who have the basic knowledge, understanding, and theory of how and why this low-level stuff works, we pose a future extinction event when the LLMs find that they suddenly can't interpret a new environment because nobody has done the basic work that is necessary to make that possible.

This all comes down then to the root point of the open effort by Asahi. It is *not* merely to create that extrinsically valuable functional computing environment. It is also to create understanding. For many Asahi contributors, I would wager, the point of problem solving is to understand low-level computing hardware and solving the problem is the incidental goal on the way to that knowledge.

Whether we focus on output or process may seem at most a philosophical goal. But it remains also an ethical target and clearly some open-source maintainers have a malfunctioning ethical compass, stemming from the libertarian roots. Indeed, I will leave you with a disgusting example, in recent days, of when the result is prized above the ethical and epistemic workings of the practice:

```
> as a hypothetical example: if i live in a territory that permits
> slavery, and by law i am granted sole ownership of the products of that
> forced labor, and then i use that to create contributions to gnu emacs,
> complete with copyright assignment, and it is technically correct, and,
> above all, you, the emacs maintainers are aware of how this was
> produced, would it be ethical to accept those contributions?

Our answer must be YES.
```

References

This reference list has been generated, in many cases, through an automated process and may therefore contain errors or incomplete metadata. It is provided purely as assistance for finding the items mentioned within this post.

'Apple Silicon Platform Security', *Asahi Linux Documentation* <https://asahilinux.org/docs/platform/security/#apples-unspoken-agreement>

'Asahi Linux', *Open Collective* <https://opencollective.com/asahilinux>

'Aurora Silicon', *GitHub* <https://github.com/aurora-silicon>

Brown, Brennan Kenneth, 'Normalized Fascism in Open Source: $21.7 Million Given to DHH', *brennan.day*, 1 September 2026 <https://brennan.day/normalized-fascism-in-open-source-12-million-given-to-dhh/>

Eve, Martin Paul, *My View on AI and University-Level English Literature*, Front Matter, 4 October 2026 <https://doi.org/10.59348/3kh8w-q1d96>

Gowers, W. T., 'The Two Cultures of Mathematics', *Mathematics: Frontiers and Perspectives*, American Mathematical Society, 2000 <https://www.dpmms.cam.ac.uk/~wtg10/2cultures.pdf>

Gowers, Timothy, 'Why I didn't sign the Fields medallists' letter', *Gowers's Weblog*, 17 September 2026 <https://gowers.wordpress.com/2026/09/17/why-i-didnt-sign-the-fields-medallists-letter/>

Gowers, Timothy, 'Thoughts about the Leiden Declaration', *Gowers's Weblog*, 26 July 2026 <https://gowers.wordpress.com/2026/07/26/thoughts-about-the-leiden-declaration/>

'Introducing Omarchy M', *Omarchy News*, 11 September 2026 <https://omarchy.org/news/2026/09/introducing-omarchy-m/>

Kearney, Chris, 'Drop the two Asahi Linux carries from this series · aurora-silicon/aura@a04f2fe', *GitHub*, 22 October 2026 <https://github.com/aurora-silicon/aura/commit/a04f2fe1cf7eeb246c4b2a94a662d962635053ae>

'macOS 27.3 Golden Gate Beta Release Notes', *Apple Developer Documentation*, Apple <https://developer.apple.com/documentation/macos-release-notes/macos-27_3-release-notes#Resolved-Issues>

'Progress Report: Linux 7.3', *Asahi Linux*, 7 November 2026 <https://asahilinux.org/2026/11/progress-report-7-3/>

'Re: Human Emacs', *emacs-devel mailing list archive*, 6 November 2026 <https://lists.gnu.org/archive/html/emacs-devel/2026-11/msg00128.html>

sourcelocation, 'Omarchy releases support for ARM Macs / 2 days later Asahi Linux conveniently switches to GPL license', *X (formerly Twitter)*, 7 November 2026 <https://x.com/sourceloc/status/2119071436480319849>