

AN ALTERNATIVE TO BIG TECH AI · ESSAY 2 OF 7

# AI That Stays in Its Lane

The premium is a machine that cannot cross the line into the judgements that are yours to make.

# The unease isn't only about privacy

What people fear is a machine deciding things it has no business deciding.

- An off-switch answers surveillance — and nothing of overreach.
- On a Big Tech stack there is no lane, only a model that reaches as far as it's allowed.
- The organising question: which decisions must never leave human hands?

# The line that isn't a matter of taste

Some judgements can be systematised; values, meaning and agency cannot.

- Automating a value smuggles in someone else's value as the default.
- The boundaries, plainly: values verified not automated; agency respected not simulated.
- The decisions a community most needs help with sit right on this line.

# A boundary you only hope for is not a boundary

A guardrail written as a wish doesn't constrain what the system actually does.

- Most 'responsible AI' is a sign asking you not to cross, not a fence.
- We built the line as architecture: reserved decisions are gated to humans by the code path.
- A check runs before an action takes effect — and can't be talked out of its job by a prompt.

# What the machine does instead

It proposes; a person disposes — and the seam between them is visible on purpose.

- Surfaces history, summarises threads, drafts routine text — real work taken off the plate.
- Stops at the judgement itself; never crosses into what's reserved.
- Every assist is logged: what was used, which model, where it ran.

# Why bounded is the more useful machine

A machine bounded by construction can be trusted close, not kept at arm's length.

- The bound is what lets you relax and hand it the tedious work unsupervised.
- We don't claim it out-thinks the biggest models — we claim a better arrangement.
- A community needs a capable assistant it can trust not to govern it. That's shipped.

# Who needs the boundary most

The people who carry responsibility they cannot delegate.

- Boards, where a decision blamed on 'the system' is a decision no one is accountable for.
- Clinicians and caseworkers, where overreach is harm, not inconvenience.
- The principled — who wanted not the machine's absence but the certainty it stays in place.

# The principle the rest of the series stands on

Knowing its place is the design discipline that makes everything else safe to own.

- The unrewritable record works because the AI can't author a director's reasons.
- The in-house model is governed by the same line about what it may decide.
- A system worth trusting knows which of the two it is: instrument, not authority.

## See it for yourself

- The demo Villages — [mysovereignty.digital/demos.html](https://mysovereignty.digital/demos.html)
- The free Governance course — [mysovereignty.digital/course.html](https://mysovereignty.digital/course.html)
- Read the essay — [agenticgovernance.digital/papers/ai-that-stays-in-its-lane.html](https://agenticgovernance.digital/papers/ai-that-stays-in-its-lane.html)

You don't have to choose between Big Tech AI and no AI. You can hold your own.