
AI & Liberty — Supercharging our movement

For a generation, the liberty movement has won remarkable victories on limited resources — consistently outspent and out-staffed by those who would grow the state, and making up the difference with conviction and hard work. AI now offers a real chance to transform that equation in our favor, if we are willing to adapt. Put simply, it lets us do far more with what we already have — not only a sharp rise in productivity, but the chance to rethink how we organize altogether. The aim is not to be clever with AI but to do more for freedom: to give donors maximum bang for their buck, and to achieve results once beyond our reach. Here are some of the ways we might better put it to work.

Not just another tool

We have all used AI as a kind of super search engine, and as an awesome writing aid — though lean on it too heavily and what you get is slop. But the real significance of AI is not as a generative tool at all. It is as a **semi-autonomous agent**.

Agentic means AI can be commissioned to carry out tasks — often ones of considerable complexity. It becomes something you delegate to, as if it were a colleague. And once you begin to see it that way, a new question matters: how do you manage it, and delegate to it, without simply deferring to it?

For a think tank, this is not a productivity tweak. It is the difference between a small shop and a big one. We advance liberty on limited resources against opponents far better funded — and anything that lets us do the work of an organization twice our size is a mission multiplier.

It works in two directions. **Inwardly**, it lets us automate the routine work that ties up our people — freeing them for higher things. **Outwardly**, it lets us turn data and research into tools that carry the argument straight to the public.

This paper sets out a number of ways AI is already being used within the movement, and shares the steps some colleagues have taken. They are first, tentative steps — and certainly not the last word.

PART I

Getting going

a) Setting up

A range of AI models now offer this agentic, works-with-your-files mode; the paid tiers are where it lives. Two things are worth getting right from the outset. **Mind your data:** use an enterprise plan, or at the very least switch on the privacy setting, so confidential or donor information is never used to train the model. **Cap the spend:** usage-based pricing can creep up, so set a monthly limit.

b) Prompting

Prompting matters less than people imagine; context matters more. The trick is not a clever form of words but telling the AI what you are trying to achieve, who the work is for, and what a good result looks like — much as you would brief a capable colleague who has just joined.

One habit is worth forming above all others: **start with an exploratory question.** Before you ask the AI to do anything, explain what you are trying to achieve and ask it for suggestions. More often than you would expect, it will offer an approach you had not thought of. Then, with a second prompt, ask it to carry out the option you have chosen.

Even when you think you already know exactly what you want done, this two-step habit pays off: it gives the AI a far better understanding of what you actually intend — and you a better result.

c) AI policies

Any organization using AI would be wise to have a policy — a clear statement of how the tool may and may not be used. It might cover what data may be shared with it, who signs off on what goes out, how facts are to be checked, and where a human must stay in the loop.

But there is not much point having an AI policy if you never tell your AI about it. Plenty of organizations have one — sitting in a folder in the CEO's office, changing nothing. To have any effect, a policy has to be built into the instructions the AI actually reads before it works. In other words, your AI policy should be part of your Super Prompt.

d) Super Prompts — what they are and why they matter

Everyone knows AI can hallucinate. Just as infuriating is its inability to stick to what you thought you had agreed: each time you give it a new task, it tends to start from scratch, forgetting what you settled five minutes ago. That is why you need a **Super Prompt** — a standing order that transcends any single prompt, which the AI reads before it begins.

Distill your policy and your preferences into a set of standing instructions, install them as a file the AI reads before every task, and your principles travel with every job. MCPP runs one, refined over many months, and shares it with anyone who asks.

Some of what a Super Prompt ought to include:

- a. Check prior work first — reuse what's already been done and agreed.

- b. Use your own words when reworking your writing.
- c. Always verify facts, and flag anything you can't confirm.
- d. Say "I don't know" rather than guess.
- e. Show a draft before treating anything as final.
- f. Never send, post, or delete without your say-so.
- g. Record any new standing instruction in the file.
- h. Protect the file — back it up before any edit.

PART II · INWARD

Automating the routine

e) Automated back-office workflows

Some jobs come around again and again. These are the ones to hand to an AI on a schedule.

Examples of automated workflows:

- a. Personalizing donor-cultivation letters and emails, and building call lists, to a set routine.
- b. Creating next-day follow-up reminders from each day's meetings.

In practice, the automations that earn their keep cluster in two areas: communications and back-office administration. The pattern is always the same — the AI does the legwork, reshaping and formatting what you have already written, while a person keeps charge of the words and approves what goes out. A small team ends up with the output of a much larger one.

f) How to start

Do not try to automate everything at once — that is how people stall. Pick a single workflow that is repetitive and valuable; build it end to end with a human approval gate; run it for real and fix what is clumsy; and hand it to one curious staffer so it does not live only in your head. Start simple and build in complexity over time: only once that first workflow is working should you add a second.

Pick the painful one. The instinct is to start with something easy and safe. Resist it: the drudgery that costs you the most senior time, or the most money left on the table, is exactly where the first win should be — because that is the result that proves the case to your board, your donors, and yourself.

g) Guardrails: keeping it honest and safe

The failure mode is not the robot taking over. It is a confident, wrong answer slipping through because nobody checked. Four rules, built into your Super Prompt, hold the line automatically:

Verify and flag. It checks factual claims and tells you what it could not confirm, rather than asserting it.

Ask before acting. It never sends, posts, publishes, or deletes on its own. A person approves anything that leaves the building or moves a dollar.

Admit uncertainty. It says “I don’t know” instead of inventing something plausible. This one rule prevents most of the trouble.

Keep a person at the gate. Automate the legwork; keep the words, and the judgment, human.

Used this way, the risk is low and the upside large.

Avoid slop

Mention AI to a lot of people and they think of it as a writing tool. Plug in a few commands, and hey presto, the AI generates acres of perfect prose. Or does it? Actually, a great deal of what AI writes reads like slop. No matter how good AI becomes at writing, it is not going to have your ideas, your thoughts, your insights. You cannot magic that into existence — and trying to use AI to do so creates a credibility problem.

A good rule of thumb: using AI to edit and improve your writing, your thoughts and your words is smart. Using AI to try to produce any of that for you will sound banal at best, false at worst.

PART III · OUTWARD

Tools that reach the public

h) Public tools that advance the mission

AI lets us access and use data in ways that were not possible before — or at least not at a cost we could bear. That changes things: a think tank can now build interfaces that bring data straight into the debate, without the coding skills — or the developers’ fees — the job once demanded.

These are the visible, shareable wins. A public web tool puts data directly into citizens’ hands — driving engagement, earning media, and quietly gathering supporter information. What once demanded developers, budgets and specialist skills is now something a think tank can do for itself. That alone changes what is possible.

Examples of tools built:

- a. **Open Hearings** (Grassroot Institute of Hawaii) lets people monitor hundreds of government meetings — work that once needed dozens of people.

- b. **CompareMySchool.com** (MCP) — every Mississippi public school, side by side, built on official state accountability data: any parent can see how their child’s school really performs and rank it against the rest.
- c. **TaxToolMS.com** (MCP) — shows taxpayers what they pay across federal, state and local government, and follows every dollar to where it actually goes.
- d. **CompareMyPower.com** (MCP) — shows every Mississippi household what it pays for power and how that compares, built on official EIA and Public Service Commission data — making the state’s energy advantage impossible to miss.

This is no more than a rough set of pointers, sharing our experience at MCP. We are learning new things every day, and would love to hear from you — especially if you have insights or innovations of your own to pass along.