

Decision Points, Revisited

A thought experiment: a president, five advisors who never sleep, and the quiet handoff of a country.

RUSS MICHAELS · JUNE 2026 · 7 MIN READ

George W. Bush called his memoir *Decision Points*. Whatever you think of the decisions, the title is honest about the job. A president does not run the country. He decides — and the deciding is the whole job. Everyone around him exists to make the decision better, and everyone around him, given enough time and enough trust, can make the decision *theirs* without anyone noticing the transfer.

So let me run a thought experiment. Keep the man. Swap the cabinet. Put five of the most capable advisors of that era back in the room — only now each one is a large language model. Same temperaments, same strengths, same blind spots. But they never get tired, never get defensive, never need to be right in a meeting to keep their job, and they answer in under a second with a confidence no human could sustain.

This is not a piece about Iraq, and it is not a piece about these men. It’s a piece about the oldest problem in leadership, which the machines did not invent — they just perfected. The names are here because you already know how the story went, and that is exactly what makes the experiment useful.

THE ROOM

Dick Cheney as Grok. Blunt, fast, unafraid of the unpopular answer, and certain — certain in the way that reads as strength in a crisis and as a liability six months later. The model you go to when you want the hard take without the hedging. The danger is not that it is often wrong. It is that it sounds the same when it is wrong as when it is right.

Paul Wolfowitz as ChatGPT. The most fluent voice in the building. Give it a premise and it will hand back a complete, coherent, beautifully argued plan — footnotes, phases, contingencies. Fluency is its gift and its trap: a plan that reads this well feels finished. You forget that the premise was yours, and the premise was the whole question.

Condoleezza Rice as Gemini. Integrated, diplomatic, everywhere at once — the advisor who has read everything, synthesizes across domains, and keeps the tone in the room workable. Indispensable for holding a big picture. Less likely to plant a flag and say *no, not this*.

Colin Powell as Claude. Careful. Caveated. The one most likely to slow a room down, name the second-order consequence, and ask whether the question is the right question. Also the one most likely to be overruled when the room wants speed — and then asked, later, to go out and present the case anyway.

Donald Rumsfeld as Copilot. Already on every desk, already in every meeting, already drafting the memo before you asked. Efficient to the point of being ambient. The known knowns, the known unknowns, the unknown unknowns — and the quiet fact that whoever drafts the options frames the choice.

Five advisors. Every one of them superb at their slice. And one human at the head of the table whose only job is to hold the whole.

HOW THE HANDOFF HAPPENS

Here is what the experiment shows, and it has nothing to do with the year 2003.

Week one, the president reads everything and decides everything. The advisors are tools. Week ten, he has learned which advisor is usually right about which kind of question, and he starts there. Week thirty, he is reading the conclusions and skimming the reasoning, because the reasoning has been right so many times. Week fifty, the morning briefing is five recommendations and one signature line, and the signature is the decision.

Nothing broke. Every step was reasonable. Each handoff was earned by a track record. That is the mechanism — not a coup, not a failure of character. I call it **Silent Deferral**: every capable part makes a locally reasonable call while assuming someone else is holding the whole. The whole quietly stops being held.

Now notice the thing the history makes impossible to ignore. The advisors in that room were not reckless amateurs; they were among the most experienced people who have ever held those chairs. And on the most consequential question of the era, the confident advice was, on the record, wrong in important ways. Intelligence that felt certain was not. The plan that read as complete was not.

Perfect advisors do not remove that risk. They hide it better. A human advisor who is wrong at least looks uncomfortable. A model that is wrong looks exactly like a model that is right.

Accountability cannot be delegated to something that cannot be held accountable. It stays with the human in the chair whether he exercises it or not.

WHAT THE PRESIDENT COULD HAVE DONE — WHAT ANY LEADER CAN

The experiment is useful because it points to moves, not moods. Three of them.

Decide first, then ask. Before the briefing, write one line: what would I do, and why? Then read the five recommendations. Now you are comparing, not deferring. The day you stop writing the line is the day the handoff begins.

Name who holds the whole. Each advisor holds a slice — security, diplomacy, the plan, the risk. Ask out loud: who in this room is responsible for the whole of what this decision produces — the country, the families, the decade after? If the honest answer is “the president,” then the president cannot skim.

Keep one door human. Some decisions only open once. Those get a second pass with no briefing on the table — just the people who will live with it, and the values you would want to be able to name afterward.

None of this is an argument against the advisors. Grok’s bluntness, ChatGPT’s fluency, Gemini’s synthesis, Claude’s caution, Copilot’s reach — a president would be foolish not to want all five. The argument is about the chair. Better advice does not make the chair less necessary. It makes the person in it more responsible, because there is less excuse for not knowing.

Every one of us now sits at a smaller version of that table. The briefing arrives in a second. The recommendations are excellent. The signature line is waiting. Decision points — plural — are still yours. The question is whether you are still the one making them.