



April 1, 2026

**From:** Will Falk  
**To:** Healthcare AI Watchers  
**Re:** GOVERNING GENERATIVE AI: FROM PERMISSION TO COMPETENCY TO MEASUREMENT

---

Most debate about generative AI in healthcare assumes the danger is under-regulation. That assumption is wrong.

In 2026, the greater risk is over-regulation that stops much needed system innovation. That includes approval frameworks that are not yet thought through or that were built for drugs and devices being forced onto software tools that support better clinical practice.

No jurisdiction has a stable regulatory framework for generative AI in healthcare that can be readily copied. The United States, United Kingdom, and Europe all signalled that GenAI would travel through existing medical device pathways. Each has retreated. Many have settled for self-certification, which is largely performative. There are no global best practices to import.

Canada, often criticized for moving slowly, now looks prudent. Health Canada avoided premature rigidity. It **opted** instead for AI in Healthcare Principles. OntarioMD and Canada Health Infoway built safe-harbour procurement rules that let clinicians adopt ambient scribes and use clinical decision support within existing scopes of practice. Accountability stayed human. Existing privacy, liability, and insurance structures proved adequate.

Ambient scribes reached 28 percent of Canadian physicians last August according to a CMA-CFIB survey. Self-reporting by the largest clinical decision support (CDS) vendor shows 34-percent adoption in February. No integration of personal health information or electronic medical records yet, but widespread adoption. Physicians report relief from forced digital health data clerk drudgery and a large reduction in unpaid “pajama time.”

The tools are now in wide use and spreading fast towards being the standard of practice. Clinicians and patients love the technology. Dozens of studies and more emerging every day that support this. Users are voting with their feet because we allowed it and didn’t nanny state them.

That was the right choice because these AI tools are complements, not substitutes. They draft notes. They summarize evidence on a second screen. They prepare notes and referrals.

They do not diagnose, prescribe, or override professional judgment. The regulatory question that matters is not “does this system use AI?” It is “does it substitute clinical judgment?” A supervised AI complement should not face the same approval burden as an autonomous diagnostic system. Canada has recognized this simple fact.

Forcing complements into drug or device frameworks would do three negative things in the name of safety.

First, it would delay tools that are already cutting documentation time and reducing physician burnout. Second, while it would not stop AI use, it might push it underground or into grey usage areas. The UK Royal College of Physicians documented this in late 2025: Without approved second-screen CDS tools, British clinicians turned to ChatGPT and other foundation models that have no guardrails and are less safe, not medically curated, and open to anyone. Third, it would lock Canada’s rules around a technology that moves faster than any approval cycle.

In addition to these healthcare specific negatives, over-regulating would give up Canada’s head start on generative AI technologies. We have world leading technologists; we need to support entrepreneurs not hamstring them. It is bad industrial policy to shut down innovation. And in an area like healthcare with almost 6 million Canadians without primary care it is unconscionable. We need both the new industry and the new healthcare capacity.

We can manage these tools safely. Healthcare already knows how to govern intelligence that learns on the job. Nobody approves a medical student once and walks away. Even fully licensed physicians have continuing medical education requirements. Supervisors observe, evaluate, audit, and grant autonomy as competence grows.

A competency-based framework could apply the same logic to artificial intelligence as it already does to humans. Measure error rates. Track override frequency. Test for bias across populations. Tie the regulatory burden to how much clinical autonomy the system actually exercises. Regulate performance in context, not compliance against fixed pre-market specifications. And certainly against dystopian edge cases.

Canada should take the win. Its approach has been incremental, principled, grounded in professional accountability. We let complementary AI reach clinicians while jurisdictions that reached for tighter control produced neither safety nor clarity. The task now is to formalize what works, not to import what has failed elsewhere.

---

*Will Falk is a Fellow at the C.D. Howe Institute and three other Canadian think tanks and universities and a contributing editor to [Canadian Healthcare Technology](#), where a version of this Memo first appeared.*

*To send a comment or leave feedback, click [here](#).*

*The views expressed here are those of the author. The C.D. Howe Institute does not take corporate positions on policy matters.*