

AI Could Make Governments Slower

Sep 11, 2026 **SAMI MAHROUM**

— Please log in to listen to this commentary

0:00 / 0:00 1x Audio unavailable

ABU DHABI—AI-generated writing has become ubiquitous since ChatGPT's release in 2022, and institutions are struggling to keep up. Nowhere is the pressure more evident than in the courts. In the United States, the share of federal lawsuits filed without a lawyer rose from an 11% average in 2005–22 to 16.8% in 2025, with researchers linking the surge to the advent of AI. The same study found that the share of federal civil complaints containing machine-written text rose from 1% in 2023 to 18% today.



Politics

Have the Terrorists Won?

Timothy Snyder thinks the 25 years since the 9/11 attacks have played out much as the attackers would have wanted.



A government agency that can process twice as many public comments may soon find itself receiving twice as many. The safety researcher David Woods calls this the law of stretched systems: when a system is pushed to the limits of its capacity, each

In July, the US Government Accountability Office warned that the Internal Revenue Service risks being swamped by AI-generated public comments on tax regulations. Because large language models can generate virtually unlimited variations on the same text, it noted, nearly identical comments are increasingly difficult to detect.

Research funders are being similarly inundated. Applications for grants across 12 funding organizations in the United Kingdom and Europe have risen by 57% since 2022, fueled by a surge in AI-assisted submissions.

In response to the flood of AI-generated writing, governments have begun automating their end of the process. An OECD survey of 58 national tax authorities found that 69% were using AI in 2023, up from just 9% in 2016. As AI's capabilities have advanced, so have its uses, with tax authorities now relying on the technology to support analytical work and case selection, improve taxpayer services and correspondence through virtual assistants, and automate high-volume, repetitive tasks. Make your inbox smarter. Select Newsletters But greater capacity invites greater demand.

improvement is quickly exploited to operate at greater scale. Woods pairs that concept with what he calls “the substitution myth,” which holds that automation simply replaces human workers when, in practice, it changes the nature of the work and often creates more of it.

Human oversight, however, cannot be scaled as easily. The European Union’s AI Act requires high-risk systems to be designed so that “natural persons” can effectively oversee them. EU data protection law also gives individuals the right not to be subject to decisions based solely on automated processing.

Nor can companies and organizations satisfy these requirements by putting a human at the end of the process. In its 2023 SCHUFA ruling, which concerned automated credit ratings, the European Court of Justice held that what matters is whether the final decision depends on the automated assessment, even if a human formally retains decision-making authority. A signature that merely ratifies an automated decision does not amount to meaningful oversight.

Put these two forces together, and a bottleneck emerges. Everything upstream of the responsible person—search, sorting, drafting, checking, routing—gets faster, while the person does not. The deluge of AI-generated content boosts demand for human judgment, but the rate at which decisions can be made remains unchanged.

A recent engineering study illustrates the problem. Modeling nuclear reactor licensing, the authors compared today’s manual process with two automated alternatives. They found that the projected timeline even for the more integrated automated systems—roughly 15 months—did not change even as the assumed capabilities of AI improved from conservative to optimistic. The bottleneck was the queue of senior reviewers responsible for final approvals on the most complex cases, who could complete only about three a day. Making the machines twice as capable, in other words, would do little to speed up the review process.

Governments have so far used AI for searching documents, summarizing cases, sorting comments, and drafting correspondence. These are natural tasks for AI because they can be automated without delegating consequential decisions to machines. But when decision-making capacity does not expand at the same rate, automation merely shifts the bottleneck rather than removing it, potentially even making it worse.

The alternative is not to let machines decide, which the law prohibits in certain cases and which is generally unwise, but to use them to reduce the cognitive burden of reaching a decision while leaving the judgment itself to a human.

British psychologist Alan Baddeley’s model of working memory offers a useful way to think about this. Some parts of our working memory hold and process verbal and visual information, while another part, which Baddeley calls “the central executive,” directs and allocates our limited attention. AI can already read, organize, compare, and present information, but deciding what deserves attention and exercising judgment remain distinctly human functions.

The goal of AI adoption, then, should be to make it easier for public officials to make decisions by developing tools that identify trade-offs and uncertainties, present competing arguments fairly, and acknowledge their own limitations. Alas, such a system is harder to build than a summarization tool and makes for a less impressive sales pitch, since its output would look more like better briefs than an excuse to eliminate entire departments.

This problem is not new. In 1983, cognitive psychologist Lisanne Bainbridge warned of the “ironies of automation”: as machines take the easier cases, humans are left with the difficult ones. While Bainbridge was writing about industrial process-control systems, the same logic applies to AI. Government agencies measure applications processed, cases closed, response times, and, increasingly, tasks automated. They don’t measure how much more decision-makers are being asked to handle.

The more work AI models can process, the more they leave for humans to decide. Unless governments account for this imbalance, they may discover that they have automated everything but gained little.

SAMI MAHROUM

Sami Mahroum, Founder of Spark X, previously held posts at INSEAD, the OECD, and Nesta.

<https://prosyn.org/aNojY64>

[Support Project Syndicate](#)

[Subscribe](#) [Upgrade](#) [Donate](#)

[Make your inbox smarter. Select Newsletters](#)

© Project Syndicate - 2026 [Privacy Policy](#) [Terms & Conditions](#) [Cookie Policy](#)