

By YANG XUE

As AI becomes increasingly involved in decision-making, enforcement, and governance, concern often centers on whether it might violate the law. Christoph Kletzer, a professor of law and vice dean of International and Executive Education at King's College London, argues that the greater danger may lie in the opposite direction: AI could enforce the law too strictly. Traditional legal orders, in his view, have always left a gap between norms and their enforcement—a space in which human judgment, discretion, and mercy can operate. As automation promises increasingly “frictionless” enforcement, preserving that space may become a central task for public law.

Turning to the institutional design of global AI governance, Kletzer offers a series of broader judgments from the perspective of legal philosophy: Authority should develop from the bottom up through reciprocity and inclusion; regulatory paradigms across countries need not converge—what matters is a shared descriptive vocabulary; and law is not simply lagging behind technology, but operates across multiple layers moving at different speeds.

Preserving space for legal discretion

CSST: As AI systems become increasingly capable of making or implementing decisions with direct physical consequences, does this challenge some of the conceptual foundations of existing legal systems?

Kletzer: Many people worry that AI systems will break the law. My worry runs the other way. I think they will enforce it far too well. Seen that way, the robot that disobeys or the self-driving car that crashes is a small problem. The real issue is the system that applies the law too well. When software denies the loan, stops the vehicle, flags the traveller, or freezes the account, enforcement has stopped passing through human hands altogether.

The loss here is easy to miss. Between a norm and its application there has always been a gap: the moment in which a human being reads the rule against the case. Surprisingly, much of what we treasure in legal civilization lives in that gap. Judgment, interpretation, discretion, mercy, and even the disobedience that puts a rule to the test. The US Constitution, for instance, deliberately makes decisions hard, slow, and inefficient by

Perils of perfect AI enforcement: Why law needs its gaps

balancing powers between political actors. In one reading, these frictions of enforcement—the slowness, the cost, the need to justify a decision before a judge—are inefficiencies to be regretted. In another, they are an important constitutional safeguard. Until now, no one had to design them, because nobody could remove them. Automation does just that. Perfect, instantaneous, frictionless enforcement is something no legal order in history has ever had to survive.

So what is under the greatest pressure is a very old distinction: the one between what the law says and what actually happens, between the validity of a norm and its efficacy. Code has no gap. I certainly do not want to draw any Luddite conclusion from this. Automation is coming and much of it will be welcome. But the task of public law in this era is to rebuild the gap deliberately: rights of contestation, human checkpoints in the chain of force, the guaranteed possibility of demanding justification. What friction gave us for free, constitutions must now supply by design.

Authority, reciprocity, and China's role in global AI governance

CSST: In global AI governance, there appears to be a growing gap between power, which is concentrated among a relatively small number of technology companies and states, and legitimate authority, which depends on broader international recognition. In your view, where should the authority for an effective global AI governance system come from?

Kletzer: In the international order, authority accretes from below, through use. Most of the world's institutions that actually work were created that way, and almost none of them were planned. The medieval law merchant grew the same way, out of traders recognizing one another's customs centuries before any sovereign got round to blessing the arrangement. Authority of this kind is earned transaction by transaction. And one feature keeps it honest: recognition can be withdrawn.

So, where should the authority come from? From reciprocity, and



Christoph Kletzer Photo: PROVIDED TO CSST

from inclusion. Reciprocity supplies the discipline: Whoever games the shared tests or falsifies the shared reports loses recognition, and that loss has to cost something real, or the whole arrangement is theatre.

CSST: China has proposed the Global AI Governance Initiative and, more recently, supported the establishment of the World Artificial Intelligence Cooperation Organization (WAICO). Having participated in the 2026 World AI Conference in Shanghai, how do you assess China's efforts to advance global AI governance? What role do you think China can play in shaping a more inclusive and representative framework for international AI cooperation?

Kletzer: I can give you the assessment of a legal philosopher, which is an assessment of institutional design. On design, two things struck me in Shanghai.

The first is the noun. Look at what the proposed body is called: a cooperation organization. Since I do not expect a world regulator to arrive, in AI or anywhere else, I think the shape is right. Order among sovereign equals has always been built sideways, through networks of recognition. Any institution that wants to stand a chance of success in this field has to be honest about that fact.

The second is the emphasis on inclusion, where a point often goes missing. Governance by mutual recognition carries a hard precondition:

You can recognize another jurisdiction's tests only if that jurisdiction can run tests. Capacity building across the globe is therefore a structural requirement. Whoever builds evaluation capacity in a hundred countries creates the very possibility of a governance network that reaches past the advanced few.

The test for any such organization is openness under strain, whether rivals can sit inside it, and whether recognition within it must be earned and can be withdrawn rather than being conferred once and forever. Should WAICO become the room in which those disciplines operate, it will have achieved more than a decade of treaty drafting.

Building common language across differences

CSST: Different countries and regions are developing distinct approaches to AI regulation, reflecting differences in legal traditions, political systems, technological capabilities, and social priorities. How should the international community manage these differences while avoiding regulatory fragmentation?

Kletzer: This was my subject in Shanghai, so I can be brisk. The word fragmentation smuggles in a fiction: That unity was the natural condition and we have somehow lost it. But there was never a unity to fragment. What exists are different governance paradigms. Europe writes a comprehensive code in advance. America lets the technology run and litigates afterwards. China proceeds by targeted measures, piloted, tested and revised. Different answers to one question: Who may act, at what moment, on what evidence. They are not converging, and asking them to converge is like asking every country to adopt the same electrical plug. It just will not happen.

Lawyers have known this for centuries. Private international law connected the courts of the world by asking one question of every case: Whose law applies. It never had to ask what the common law of mankind should be. So the work ahead is adapter-building. Conflict-of-laws rules for

data and for models. A shared vocabulary for describing AI failures, so that an incident observed in one country is legible in all the others. Rivals cannot share rulebooks. Rivals can share descriptions, a vocabulary. Everything I propose for AI is a variation on that discovery.

CSST: AI is evolving at a pace that often outstrips law and international regulatory frameworks. From the perspective of legal philosophy, how should we understand and address this temporal mismatch? Does effective AI governance require new approaches to the pace and adaptability of law, or are there good reasons for law to remain deliberately slower than technological change?

Kletzer: Both, and the trick lies in knowing which layer you are standing on. Law runs on several clocks, geared together. The technical layer—the standards, test protocols, and reporting formats—must move at something close to the speed of the technology, and it can, since specialists administer it and revise it without much fuss. The layer of principle should be slow, and confidently so.

The lament about law falling behind is very old and has always been half wrong. Law never kept pace with the sea, with finance, or with flight. It governed all three anyway, and it did so by declining to chase them. A legal order works by standing where the traffic must pass: At the port, at the exchange, at the border, at the point where a machine's decision falls on a particular person. Those crossing points move far more slowly than the technology does, which is why law standing there always looks slower than the traffic and always turns out to have been waiting for it.

What is new and rather encouraging is that the world is currently running several experiments in fast-layer adaptability under slow-layer stability, and running them from opposite ends. Britain holds its principles constant and lets sectoral regulators adapt beneath them. China holds strategic direction constant and lets targeted measures iterate. What strikes a legal philosopher is the structure the two share: stable frame, revisable content. That is what adaptive law actually looks like. Legislation at the speed of code is a fantasy, and the real hazard of this moment is its opposite number, law made at the speed of panic. Statutes cannot be retrained overnight the way models can.

