

Kluwer Arbitration Blog

Soft Law for Anxious Minds and Imaginary Problems: Why AI-Arbitration Guidelines Arrived Too Soon

Leonardo F. Souza-McMurtrie (University of Cambridge) · Saturday, June 14th, 2025

The thirty-three identical towers that once loomed over St. Louis, Missouri, collapsed in a grey cloud of dust. They were once the future. In the 1950s, rationalist architect Minoru Yamasaki envisioned these high-rises as a utopia in raw concrete. Hundreds of families were evicted from their self-built homes to fulfil his vertical dream. Every detail had a purpose. Walkways stretched along each floor—his “streets in the sky”—designed to create community. Elevators skipped floors to force residents to meet in the stairwells. Yamasaki poured his ideology into steel and cement.

But habit cannot be fabricated. Community requires spontaneity, time, and repetition. Yamasaki’s drawings made no room for these. His social logic was abstract, ideal, and imposed from above. Within four years, the corridors became dangerous, the skip-stop elevators helped thieves more than neighbors, and the towers became one of the city’s most dangerous places. Dystopia. The demolition came to be known as *the day modern architecture died*.¹⁾

There are limits to what theoretical purity alone can achieve when ruling the unknown, and there are consequences when we try. I am not enthusiastic about the AI-Arbitration soft law issued between 2024 and 2025. What follows is a meditation on the virtues of waiting.

The AI Guidelines

From 2023 to March 2025, many have attempted to regulate the use of artificial intelligence in arbitration: the [American Arbitration Association–International Centre for Dispute Resolution \(AAA-ICDR\) Principles](#), the [Silicon Valley Arbitration & Mediation Center \(SVAMC\) Guidelines](#), the [Chartered Institute of Arbitrators \(CIArb\) Guideline](#), and the [Vienna International Arbitral Centre \(VIAC\) Note](#). All stress fairness, transparency, competence, and procedural integrity.

At first glance, they all appear to agree. AI is treated as both promising and risky. Arbitrators must retain control. Confidentiality must be preserved. Parties should understand the tools they are using. Human judgment must not be outsourced.

Yet this apparent consensus is misleading.

For instance, the guidance on disclosure diverges significantly. SVAMC states that disclosure is

not generally necessary, treating it as a case-by-case exception. CIArb implies that disclosure may be required, particularly where AI impacts evidence or procedural fairness. According to CIArb, arbitrators may also impose disclosure themselves. VIAC takes the middle path: disclosure should be discussed at the case management conference but is not assumed. AAA-ICDR avoids the issue altogether. What seems like a shared principle dissolves into three divergent recommendations and one omission.

The same can be seen in the approaches to the use of AI by arbitrators. CIArb is cautious. Arbitrators are encouraged to consult the parties before using any AI tool, and refrain if the parties object, unless the tool is explicitly ‘low risk.’ SVAMC is more permissive. Consultation is only needed when relying on AI-generated content outside the record. VIAC leaves it to the arbitrator’s discretion. They may inform the parties and invite comments but are not required to do so. AAA-ICDR offers high-level principles, but no procedural rules.

Thus, depending on the guideline followed, an arbitrator using the same AI tool might require party consent, mere notification, or no disclosure at all. It depends not on the facts, but on which guidance they prefer.

Even the definition of AI resists consensus among the different guidelines. CIArb and SVAMC offer definitions which do not align. VIAC, perhaps wisely, avoids defining it altogether.

What remains is the appearance of agreement without its substance. The guidelines repeat values (non-delegation, competence, transparency, impartiality, etc.) that no one would dispute. But they offer little in the way of action. When they do, each standard is paired with an open-ended exception. Disclosure should happen, unless it doesn’t; arbitrators should consult, unless they don’t and so on, in a cluster of standards that are, perhaps, too soft even for soft law. But ambiguous enough to allow for plausible deniability.

This is not the fault of the institutions or the drafters. The problem is more fundamental. The AI guidelines fail to present best practices because, in truth, there are none to draw from yet. The technology is too new, still developing, and barely any court or arbitration cases exist. These guidelines can only respond to imaginary problems and, thus, are limited to Yamasaki-like solutions.

Good Soft Law Comes From Experience

Soft law begins with observation: We watch what is done, find the pattern, and then name the rule.

The [IBA Guidelines on Conflicts of Interest](#) are a good example. Far from being a top-down imposition, they were built on decades of experience. [Drafted between 2002 and 2004 by a working group of 19 practitioners, they drew on 13 national reports detailing the actual laws and customs of jurisdictions as diverse as Switzerland, Mexico, and Singapore.](#) The group compiled real-world issues (bias standards, disclosure practices, the legal status of arbitrators) and drew up lists of recurring practical scenarios. These were categorized by relevance and risk into the now-familiar Red, Orange, and Green Lists.

Institutions like the ICC, LCIA, ASA, and SCC were invited to comment. Drafts were debated in public sessions in Durban and San Francisco. The guidelines were even stress-tested by arbitrators

trained to spot edge cases. Every detail, from the 30-day disclosure window to the distinction between waivable and non-waivable conflicts, was drawn from actual cases.

The IBA Guidelines refined the approach pioneered by the 1987 [IBA Rules of Ethics for International Arbitrators](#), which themselves were shaped by the lessons of 20th-century arbitration, particularly in the post-WWII era. It was prudence building on prudence. Their strength lay in distilling the productive chaos of practice into standards for action. As a result, by 2015, approximately 60% of practitioners regarded the IBA Guidelines as highly effective (see [here](#)).

The AI guidelines invert the order: they prescribe norms in the absence of experience.

This is where complexity theory helps.²⁾ In well-functioning systems, information flows from practice. Systems evolve through fragmented trial and error so that the norms that endure are those shaped by the very practices they govern. Arbitration is no exception. Like any system, it evolves in ways that are impossible to predict. The impulse to regulate now stems from a deep, unstated belief that we already understand what is to come. But as Niels Bohr famously observed, prediction is very difficult, especially about the future. Premature regulation risks disrupting this process, fixing what isn't broken, and in doing so, preventing something better from coming into existence.

The IBA Guidelines, by contrast, are rich in detail because they drew on a rich body of practice: Should co-counsel in the same chambers be treated as a conflict? Which institutional affiliations matter? Can a law firm's prior contact with a party be imputed to an arbitrator? These were not hypothetical questions. They arose from real disputes, and the resulting norms reflect that grounding.

The same holds if we borrow from Friedrich Hayek: knowledge is dispersed. No committee, however expert, can match the insights of practitioners responding to the day-to-day. Common law did not begin with principles; it began with decisions. The *lex mercatoria* grew out of custom, not doctrine. So did international law. For all its abstract elegance, even civil law began with compilation: the Justinian Code organised what people already did.

Systems that thrive do so from the ground up.

Premature soft law is risky. If not grounded in real practice, guidelines become either too open-ended to matter or impose unnecessary duties. They may require procedures few can justify, or focus on problems that aren't relevant, and miss the ones that are. For example, the "black box" problem of AI (see previous blog posts [here](#) and [here](#)) is mentioned in some guidelines, but is it a real problem for arbitration, or simply mentioned just in case? Only time can tell.

The result: ambiguous duties, inconsistent standards, and less flexibility. Such guidelines may offer a false sense of security, as if there are best practices where none yet exist.

As of 2025, no body of practice supports authoritative guidance on AI in arbitration. There are simply not enough cases. These guidelines attempt to control what does not yet exist, with standards not yet tested. And in doing so, they miss what soft law does best, which is turning memory into insight.

What Does It Mean To Be “Soft”?

Soft law has no coercive power, enforcement mechanism, or claim to democratic legitimacy. Technically, it carries no more legal force than this very blog post. Its authority lies not in its form, but in its substance. It reflects what people do, and what experience shows to be wise.

The AI arbitration guidelines are something else entirely. They may serve a signaling function, and their intent is laudable. But we must be careful not to confuse them with lived, experiential soft law.

I understand the generalized fear of AI among lawyers and how that fear pressures institutions. But anxiety is precisely the search for solutions to problems that do not yet exist—problems in a future that may never come. Imaginary worries feel urgent: “*What if everyone is using AI but there are no cases because no one discloses it? We need to do something now.*” Yet this same reasoning could have applied to conflicts of interest in the 1960s. Back then, too, the undisclosed evil of the unknown could have justified premature regulation. But the IBA waited. Practice emerged in cases, then commentary, then public discourse, and only then did soft law follow. We are better for it.

Speculation has its place, and I’ve indulged in it too, as you can see [here](#), [here](#), and [here](#). It belongs to the collective brainstorming from which practice will eventually distil what matters. But regulation, even at its softest, should be the product of this process, not another voice in the cacophony.

My point is simple: anxiety is a poor legislator. It pushes rules without grounding and may end up strangling that which it seeks to control.

A different posture would serve us better: curiosity over control. Not passivity, but attentiveness. In short: softness. The most responsible move may be to let practice evolve in all its messiness and contradiction—and only then ask: *what have we learned?*

Good regulation takes time. Sometimes, the wisest action is simply to pause, to breathe, and to let things be.

To make sure you do not miss out on regular updates from the Kluwer Arbitration Blog, please subscribe [here](#). To submit a proposal for a blog post, please consult our [Editorial Guidelines](#).

2024 Future Ready Lawyer Survey Report

Legal innovation: Seizing the future or falling behind?

Download your free copy →



Wolters Kluwer



References

“Modern Architecture died in St. Louis, Missouri on July 15, 1972 at 3.32 p.m. (or thereabouts) when the infamous Pruitt-Igoe scheme, or rather several of its slab blocks, were given the final coup de grâce by dynamite.” Quote by architectural historian Charles Jencks, in his 1977 book ‘The

?1 Language of Post-Modern Architecture’ commenting on the demolition of the Wendell O. Pruitt Homes and William Igoe Apartments, known as the Pruitt–Igoe joint urban housing projects. I want to thank Ada Souza-McMurtrie, Doctoral Researcher in Architectural History at the University of Edinburgh for the advice on the points of architecture and urban planning.

The evolutionary nature of systems through trial and error is a theme in Karl E. Weick, Sensemaking in Organizations (Thousand Oaks, CA: Sage, 1995). The unpredictability of such

?2 evolution is detailed in the study of complex adaptive systems, for which John H. Holland provides a great introduction in Hidden Order: How Adaptation Builds Complexity (Reading, MA: Addison-Wesley, 1995).

This entry was posted on Saturday, June 14th, 2025 at 8:03 am and is filed under [Arbitration Institutions and Rules](#), [Artificial Intelligence](#), [Soft Law Instruments](#), [Technology](#)

You can follow any responses to this entry through the [Comments \(RSS\)](#) feed. You can leave a response, or [trackback](#) from your own site.