Oxford Institute
of Technology and Justice

PERSPECTIVES

BRAZILIAN CHOICES IN REGULATING GENERATIVE AI

Isabella Ferrari
Regional Federal Court of the 2nd Region in Rio de Janeiro

BRAZILIAN CHOICES IN REGULATING GENERATIVE AI

By Isabella Ferrari

A Perspective on the Brazilian National Council of Justice Resolution 615, of 2025

"Blessed are those who hunger and thirst for justice, for they will be satisfied."
— Matthew 5:6 (New Living Translation)

I. INTRODUCTION

For decades, judiciaries worldwide have searched for ways to manage massive caseloads while delivering justice to those who seek it. Brazil's situation is particularly dramatic, even by international standards. The country faces a backlog of 80 million pending cases. Settlement rates hover below 8%. Brazilian judges render an average of 9 decisions per working day—a staggering output that speaks both to judicial productivity and to the relentless pressure of demand.

Against this backdrop, artificial intelligence (AI) has emerged as a potential tool for managing litigation at scale. Brazil has been at the forefront of this development. The country was among the first to regulate AI in the judiciary, adopting National Council of Justice (CNJ) Resolution 332 in 2020 to govern predictive AI tools. This foundational framework was replaced in 2025 by Resolution 615[1], incorporating comprehensive provisions on generative AI alongside updated rules on predictive systems.

Brazil's approach offers an interesting contrast to other regulatory efforts—the EU AI Act's risk-based framework, US executive orders, and evolving frameworks elsewhere. Resolution 615 reflects distinctive regulatory choices: institutional

[1] The full text of the National Council of Justice Resolution 615/2025 is available in Portuguese at this link: https://atos.cnj.jus.br/files/original155530202503146 7d4517244566.pdf

capacity-building alongside individual flexibility, accountability without technological determinism, transparency to institutions without compromising judicial independence. These choices respond to both universal concerns about AI in justice systems and Brazil's specific context, shaped by the urgent need to deliver justice at scale.

This piece examines five key regulatory choices embedded in the Brazilian National Council of Justice Resolution 615[2] and analyzes their implications for judicial practice and AI governance. It is written from two complementary perspectives that I embody myself: that of a federal judge with thirteen years on the bench—who uses AI tools daily and understands their practical impact—and that of an academic who, since 2016, has studied both the promises and the risks of AI in the judiciary, culminating in a doctoral thesis on algorithmic discrimination within the Judiciary.

Taken together, these perspectives ground a reflection on how the Brazilian judiciary is approaching AI amid uncertainty, with particular attention to some of the most sensitive and distinctive features of the regulatory sandbox Brazil has elected to create for this domain.

II. FIVE KEY REGULATORY CHOICES

1. The Dual-Track Approach: Institutional Systems and Private Subscriptions

Resolution 615 establishes a preference for court-provided solutions (Art. 19, §1). Institutional generative AI systems have been deployed across major Brazilian courts: Assis at the Rio de Janeiro State Court, ChatJT in the Labor Justice system, Logos at the Superior Court of Justice, and Maria at the Supreme Court. Although their scopes and technical architectures vary, all of these tools are designed for internal use—supporting judges and Justices, court staff, and law clerks in performing core judicial functions.

[2] The National Council of Justice (Conselho Nacional de Justiça - CNJ) is an administrative body within the Brazilian Judiciary, headed by the Chief Justice of the Supreme Court. Brazil's judicial system possesses significant self-regulatory capacity, which enabled the Judiciary itself to adopt these regulations rather than waiting for legislative action. CNJ Resolutions are administrative acts, not statutes—a deliberate choice that offers regulatory flexibility. Unlike legislation, which requires lengthy parliamentary procedures to amend, Resolutions can be updated relatively quickly as technological developments and practical experience demand. This institutional structure has proven particularly valuable in governing rapidly evolving technologies like artificial intelligence, where rigid legislative frameworks might quickly become outdated.

ChatJT, the Labor Justice system's tool, deserves closer attention. Developed with generative AI technology, its database contains all Labor Justice decisions—from first-instance trial courts through the Superior Labor Court. Importantly, ChatJT has access to case data, including procedural documents and case docket entries. This architecture allows judges and staff to interact with both the case law database and their own case files in ways that reduce the hallucinations that would occur using general-purpose tools without specialized training.

By November 2025, ChatJT had already registered over 26,000 users and more than 1 million conversations. Use of the tool is optional, yet adoption has been remarkably strong from the outset. Users can create their own custom GPT assistants to help with specific tasks. In just eight months since its March 2025 launch, the platform had been used to create nearly 17,000 assistants. When users find an assistant valuable, they can choose to share it with colleagues. By November 2025, over 429 assistants had been shared across the system. The three most-used virtual assistants are Falcão (jurisprudence search), the Legal Summary Assistant, and the Headnote Assistant[3-4]. These numbers reveal not only strong adoption but significant user participation in shaping the system itself. Generative AI inherently requires active user participation in the technology's development—and ChatJT's architecture embraces this by enabling judges and staff to become co-creators rather than mere consumers of the tool.

As is typical with Brazilian judicial AI projects, the project leader is a judge who is also a programmer, supported by the tribunal's internal IT teams or those of the National Council of Justice.

When institutional solutions are unavailable, however, judges may directly contract with commercial models—ChatGPT, Claude, and others—through private subscription (Art. 19, §2). This authorization reflects a pragmatic acknowledgment of how generative AI actually entered the Brazilian judiciary. Judges who had

[3] In the Brazilian Labour Justice system, many judicial decisions—especially those issued by appellate labour courts—are accompanied by an ementa, a concise and structured headnote that synthesizes the key legal issues, factual background, and core reasoning of the judgment. Although first-instance decisions are not usually drafted with a formal headnote, the ementas produced by the higher courts are widely used as the main reference for legal research, precedent navigation, and jurisprudential consistency. The AI Headnote Assistant is an artificial-intelligence tool that automatically generates draft ementas (headnotes or case summaries) from the full text of judicial decisions, identifying the most relevant elements and proposing a coherent and standardized summary of the case. The tool operates in an assistive manner: it analyses the decision, suggests a draft headnote, and the human judge remains responsible for reviewing, editing, and validating the final text.

[4] You can find more information about ChatJT in Portuguese at this link: https://ia.jt.jus.br/welcome/transparencia.html

subscribed to these tools for personal use discovered their utility for professional work and began applying them to judicial tasks. The National Council of Justice faced a regulatory choice: prohibit this emergent practice or establish guardrails around it.

Several considerations shaped this decision. Forbidding private subscriptions risked stalling innovation and losing ground already gained—particularly when judicial training schools had already offered courses on these tools, teaching judges their strengths, limitations, appropriate use cases, and potential pitfalls. Moreover, there were practical enforcement concerns. For users proficient with AI, identifying unauthorized use presents significant detection challenges. A prohibition might prove difficult to enforce effectively, potentially creating compliance issues without achieving its protective aims.

The dual-track approach thus emerged as a considered response to competing imperatives: the need for institutional control over AI deployment and the reality that technological innovation often outpaces institutional capacity. Rather than banning private tools while institutional solutions mature, Brazil permits regulated individual adoption. Judges who choose private subscriptions must meet specific conditions: training requirements, data protection standards, and full liability for their decisions.

2. Full Liability for Private Use: The Judge as Guarantor

Art. 19, §3, II establishes that judges using privately subscribed models remain "entirely responsible for the decisions made and the information they contain." Notably, Resolution 615 does not establish equivalent liability frameworks for institutional systems.

This asymmetry creates clear incentives: it encourages institutional adoption while ensuring individual accountability when judges venture beyond approved systems. The judge becomes the ultimate human-in-the-loop, never delegating decisional authority to AI.

AI is explicitly framed as "auxiliary and subsidiary"—a tool to assist, never an autonomous decision-maker. This rejects any notion that algorithmic accountability might displace judicial responsibility. When a judge chooses to use a private AI subscription, that choice comes with full ownership of the results.

3. Mandatory Training: Building AI Literacy in the Judiciary

Art. 19, §3, I of Resolution 615 establishes a clear precondition: before judges or court staff may employ large language models—whether institutional tools or individually contracted systems—they must first complete specific training on the *limitations, risks, and ethical, responsible, and efficient use* of AI.

This requirement represents a significant shift. Under the regulations that govern the judicial career, judges must complete a minimum number of training hours per semester, but they are free to choose among the courses offered by judicial training schools. There is *no general obligation* to receive instruction on artificial intelligence. Resolution 615 therefore introduces a *functional* prerequisite rather than a *career-wide* educational mandate: *AI literacy becomes mandatory only for those who intend to use AI in adjudicative work*.

Judicial training institutions—ENFAM for the general courts, ENAMAT for the labor courts, and their state-level counterparts—have quickly adapted by offering courses on both general-purpose AI and institutionally developed tools. These programs have seen exceptionally high demand, reflecting not only the judiciary's enthusiasm for adopting new technologies, but also a widespread recognition that such adoption requires structure, caution, and informed competence.

Brazil thus treats AI literacy not as optional professional development but as a *condition of access*. This approach acknowledges a fundamental tension: generative AI's intuitive interface creates an illusion of safety, even though the underlying systems can hallucinate, amplify bias, and produce outputs opaque in their reasoning. The policy message is unambiguous—the power of AI demands proportional responsibility.

Even in practice, this training is far from superficial. Judicial AI courses typically involve sustained, in-person, hands-on instruction, designed to expose users not only to the strengths of these tools but also to their operational and ethical limits[5]. What emerges is a model in which competence precedes deployment, and institutional safeguards are built not only into software but into human capacity.

[5] For example, in my own case, I completed two separate multi-day intensive trainings—each consisting of full-day sessions from 8 a.m. to 6 p.m.—focused on the practical capabilities and limitations of the tools.

In this sense, Resolution 615 reframes AI literacy as a form of judicial prudence: a skill that is neither decorative nor discretionary, but integral to accountable decision-making in a digital judiciary.

4. Data Protection: Reconciling Public Process with Privacy

Art. 19, §3, IV prohibits using private LLMs to process confidential documents or data protected by judicial secrecy, except when properly anonymized or when technical safeguards ensure effective protection.

Brazil's procedural system operates on a principle of publicness—court proceedings are generally open. Yet sensitive personal data routinely appears in case files: health conditions, financial information, family matters.

Resolution 615 establishes a compatibility framework rather than a tension. Public case management can employ AI, but sensitive data requires either institutional systems with proper safeguards or anonymization before exposure to private models. This allows the judiciary to embrace AI's efficiency gains while respecting data protection principles that govern all Brazilian institutions.

5. Transparency and Discretion: The Registration vs. Disclosure Question

Art. 19, §6 creates an interesting distinction: judges *may* mention AI use in their written decisions at their discretion, but the court's internal system *must* automatically register such use for statistical, monitoring, and audit purposes.

The case for disclosure seems straightforward. Mentioning AI use could promote transparency and help normalize these tools for the public, reducing fear and mystery around judicial AI.

But there's a counter-argument rooted in judicial tradition. Brazilian procedure does not compel judges to detail their entire decisional process. A judge need not disclose whether they discussed the case with colleagues, consulted particular secondary sources, or used specific research methods. Requiring AI disclosure could create a false dichotomy—suggesting AI-assisted decisions are fundamentally different from "human" decisions, when in practice all judicial decisions involve multiple tools and inputs.

Brazil's solution threads this needle. Mandatory backend registration enables institutional oversight and quality control. The system knows when AI has been used, creating accountability and allowing empirical study of AI's impact on judicial work. Yet judges retain discretion over public disclosure, avoiding the risk that disclosure requirements might stigmatize AI use or suggest such decisions lack legitimacy. This balances transparency to the institution with independence from potential public pressure about methodology.

This institutional—public distinction also raises a practical question: should individual judges voluntarily disclose their own AI use? At present, I do not include such disclosures in my decisions, although I make routine use of generative AI for grammar review, stylistic restructuring, and occasional brainstorming. My hesitation is not about transparency, but about systemic coherence. Because neither my court nor the broader judiciary has yet adopted a consistent disclosure practice, isolated, judge-by-judge disclosures risk producing uneven expectations and potentially distorting litigants' perceptions in a single jurisdiction. In this sense, AI disclosure functions less as an individual ethical choice and more as a systemic design question: its benefits depend on coordinated institutional adoption. Until such guidance exists, I view alignment with my court's current practice as the option least likely to introduce unintended disparities for litigants.

III. CONCLUSION: A PRAGMATIC REGULATORY MODEL

Resolution 615 reflects distinctively Brazilian regulatory priorities shaped by the country's litigation crisis and its tradition of judicial regulation. The framework builds institutional capacity while permitting individual flexibility. It demands accountability without succumbing to technological determinism. It creates transparency for institutions without compromising judicial independence in individual cases.

These choices offer lessons for other jurisdictions grappling with AI in justice systems. The dual-track model may prove more sustainable than either purely institutional or purely market-driven approaches. The emphasis on training as prerequisite recognizes that AI adoption without understanding creates new risks rather than solving old problems. The distinction between internal registration and public disclosure suggests ways to enable oversight without creating perverse incentives.

As institutional systems mature and case law develops around AI-assisted decisions, Brazil's framework will face practical tests. The extraordinarily high demand for training suggests the judiciary recognizes both the promise and the challenge of this

technological moment. Whether AI can help satisfy the hunger for justice at scale remains an open question—yet in confronting 80 million pending cases, the search for tools that might help deliver justice cannot wait for perfect solutions. While we cannot promise divine justice, we can strive to make human justice more accessible, more timely, and more responsive to those who seek it. Brazil's regulatory approach—pragmatic, flexible, accountability-focused—offers one path forward in that effort.

Regulation is itself a form of institutional learning. Resolution 615 represents Brazil's latest contribution to an ongoing global dialogue between law, technology, and judicial practice—building on five years of experience since the pioneering 2020 framework, and responding to the urgent need to deliver justice to 80 million pending cases.