WORKING PAPER Nº IDB-WP-01856

AI and Judicial Productivity: The Impact of MIDAS on the Courts of Fortaleza, Brazil

Gastón Pierri
Marcelo Fontenele
José Luiz Nunes

Inter-American Development Bank
Institutions for Development Sector
Institutional Capacity of the State Division

July 2026

AI and Judicial Productivity: The Impact of MIDAS on the Courts of Fortaleza, Brazil

Gastón Pierri
Marcelo Fontenele
José Luiz Nunes

Inter-American Development Bank
Institutions for Development Sector
Institutional Capacity of the State Division

July 2026

AI and Judicial Productivity: The Impact of MIDAS on the Courts of Fortaleza, Brazil*

Gastón Pierri† Marcelo Fontenele‡ José Luiz Nunes§

Abstract

This paper presents preliminary results from a pilot study conducted in the courts of Ceará, Brazil. The study evaluates the impact of introducing a tool that uses natural language processing and machine learning techniques to cluster judicial acts by textual similarity on clerk productivity, measured as the number of case files a clerk can produce in a day. Estimates indicate that treatment-group clerks produced approximately 10 more case files per day than control-group clerks, a statistically significant difference equivalent to a 37% increase relative to the control group mean. The results are robust to the exclusion of outlier observations and exceptionally productive clerks.

JEL codes: K40, O33, D73, C93.

Keywords: Judicial productivity; Artificial Intelligence; Natural Language Processing; Digital government; e-justice.

*The authors gratefully acknowledge the TJCE, whose commitment to innovation and digital transformation made this research possible within the framework of the PROMOJUD. We are especially grateful to President Heráclito Vieira de Sousa Neto, José Tarcílio Souza da Silva, Sirley Cíntia Pachêco Prudêncio, Marcelo Roseno de Oliveira, and Rafaella Lopes Ferreira for their leadership and institutional support throughout the development of the program and this study. We also thank all the professionals involved in PROMOJUD across its different stages of design, implementation, and monitoring. In particular, we acknowledge the contributions of the teams from Setin, Segov, Sejud 1, and the Project Management Unit (UGP), as well as the judges, court staff, project managers, and technical personnel whose expertise and dedication were essential to the design, implementation, and execution of the pilot study. At the IDB, we thank Paula Acosta Marquez, Mariano Lafuente, Martina Bergues, and Taís Pelinson for their technical guidance and valuable contributions to the design and implementation of this research. The views expressed in this paper are those of the authors and do not necessarily reflect those of the TJCE, the IDB, or their respective institutions

†Specialist in State Modernization, Institutional Capacity Division (ICS/IFD), IDB

‡First-Instance Court Clerk at the TJCE

§Researcher at the FGV Law School in Rio de Janeiro

1 Introduction

Estimates from 2019 indicated that around 1.5 billion individuals were unable to resolve their legal problems, despite living in contexts with a functioning justice system and institutions (World Justice Project, 2019). In Brazil, the State Justice system closed 2025 with nearly 58 million pending judicial processes, despite having adjudicated and closed more than 30 million during that year[1]. In the state of Ceará alone, the state Court of Justice (TJCE) received 692,630 new judicial processes and ended the year with more than one million pending.[2]

These backlogs are not merely an administrative problem, as the efficiency of the judicial system has significant economic repercussions. The literature shows that delays in resolving judicial proceedings hamper contract *enforcement* and insolvency resolution, generating a negative effect on firms' credit and investment (Djankov et al., 2008; Ponticelli and Alencar, 2016; Visaria, 2009; Chemin, 2012). Similarly, reductions in judicial processing times have been associated with greater entrepreneurship (Chemin, 2009), firm performance (Chakraborty, 2016), and the country's economic growth and development (Amirapu, 2021; Djankov et al., 2025). The need to increase judicial productivity becomes even more relevant in a context in which artificial intelligence tools are also reducing the costs of access to justice and potentially increasing the volume of litigation faced by courts (Shah and Levy, 2026).

This paper analyzes whether providing judicial clerks with a tool based on natural language processing (NLP) that clusters similar rulings, allowing them to process cases in batches rather than one by one, increases their daily productivity relative to the traditional workflow. To this end, we carried out a pilot study to examine the consequences of introducing the MIDAS system (Mecanismo Identificador de Actos Similares — Similar Acts Identification Mechanism) at the First-Degree Judicial Secretariat (SEJUD) of the Court of Justice of Ceará (TJCE), Brazil[3]. MIDAS is a tool that uses NLP and machine learning techniques to cluster judicial rulings by textual similarity. Whereas in the traditional workflow clerks must prepare each case file one by one, with MIDAS clerks can prepare multiple case files at once, thereby reducing the average time per case and increasing total output volume. This increase in productivity should,

[1] The Brazilian judiciary is divided into several branches, including Federal, Labor, Military, and State. The latter has the broadest jurisdiction and accounts for the largest share of citizen interaction with the justice system.

[2] Source: https://justica-em-numeros.cnj.jus.br/painel-estatisticas/

[3] In recent years, the Brazilian judiciary has become a world pioneer in the adoption of AI and automation, with 140 initiatives underway across 62 courts as of 2023 (da Silva and de Oliveira, 2025).

in principle, contribute to reducing the *backlog*, to the extent that it does not simply shift the bottleneck to another stage of the judicial process.

The study was conducted between June and August 2025 within the framework of the Modernization Program for the Judiciary of the State of Ceará (PROMOJUD), with support from the Inter-American Development Bank (IDB). The study compared the daily productivity of 62 clerks assigned to two groups: a treatment group that processed case files with the support of MIDAS, and a control group that followed the traditional workflow. The central element of the design was randomization at the level of the ruling or judicial act: within each *cluster* generated by the system, rulings were randomly reordered using a Python *script*, and a fraction was extracted to be processed individually by the control group. This mechanism ensures that the judicial acts assigned to both groups come from the same pool of rulings and are comparable in their observable and unobservable characteristics in expectation.

The results indicate that treatment-group clerks produced on average 10 more case files per day than control-group clerks, a statistically significant difference equivalent to a 37% increase relative to the control group mean. Furthermore, these results are robust to the exclusion of potentially outlying observations, which reinforces the robustness of the conclusions. In terms of workload, our estimates suggest that access to the tool reduces case file preparation time from approximately 16 to 12 minutes per case file, a gain of 4 minutes per case that, accumulated over the workday, allows clerks to process one-third more case files in the same amount of time. Furthermore, the tool appears to expand the right tail of the productivity distribution: days on which a clerk analyzes more than 100 case files are notably more frequent among those who used MIDAS. However, since the treatment group included both clerks with full adherence to the system and clerks with partial exposure, the estimated effects should be interpreted as a lower bound on the tool's true impact.

This paper contributes to the existing literature on the use of artificial intelligence to improve efficiency in the judicial sector (Aidid and Alarie, 2023; Casey and Niblett, 2019; Volokh, 2019). Prior studies have shown how these tools can be used to automate routine tasks, such as analyzing and processing evidence, conducting investigations, and classifying legal documents (de Oliveira and Nascimento, 2021; Oliveira and Sperandio Nascimento, 2025; Razmetaeva and Razmetaev, 2021; Solovey et al., 2025), allowing court staff to devote more time to other activities (Alarie et al., 2018; Borgesano et al., 2025). The literature also highlights their potential to assist in

judicial decision-making and, in some cases, to partially substitute judges' functions (Chen et al., 2022; Volokh, 2019). Some countries already report productivity gains derived from their use, including Argentina and Brazil (de Sousa et al., 2022; OECD, 2025). Our study contributes to this growing body of research by being, to our knowledge, the first to estimate the impact of an NLP-based clustering tool on judicial staff productivity through a pilot study with randomization at the level of the judicial act.

The remainder of the paper is organized as follows: Section 2 describes how the intervention works, the MIDAS system, and its position within the Judicial Secretariat's workflow. Section 3 then presents the study design, while Section 4 describes the data used. Section 5 presents the preliminary results of the pilot study, as well as the lessons that should be considered in future studies. Finally, Section 6 discusses the study's main methodological limitations and concludes the paper.

2 Institutional background

In recent years, the TJCE has been developing multiple digitalization and digital transformation projects, including automation tools to increase the efficiency of the work of its clerks and magistrates. Part of these efforts fall within the Modernization Program for the Judiciary of the State of Ceará (PROMOJUD), developed with the support of the Inter-American Development Bank (IDB).[4] PROMOJUD aims to increase efficiency and user satisfaction with the justice system by advancing the TJCE's digital transformation. Its key indicators include reducing case processing times, increasing clerk productivity, reducing the stock of pending judicial processes, reducing the average cost per closed judicial case, and providing digital training for magistrates and clerks. In turn, PROMOJUD's monitoring and evaluation plan includes an evaluation design to collect evidence on the effects of these measures on judicial productivity, within which the present study is situated.

[4]Project BR-L1560, co-financed by the IDB, approved in April 2021 and whose implementation began in January 2022, amounts to US$ 35,000,000 and is part of the "Brasil Mais Digital" credit line for US$ 1,000,000,000. Information on the project is available at: https://www.iadb.org/en/project/BR-L1560.

2.1 The SEJUD

One of the relevant organizational changes at the TJCE was the creation, in 2021, of the First-Degree Judicial Secretariat of Ceará (SEJUD). This body was created to centralize the case files and activities that were previously carried out within each judicial unit.[5] The internal organization of the SEJUD can be seen in Figure 1.

Figure 1. Internal organization of the First-Degree SEJUD

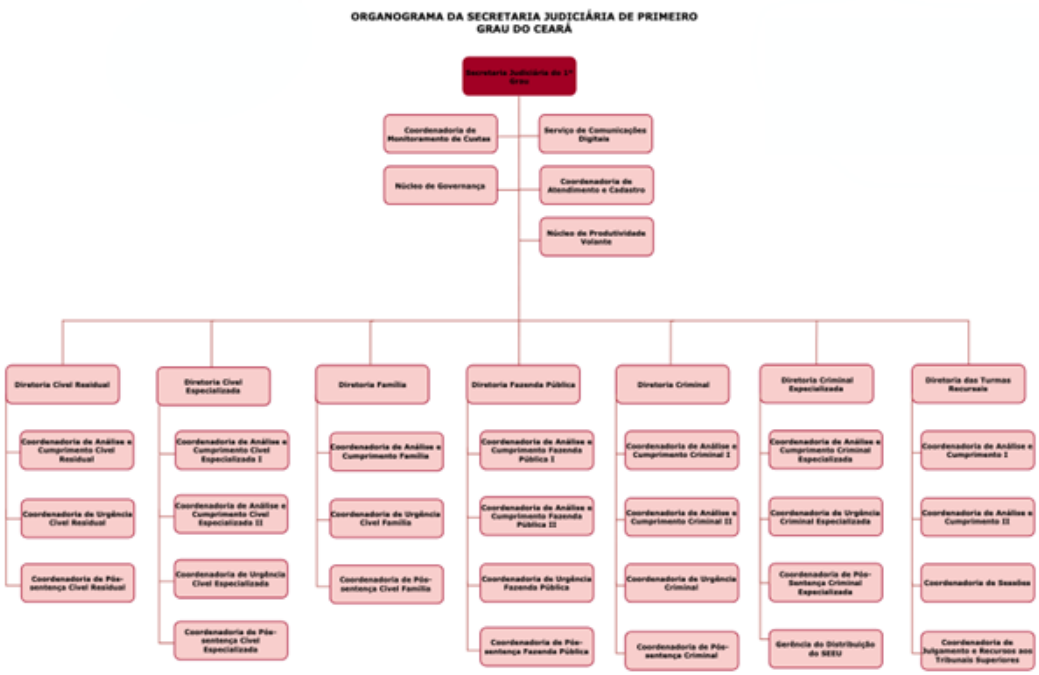

Source: Court of Justice of Ceará.

The First-Degree SEJUD of Fortaleza is currently responsible for the secretariat case files of 104 judicial units, under the supervision of a Judge of Law appointed by the Court's Presidency. Its management also comprises a First-Degree Judicial Secretary, Directors, Coordinators, and operational Supervisors. Since its creation, the unit has progressively increased its workload: the average number of case files per clerk rose from 2,222 in 2022 to 2,456 in 2024, the year in which the unit produced a total of 8,216,907 case files.

The main activity of the SEJUD consists of producing *case files*: administrative acts that must be carried out following each judicial ruling in order to enable its enforcement. It is important to note that the term *case file* does not refer here to the complete judicial record of a case, but rather to the specific administrative document that the SEJUD must generate in response to each

[5] The clerks in each judge's office were responsible for analyzing the rulings issued in their unit and producing the administrative acts necessary for their enforcement.

of the judge's decisions (for example, a summons letter, an attachment order, or a notification to the parties and/or competent authority). These are therefore acts of a largely procedural nature, whose content is to a great extent determined by the type of ruling that gives rise to them, which makes them repetitive and ideal for batch processing.

Figure 2. Simplified flow of a first-degree judicial proceeding.

Source: Authors' own elaboration.

Figure 2 illustrates, in simplified form, the flow of a judicial action at the first degree. The process begins with the receipt of an initial claim, after which the defendant must be notified. The SEJUD's role is to identify and draft the summons letter so that the party becomes aware of the existence of the proceedings and can present their defense. Throughout the course of this process, the judge will issue multiple rulings. After each of these, the SEJUD determines whether it is necessary to prepare a case file for its enforcement. During the pre-judgment phase, multiple acts by the judge and statements by the parties follow one another, until the judgment deciding the merits of the case is issued. MIDAS intervenes precisely at this stage: after each decision or order from the judge, and before the preparation of the corresponding case file.

As an example, consider the initiation of a judicial action. After the filing of the claim that opens the proceedings, the judge orders notification of the defendant through a summons order. The SEJUD identifies the type of ruling, determines the relevant case file, and drafts the summons letter, which is sent to the corresponding party. From that point on, that party is deemed to have legal knowledge of the proceedings, and the relevant deadlines begin to run. Upon completion of this analysis, the process returns to the judge for follow-up.

2.2 The MIDAS system

Another development that took place within the PROMOJUD framework was the creation of MIDAS (Mecanismo Identificador De Actos Similares — Similar Acts Identification Mechanism), a computational system whose objective is to cluster judicial acts based on the similarity of their textual content. Using natural language processing algorithms, the system generates a representation of the content of each judicial ruling pending analysis. These rulings then pass through unsupervised *machine learning* algorithms, which compare them and identify similar documents, grouping them into *clusters*, typically composed of the same type of judicial act and which, consequently, require the preparation of the same type of case file. Based on this grouping, SEJUD clerks can prepare case files in blocks of judicial rulings, thereby optimizing the work process. This process is illustrated in Figure 3:

Figure 3. Judicial ruling analysis flow before and after the implementation of MIDAS at the SEJUD.

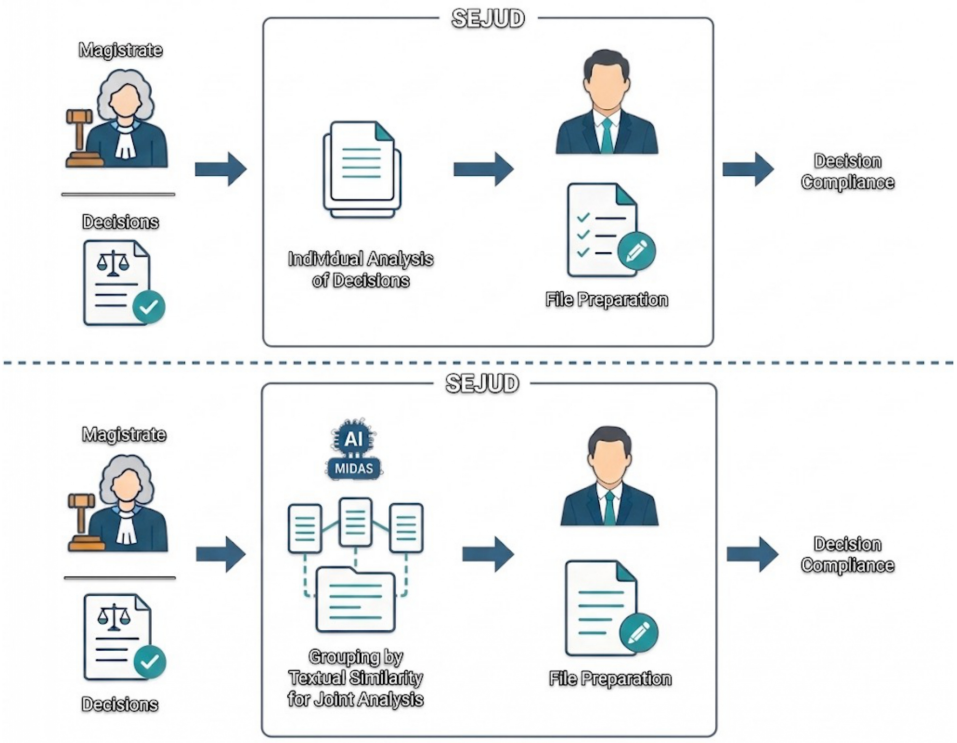

Source: Authors' own elaboration.

Before MIDAS, each clerk had to analyze rulings individually and in isolation: for each one received, the clerk would start the process from scratch and generate a separate case file, repeating the same cycle as many times as there were rulings. The result was a cumulative burden of repetitive, disconnected tasks. The MIDAS system transforms this workflow by grouping rulings before they reach the clerk, who now only needs to execute that procedural sequence once for

the entire set of rulings in the *cluster*, rather than *N* times for *N* individual rulings. This should ease clerks' workload by eliminating unnecessary procedural repetition when multiple rulings require the same administrative response. It is worth noting that, unlike other AI systems applied to the judicial sphere, potential errors in MIDAS do not directly affect the outcome of the case or the rights of the parties involved. An imperfect grouping may result in a loss of efficiency, but it does not replace or influence the final legal decision, which remains under the clerk's control.

To further illustrate how the system works, consider the scenario in which a judge issues, over the same period, ten rulings ordering the summons of different defendants in debt-collection cases. Each of these rulings differs with respect to the data of the parties involved (names, amount owed, addresses, etc.), but shares the same textual structure insofar as they are all summons orders with the same legal basis and the same type of administrative burden required. Under the traditional workflow, the SEJUD clerk would have to process each of these rulings individually: opening the case file, identifying the type of act, selecting the summons letter template in the PJe system, filling in the data of the parties involved, and registering the act, repeating this sequence for each ruling. With MIDAS, the system would detect the textual similarity among these ten rulings and group them into a single *cluster* before they reach the clerk. The clerk performs the analysis only once and generates the ten case files in a batch, filling in only the variable fields (party data) for each case. Processing time is thus reduced not because the clerk analyzes each case less carefully, but because it eliminates the nine unnecessary repetitions of an identical procedural sequence.

3 Study design

With the aim of producing preliminary evidence of the tool's impact to evaluate its expansion at the state level, a TJCE team designed and implemented a pilot study with an experimental design, in collaboration with the IDB's impact evaluation team. The study was conducted between June 30 and August 14, 2025, and involved 62 clerks distributed across treatment and control groups. The study design and the adjustments made during its implementation are described below.

3.1 Pilot design and groups

The study involved 62 First-Degree SEJUD clerks, distributed into two groups: a **treatment** group (37 clerks) and a **control** group (25 clerks). Within the treatment group, 15 clerks worked exclusively on rulings clustered by the MIDAS system, while the remaining 22 clerks worked both on rulings clustered by MIDAS and on tasks within the traditional PJe system workflow. This variation in treatment intensity within the group that had access to the MIDAS system implies that the estimated effects should be interpreted as potentially conservative relative to the tool's true impact.

Group composition was proportional across the Specialized Civil, Residual Civil, and Public Treasury Directorates of the SEJUD, with the aim of ensuring representativeness. Although the assignment of clerks to groups was not carried out through a randomization mechanism, no systematic selection criterion was followed either. In contrast, the assignment of *rulings* within each group did incorporate a formal randomization mechanism, which we describe below.

3.2 Randomization of rulings

The only formal randomization element in the design operated at the level of the ruling, not the clerk. Although MIDAS clusters and orders rulings chronologically by the date of the judicial act, a randomization tool implemented via a Python *script* — previously validated by the project's academic team — was incorporated into the workflow. This *script* randomly reordered rulings within each *cluster*, preventing chronological order from systematically influencing the distribution of cases among clerks. Additionally, the code randomly extracts a proportion of the rulings that would have been assigned by MIDAS for processing under the traditional workflow.

This mechanism ensures that the rulings analyzed by clerks in both groups come from the same pool of judicial rulings and are comparable in their observable and unobservable characteristics in expectation. Nonetheless, the fact that the assignment of clerks to groups was not random introduces a potential source of heterogeneity. This limitation is addressed in two ways. First, through random effects at the clerk level in the main specification, which control for baseline productivity differences among clerks. Second, through controls for workload composition in the robustness specifications.

3.3 Implementation and work coordination

The project was selected by the IDB to be part of the 2025 Impact Evaluation Design Week (IEDW), held in Bogotá between June 9 and 13, 2025, an instance that helped strengthen the design and train key staff for its execution.[6] Implementation was carried out by members of the TJCE under the supervision of an academic researcher from FGV, one of the authors of this study, who also took part in the training.

Clerks were not informed of the details of the experiment or of their group assignment. They were only told that they were participating in a project and were asked to minimize absences as much as possible. Each clerk received daily from their manager a list of rulings on which they were to act. No personal information was collected about the clerks or the parties to the rulings. SEJUD managers also kept a daily log of the case files produced, which served as the primary source for extracting data from the database that stores the status of each ruling within the court.

4 Data and descriptive statistics

At the end of data collection, information was obtained on 43,614 case files produced by the SEJUD, originating from 30,024 distinct processes, of which 30,393 case files correspond to the treatment group and 13,221 to the control group.

From these data, the variables used in the analysis were constructed. All collected data underwent a pseudo-anonymization process, protecting the identity of the participating clerks. Researchers had access only to a sequentially generated identifier used to distinguish clerks. At the end of the study period, the data were shared with the authors of this paper, who were part of the project's impact evaluation team together with the IDB, for statistical analysis and the preparation of results. Table 1 describes the variables available for analysis:

[6] See: https://cursos.iadb.org/en/node/1861.

Table 1: Variable description

Variable	Description
Clerk code	Sequentially generated pseudo-anonymized identifier. Contains no personal information about the clerk.
Experimental group	Group to which the clerk belonged. Values: Treatment, Control.
Judicial case number	Identifier code for the judicial process according to the standardized numbering of the National Council of Justice.
Judicial unit	Court to which the process belongs. Takes 77 possible values, corresponding to the units served by the SEJUD.
Document type	Type of case file produced from the judicial ruling. Also determines the urgency of the analysis. 16 distinct values were identified.
Priority	Urgency assigned to the case file. Values: Normal or Urgent.[7]
Judicial act type	Type of ruling that gave rise to the case file. Values: Dispatch, Judgment, Decision, Ordinary Act.
Judicial act date	Date on which the ruling was published and recorded in the procedural system.
Case file date	Date on which the case file was produced and recorded in the procedural system by the clerk.
Derived dependent variable	
Case files per clerk-day	Number of case files produced by each clerk on each working day. Constructed by grouping the data by clerk code and case file date, so that each observation corresponds to a clerk-day.

Table 2, in turn, presents the distribution of rulings assigned to the different groups, distinguishing by priority level:

Table 2: Distribution of rulings by priority and group

	Treatment	**Control**	**Total**
Average rulings per clerk	821	529	703
Normal-priority rulings	29,713	7,422	37,135
Urgent-priority rulings	680	5,799	6,479

As can be seen, there is a clear divergence between the groups with respect to the assignment of urgent-priority rulings. This is because urgent case files needed to be processed as quickly as possible to preserve the proper functioning of the unit, and there was insufficient time for MIDAS to accumulate and cluster them, so they were assigned predominantly to the control group.

[7] The distribution of case files by priority and experimental group is presented in Table 2.

Finally, Table 3 shows the balance between both groups with respect to act type and the number of days until case file closure:

Table 3: Sample description and balance between groups

	Full sample		**Normal priority only**	
	Control	**Treatment**	**Control**	**Treatment**
Act type: Ordinary act (%)	< 1%	< 1%	< 1%	< 1%
Act type: Dispatch (%)	46.7	62.2	63.5	61.9
Act type: Decision (%)	48.0	21.2	28.4	21.1
Act type: Judgment (%)	4.8	16.0	7.3	16.3
Days to case file (mean)	2.9	7.9	4.3	8.0
Days to case file (median)	1	6	2	6

Again, a certain imbalance is observed between the case files produced by both groups. First, the median number of days to case file closure is substantially lower in the control group (1 day) than in the treatment group (6 days). This could be explained by the higher proportion of urgent rulings in the control group, which, almost by definition, tend to be resolved within very short timeframes, often within the same day. Second, relevant differences are also observed in the composition by act type. The control group has a higher proportion of "Decision"-type acts (48% vs. 21%), and a lower share of "Dispatch"-type (47% vs. 62%) or "Judgment"-type (5% vs. 16%) acts compared to the treatment group. Once again, these differences may be associated with the imbalance in the priority level of the processes.

To assess the extent to which these differences are associated with the overrepresentation of urgent rulings in the control group, Table 3 also reports the same statistics restricting the sample to normal-priority rulings only. Although the difference in days to case file is partially reduced when excluding urgent cases, some differences in the composition of procedural acts across groups persist. This suggests that the observed imbalance is not explained exclusively by the different proportion of urgent case files.

Taken together, these results show that there are some differences between the control and treatment groups in terms of the workload received, particularly with respect to the priority level of rulings and certain types of procedural acts. For this reason, we include a specification that controls for workload composition (specifically, the proportion of urgent rulings and the distribution of judicial act types at the clerk-day level) as a robustness check. Crucially, when these controls are included, the estimated effect not only remains stable in magnitude but also

gains precision, which constitutes direct evidence that the imbalance in workload composition between groups does not affect the main results.

5 Main results

We begin the analysis by visualizing the distribution of the number of case files produced per working day for each group. Figure 4 shows the distribution of this variable for each group.

Figure 4. Distribution of the number of case files produced per working day in each group.

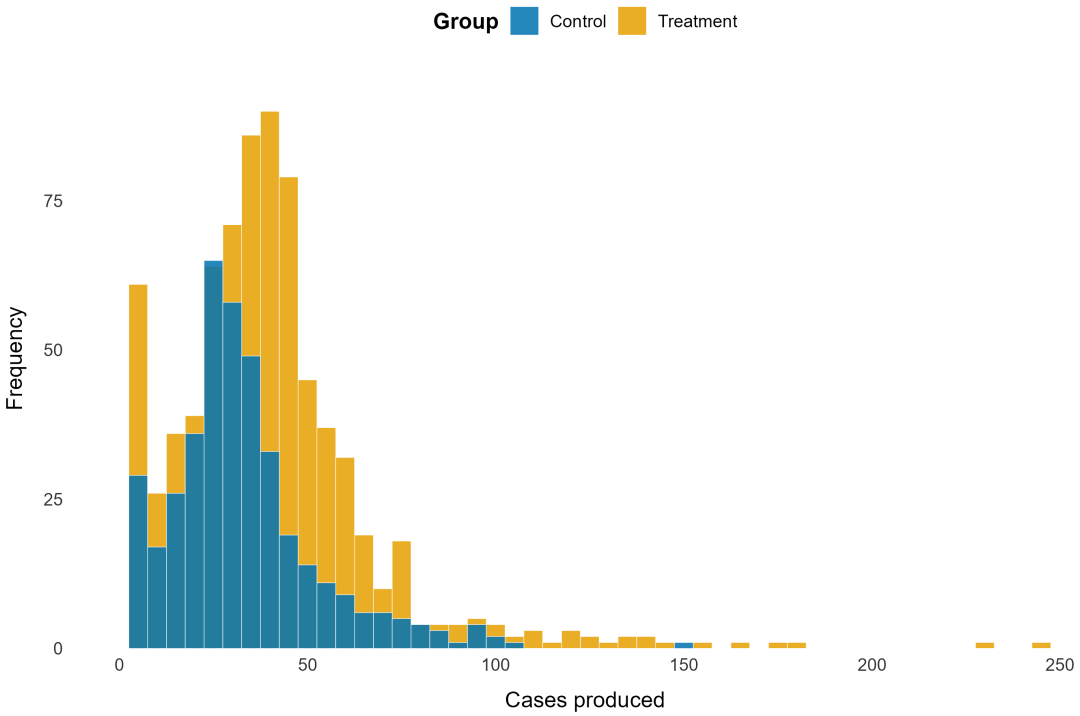

Source: Authors' own elaboration based on TJCE data.

A first visual inspection suggests that the treatment group's distribution has a more pronounced right tail than the control group's. Specifically, days on which a clerk produces more than 100 case files are notably more frequent among those who had access to the MIDAS system (24 instances vs. 2 instances). This is consistent with the expected mechanism: batch processing of similar rulings allows clerks to increase their productivity in terms of the number of case files they can prepare on any given day.

To formally assess whether there is a significant difference between the means of the two groups, the following equation was estimated:

$$y_{it} = \alpha + \beta\, Group_i + u_i + \varepsilon_{it},$$

where y_{it} is the number of case files produced by clerk i on day t, $Group_i$ is a binary variable that takes the value 1 if the clerk belongs to the treatment group, u_i is a random effect at the clerk level capturing unobserved individual heterogeneity, and ε_{it} is the error term. The coefficient of interest is β, which represents the average difference in daily productivity between treatment-group and control-group clerks, conditional on the clerk-level random effects. This model is equivalent to a one-way ANOVA with random effects nested at the clerk level. Table 4 presents the results of this exercise:

Table 4: Model estimation with clerk-level random effects

	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Group$_i$	9.89**	10.5**	8.32**	9.45**	7.86**	9.23***
	(4.69)	(4.82)	(4.02)	(3.87)	(3.04)	(2.83)
Clerk-level RE	Yes	Yes	Yes	Yes	Yes	Yes
Controls	No	Yes	No	Yes	No	Yes
# of clusters	62	62	62	62	56	56
Observations	1238	1238	1212	1212	1087	1087

Notes. Columns (1) and (2) use the full sample; columns (3) and (4) drop clerk-day observations in which more than 100 case files were produced, while columns (5) and (6) drop all observations from clerks who, on any given day, produced more than 100 case files. Controls include the proportion of urgent rulings and dummies for judicial act type. Standard errors (in parentheses) are robust and clustered at the clerk level, adjusted using the bias-reduced linearization proposed by Bell and McCaffrey (2002) and subsequently developed by Pustejovsky and Tipton (2018). * $p < 0.10$, ** $p < 0.05$, *and* *** $p < 0.01$, *respectively.*

The results in Table 4 suggest that access to the MIDAS system had a positive and statistically significant effect on clerk productivity. According to the most parsimonious specification (column (1)), treatment-group clerks produced approximately 10 more case files per day relative to control-group clerks ($F = 4.08$, $p = 0.048$), representing a 37.6% increase relative to the control group mean. Moreover, when controlling for workload composition, such as the proportion of urgent rulings or the distribution of judicial act types at the clerk level, the estimated effect remains stable in magnitude and continues to be statistically significant at the 5% level, suggesting that the previously identified imbalance in the proportion of urgent rulings does not substantively affect the main result.

Columns (3) and (4), in turn, assess the sensitivity of the results to the exclusion of clerk-day observations in which a clerk produced more than 100 case files. While the estimated effect is slightly reduced under this restriction, ranging between 8 and 9.5 additional case files per day, it remains statistically significant and of a magnitude comparable to that obtained under the full sample.

Finally, columns (5) and (6) adopt an even more restrictive criterion, completely removing the clerk-day observations of those clerks who, on any working day, recorded output above 100 daily case files (4 in the treatment group and 2 in the control group). Reassuringly, even under this restricted sample, the effect of the intervention remains positive and significant, with estimates between 7.9 and 9.2 additional case files per day. In fact, the precision of the estimators increases considerably. This latter finding suggests that those clerks exhibited substantial individual variance that was inflating the residual/between-cluster variance for the full sample.

Taken together, these results indicate that the estimated effect was not being driven by a handful of outliers or by exceptionally productive clerks, but rather reflects a robust pattern present in the sample. To gauge the practical implications of this result, assuming an 8-hour workday devoted exclusively to case file preparation, the average control-group clerk produced 30 case files, equivalent to approximately 16 minutes per case file. The increase of 10 additional case files observed in the treatment group implies that, thanks to MIDAS, that time is reduced to 12 minutes per case file, a gain of 4 minutes that, accumulated over the workday, allows clerks to produce one-third more case files in the same amount of time.

Figure 5. Mean number of case files produced daily by clerks in each group.

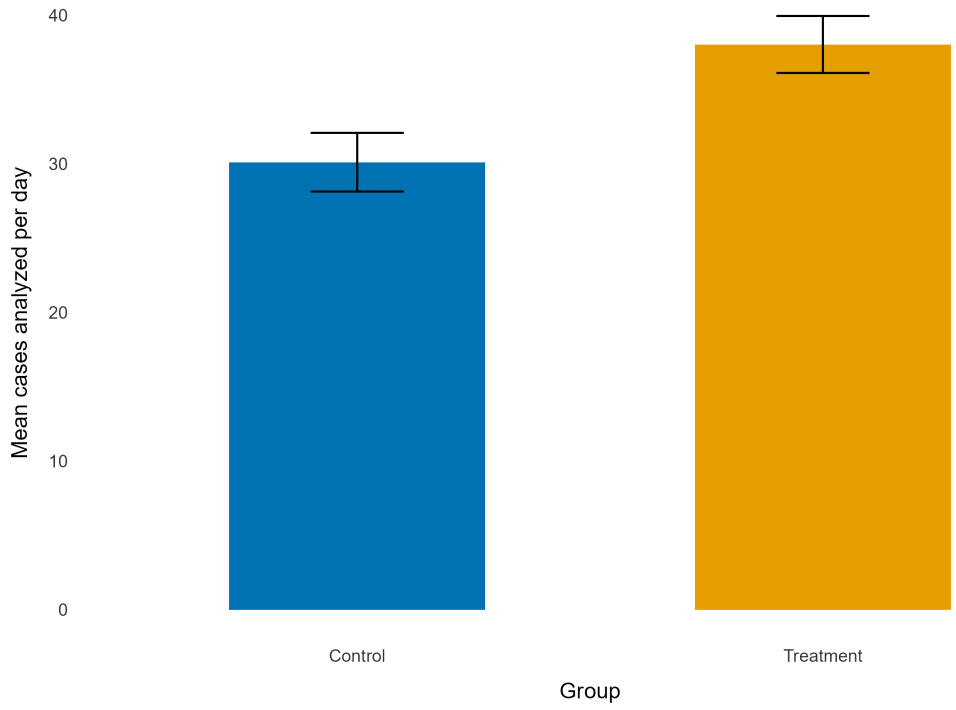

Source: Authors' own elaboration based on TJCE data.

Figure 5 presents the estimated marginal means for each group, with 95% confidence intervals. The visual difference between the groups is clear and consistent with the model's results.

Finally, in order to investigate whether the difference in productivity between clerks of the two groups is stable throughout the duration of the pilot, Figure 6 shows the average daily output (estimated from the marginal means of the mixed-effects model) of both groups by week of the study.

Figure 6. Average daily productivity of clerks in each group by study week.

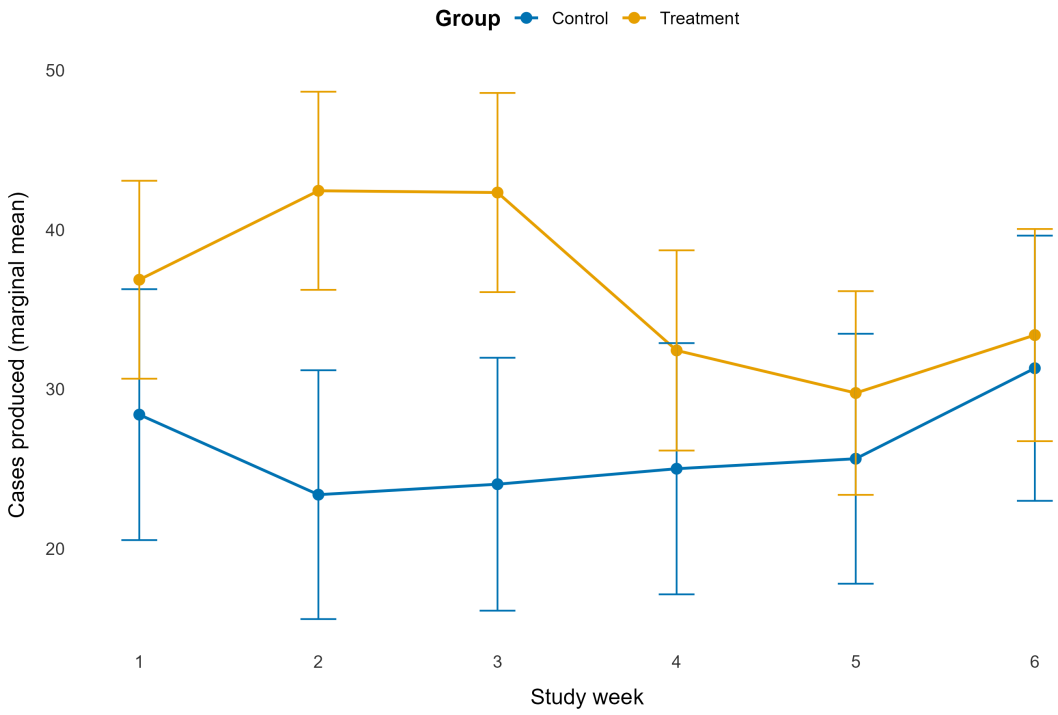

Source: Authors' own elaboration based on TJCE data.

As can be seen, treatment-group clerks exhibit higher productivity throughout all weeks of the study. However, the gap between groups is not constant: it is more pronounced in weeks 2 and 3, narrows in weeks 4 and 5, and partially recovers in week 6. In any case, the confidence intervals of both groups overlap only marginally in the weeks with the narrowest gap, and the average effect over the entire period remains statistically significant. This pattern is consistent with the interpretation that MIDAS generates persistent productivity gains throughout the pilot period.

It is important to note that the results of this study primarily capture gains in terms of the speed of case file production. The impact of the system on the quality of these case files cannot be evaluated with the available data and constitutes a natural extension for future research.

6 Conclusion

The pilot evaluation of the MIDAS system at the First-Degree SEJUD provides insight into the potential benefits of automation in supporting the delivery of justice, with indications of productivity gains in decision-enforcement workflows and greater standardization of case files. The results suggest that clerks who had access to the system processed on average 10 more case files per day than clerks who worked under the traditional workflow, a statistically significant difference. Furthermore, the tool appears to expand the right tail of the productivity distribution: days on which a clerk analyzes more than 100 case files are notably more frequent in the treatment group.

However, these results should be interpreted with caution for three reasons. First, the study is focused on the district of Fortaleza and, in particular, on First-Degree SEJUD clerks, so it is not yet clear to what extent the findings can be generalized to other contexts with similar circumstances. Second, this study focuses solely on the SEJUD's internal indicators and on how MIDAS has influenced clerks' analysis time and work efficiency, but it cannot confirm or rule out that the tool has also had an effect on the quality of these analyses. Third, the treatment group included both clerks with full adherence to the system and clerks with partial adherence, which implies that the estimated effects constitute, at best, a lower bound on the tool's true impact.

Nonetheless, evidence was found that the MIDAS system contributes to greater clerk productivity when preparing case files for the enforcement of rulings, by enabling the joint processing of similar cases and reducing rework and analysis time. These findings prompted the TJCE's decision to scale up the intervention through a statewide RCT that will include, in addition to MIDAS, a tool designed to directly assist judges in drafting their decisions. This new experiment will make it possible to address the main limitations of the present study and to produce robust, generalizable causal evidence on productivity gains and reductions in processing times.

In turn, this new experiment will make it possible to address other questions that the present study left open. First, it remains to be determined whether the productivity gains actually translate into a reduction in the *backlog* or whether they simply shift the bottleneck to another stage — for example, to judges, who would receive a higher volume of case files per day. Second, as already mentioned, the literature has warned of the existence of a *trade-off* between efficiency and quality, so future research should examine whether the increases in productivity come at

the expense of lower quality in the case files prepared by clerks. Third, it is possible that part of the time freed up by MIDAS does not translate into greater output but rather into a reduction in clerks' subjective workload. If so, the true impact on the *backlog* would be smaller than our estimates suggest. Distinguishing between these two effects would also require data on time use. Ultimately, answering these questions will be key to determining the extent to which tools like MIDAS can effectively help reduce the *backlog* and thereby improve access to justice.

References

Aidid, A. and Alarie, B. (2023). *The Legal Singularity: How Artificial Intelligence Can Make Law Radically Better*. University of Toronto Press.

Alarie, B., Niblett, A., and Yoon, A. H. (2018). How artificial intelligence will affect the practice of law. *University of Toronto Law Journal*, 68(supplement 1):106–124.

Amirapu, A. (2021). Justice delayed is growth denied: The effect of slow courts on relationship-specific industries in india. *Economic Development and Cultural Change*, 70(1):415–451.

Bell, R. M. and McCaffrey, D. F. (2002). Bias reduction in standard errors for linear regression with multi-stage samples. *Survey Methodology*, 28(2):169–181.

Borgesano, F., De Maio, A., Laghi, P., and Musmanno, R. (2025). Artificial intelligence and justice: a systematic literature review and future research perspectives on justice 5.0. *European Journal of Innovation Management*, 28(11):349–385.

Casey, A. and Niblett, A. (2019). A framework for the new personalization of law. *University of Chicago Law Review*, 86(2).

Chakraborty, P. (2016). Judicial quality and regional firm performance: The case of indian states. *Journal of Comparative Economics*, 44(4):902–918.

Chemin, M. (2009). The impact of the judiciary on entrepreneurship: Evaluation of pakistan's "access to justice programme". *Journal of Public Economics*, 93(1-2):114–125.

Chemin, M. (2012). Does court speed shape economic activity? evidence from a court reform in india. *The Journal of Law, Economics, & Organization*, 28(3):460–485.

Chen, B. M., Stremitzer, A., and Tobia, K. (2022). Having your day in robot court. *Harvard Journal of Law and Technology*, 36(1).

da Silva, F. C. C. and de Oliveira, A. F. (2025). Artificial intelligence and automation in the brazilian judicial system: analysis of emerging technology implementation and impacts on procedural efficiency. *Brazilian Journal of Information Science: research trends*, 19:e025039–e025039.

de Oliveira, R. S. and Nascimento, E. G. S. (2021). Clustering by similarity of brazilian legal documents using natural language processing approaches. In *Data clustering*. IntechOpen.

de Sousa, W. G., Fidelis, R. A., de Souza Bermejo, P. H., da Silva Gonçalo, A. G., and de Souza Melo, B. (2022). Artificial intelligence and speedy trial in the judiciary: Myth, reality or need? a case study in the brazilian supreme court. *Government Information Quarterly*, 39(1):1016–1060.

Djankov, S., Hart, O., McLiesh, C., and Shleifer, A. (2008). Debt enforcement around the world. *Journal of Political Economy*, 116(6):1105–1149.

Djankov, S., Melcarne, A., Ramello, G. B., et al. (2025). Timely justice as a determinant of economic growth. *Journal of Economic Behavior & Organization*, 238:107235.

OECD (2025). Governing with artificial intelligence: The state of play and way forward in core government functions. Technical report, OECD Publishing.

Oliveira, R. S. d. and Sperandio Nascimento, E. G. (2025). Analysing similarities between legal court documents using natural language processing approaches based on transformers. *PloS one*, 20(4):e0320244.

Ponticelli, J. and Alencar, L. S. (2016). Court enforcement, bank loans, and firm investment: Evidence from a bankruptcy reform in brazil. *Quarterly Journal of Economics*, 131(3):1365–1413.

Pustejovsky, J. E. and Tipton, E. (2018). Small sample methods for cluster-robust variance estimation and hypothesis testing in fixed effects models. *Journal of Business and Economic Statistics*, 36(4):672–683.

Razmetaeva, Y. and Razmetaev, S. (2021). Justice in the digital age: Technological solutions, hidden threats and enticing opportunities. *Access to Just. E. Eur.*, page 104.

Shah, A. and Levy, J. (2026). Access to justice in the age of ai: Evidence from u.s. federal courts. SSRN Working Paper.

Solovey, E., Flanagan, B., and Chen, D. (2025). Interacting with ai at work: Perceptions and opportunities from the uk judiciary. In *Proceedings of the 4th Annual Symposium on Human-Computer Interaction for Work*, pages 1–8.

Visaria, S. (2009). Legal reform and loan repayment: The microeconomic impact of debt recovery tribunals in india. *American Economic Journal: Applied Economics*, 1(3):59–81.

Volokh, E. (2019). Chief justice robots. *Duke Law Journal*, 68(6):1135–1192.

World Justice Project (2019). Measuring the justice gap. Technical report, World Justice Project.