Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Opening the Black Box of Eviction Court with Document-Conditioned LLM Analysis

Will von Geldern

wvg1@law.stanford.edu

Stanford University

Curtis Atkisson

University of Washington

Cheng Ren

University at Albany, State University of New York

Tim Thomas

University of California, Berkeley

Research Article

Keywords:

Posted Date: August 13th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10447706/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Opening the Black Box of Eviction Court with Document-Conditioned LLM Analysis

1. Introduction

Documents filed in civil court proceedings represent a rich source of data about litigants' behavior, procedural events, and case outcomes that are typically not measured by court administrators. The large repositories of documents held by local court systems contain critical information about case dynamics that is of interest to researchers, policymakers, and administrators, but existing research has yet to transform these plentiful repositories into structured datasets for use in empirical legal research.

This paper focuses on analyzing documents filed in eviction proceedings – an increasingly common case type that makes up more than 20 percent of civil cases filed in state and local courts (Chiappetta and White 2026). More than 7.6 million people faced evictions in a typical year prior to the COVID-19 pandemic, and the volume of filings has grown in many jurisdictions since then (Graetz et al. 2023; The Eviction Lab 2025). While these case filings generate enormous volumes of documentation, however, most datasets only measure the number of cases filed and the number of eviction judgments against tenants (e.g., Desmond et al. 2018). Efforts to better understand eviction proceedings have been hindered by the volume and heterogeneity of the documents filed in these cases, as well as the challenges of accessing them at scale.

While existing research has been unable to comprehensively analyze eviction case proceedings at their true volume, small-scale socio-legal studies have unearthed important information about evictions through manual review of case documents (Sudeall and Pasciuti 2021; von Geldern 2025; Summers and Steil 2024). The rapid development of large language models (LLMs) capable of extracting structured information from unstructured text offers a path

toward analyzing documents at a larger scale. Despite this possibility, however, the accuracy of LLM-based pipelines that analyze text extracted from civil court documents remain poorly understood.

This paper presents and validates a document-conditioned LLM analysis pipeline for measuring procedural and legal outcomes in eviction proceedings and applies it to produce novel empirical evidence on renters' participation in Washington State's groundbreaking tenant right to counsel program. The pipeline conditions LLM responses on the full OCR-extracted text of individual court documents labeled with html metadata, constrains outputs to structured JSON schemas tailored to each document type, and algorithmically collapses document-level extractions into case-level measurements. We apply this pipeline to 195,050 documents from 8,503 eviction cases filed in Pierce County, Washington between 2022 and 2024 and produce measurements of 12 procedural and legal variables including tenants' written responses, hearing attendance, legal representation, and court-ordered displacement with 98.4% overall accuracy against a manually coded validation sample of 300 cases.

The methodological contribution of this work is to assess the document-conditioned analytic pipeline. We find that this document-conditioned approach produces highly accurate estimates of case-level outcomes when compared with labor intensive hand coding of documents. Substantively, this methodology allows us to depict eviction case proceedings at a scale that has previously not been achievable. We examine the associations between procedural dynamics, tenant behavior, access to justice, and case outcomes over the course of several years during the COVID-19 pandemic and shortly after the passage of the nation's first statewide tenant right to counsel law. Findings show that the existence of a right to counsel program does not guarantee access to legal assistance for tenants, and that the associations between

representation and legal outcomes can vary substantially in different policy environments. These findings serve to advance the growing body of empirical research on civil court, evictions, and tenant right to counsel programs.

2. Background

2.1 *Computational analysis of legal documents*

AI tools have created new opportunities for analysis of documents filed in civil court proceedings that are often produced by a combination of attorneys, judges, court staff, and pro se litigants at multiple stages of the legal process. Even within a single case, documents may be highly heterogeneous and contain handwritten notes from multiple court actors, tables and other images, or typewritten text. Court documents also frequently degrade during scanning and may contain stamps or other markings from when documents are filed.

Although the possibilities for analyzing these documents at scale remain largely unexplored, some researchers have attempted to overcome this challenge. Recent research from Ren (2025) demonstrates that metadata extracted from case webpages can help identify critical document types for analysis in high-volume civil courts. Text-based LLM analysis has been shown to produce highly accurate – if imperfect – estimates in legal contexts (Breton et al. 2025; Siino et al. 2025), and specialized LLMs pre-trained on corpora of legal documents have further improved performance on domain-specific tasks (Singh et al. 2025). Another recent study explored the possibilities for systematic analysis of images in legal documents, although this approach may be less effective in document sets that are filed by diverse groups of court actors, attorneys, and pro se litigants (Westermann and Savelka 2024).

While most applications of LLM analysis to legal texts have focused on tasks performed by lawyers such as contract analysis, LLMs can also enable measurement of procedural and legal events in case proceedings. A recent study demonstrated that LLMs can extract information from employment tribunal judgments with high accuracy (Ribeiro de Faria et al. 2025), while another evaluated LLM performance on transparency tasks including jury selection audits and eviction proceedings (Shastri et al. 2024). These works highlight both the promise and challenges of using LLMs on court documents, including the need for careful validation and awareness of potential errors.

2.2 *Potential applications to eviction cases*

Given the high volume of eviction cases filed in civil courts across the U.S. each year, LLM analysis could help researchers open the "black box" of eviction court proceedings (Sudeall and Pasciuti 2021). Existing evidence on court procedures and case outcomes largely comes from studies that rely on hand-coding of case documents or limited administrative data outputs (Cassidy and Currie 2023; von Geldern 2025; Summers and Steil 2024). Reliance on manual coding has produced valuable but limited evidence from small case samples, while large-scale administrative data typically only contains limited details about each case. The scale and format of eviction court records – millions of cases filed annually with hundreds of millions of associated documents – create an opportunity for LLM-based approaches that can operate at scale while providing substantial case-level detail and maintaining accuracy.

One recent study tested this proposition using dual case studies focused on measuring details of jury selection and eviction procedures (Shastri et al. 2024). The authors evaluated

GPT-4 Turbo on a range of extraction tasks from Massachusetts eviction case files including case dispositions, settlement types, and landlord representation status. Performance varied substantially across tasks: simple categorical extractions achieved 95.8% accuracy, while the LLM was less accurate at measuring more complex tasks requiring synthesis across multiple documents or analysis of legal text, such as orders for execution (formal eviction judgments). These findings highlight both the potential and the limitations of LLM-based document analysis: this approach can dramatically reduce the labor required to process court documents at scale, but only when coupled with careful validation and appropriate skepticism about measurement inaccuracies.

For the purposes of research on eviction proceedings, LLM analysis could enable measurement of procedural and legal outcomes at a scale that was previously infeasible. Doing so, however, would require explicit validation frameworks that assess both document-level and case-level extraction accuracy. This study demonstrates the possibilities of LLM analysis for detailed measurement of procedural and case outcomes through an in-depth study of three years' worth of eviction cases filed in a high-volume urban court in Pierce County, Washington.

3. Data and methods

3.1 *Analytic pipeline*

We collected and analyzed 195,050 case documents filed in 8,503 unlawful detainer (eviction) cases in Pierce County, Washington between January 2022 and December 2024. Pierce County is the second most populous county in Washington state and has the state's second highest county-level eviction rate (Thomas 2026). These cases represent all residential and

commercial eviction cases filed in Pierce County during those years. Using a comprehensive list of case numbers provided by the Washington State Office of Civil Legal Aid, we accessed documents filed in each case through the Pierce County Superior Court's Legal Information Network Exchange (LINX) web portal. We then labeled each document using the site's HTML metadata, with underscores used to replace spaces and punctuation marks (i.e., a document labeled as Clerk's Minute Entry became clerk_s_minute_entry.pdf).

We used a pre-built optical character recognition (OCR) model available through the Azure AI portal to extract all readable text from the documents. Using this text, we created a dataset containing case numbers, document names, and text strings containing all OCR outputs from each document. Following the example of prior research that relied on hand-coding of case documents (e.g., von Geldern and Martin 2025), we identified seven document types capturing the procedural and legal trajectory of cases from filing through resolution: summonses, complaints, written responses, hearing minutes, settlement or agreement documents, dismissals, and stay or vacate orders. Each document type contains distinct information relevant to measuring tenant procedural engagement, legal representation, and case outcomes.

We developed a two-stage approach to identify case-level outcomes. First, we identified documents filed in each case using each document's metadata label. We then supplemented this document identification with LLM analysis of OCR outputs. We identified the case-level variables to extract from each document type through manual review of available information in those documents. Table 1 summarizes the relationship between variables and source documents. For each selected document, we prompted GPT-4o using document-conditioned prompts designed to extract information present in each document type. Document texts were truncated at 200,000 characters when necessary, which affected fewer than 50 complaint documents and no

other documents. We constrained LLM outputs to structured JSON objects aligned with predefined variables specific to each document type. This constraint ensured that extraction outputs could be directly integrated into downstream statistical analyses without manual parsing or post-processing.

Table 1: Source documents and measurement of case-level variables

Variable	**Source documents**
Address	Summonses, Complaints
Case type (residential or commercial)	Complaints
Alleged amount owed	Complaints
Landlord had representation	Agreements, Complaints, Dismissals, Summonses
Tenant submitted written response	Responses, Notice of Appearances
Hearing held	Hearing minutes
Tenant attended hearing	Hearing minutes
Tenant received representation	Orders appointing counsel, Agreements, Orders and Motions for Dismissal
Eviction judgment issued	Writ of Restitution, Orders staying or vacating Writ of Restitution
Case dismissed	Orders and Motions for Dismissal
Record protection order	Order for Limited Dissemination
Tenant displaced	Agreement, Order for Writ of Restitution, Order staying or vacating Writ of Restitution, and Orders and Motions for Dismissal

After extraction, we removed commercial cases (88 in total) from the dataset. We then cleaned and geocoded address data extracted from summons and complaint documents. We corrected OCR errors in city names through visual inspection of extracted data (e.g., Tacuman to Tacoma, Purellup to Puyallup) and aligned addresses from both documents to identify a valid address. Nearly all (98%) of the cases in our dataset had addresses in both summons and complaint documents; 145 cases had a complaint address only, 9 had a summons address only, and 4 had no legible address in either document. For cases with only one document type filed, we used that address for geocoding.

We geocoded addresses in two stages. First, we submitted all addresses with complete street, city, state, and ZIP code information to the U.S. Census Bureau's batch geocoding API, which matched 92.4% of cases. We then submitted the remaining unmatched addresses (predominantly cases with missing ZIP codes) to the Google Maps Geocoding API, which recovered an additional 6.3% of cases, bringing total geocoding coverage to 98.7%. We performed a spatial join between geocoded coordinates and 2020 Census block group boundaries to assign each case a block group identifier for merging with American Community Survey estimates of neighborhood median income. We verified that all geocoded addresses fell within Pierce County boundaries, and excluded all cases for which no extracted addresses geocoded to a valid location within the county.

Through this process, we removed 46 residential cases with missing address data or addresses geocoded outside of Pierce County. Using geocoded locations, we estimated travel times from each tenant's address to the Pierce County Superior Courthouse via both driving and public transit using the Google Maps Distance Matrix API, assuming a 1:00pm arrival to align with typical unlawful detainer docket timing. Public transit estimates were unavailable for 652

cases and these cases were omitted from relevant analyses, but driving time estimates were successfully calculated for all cases in the final dataset. The final dataset for analysis comprised 8,380 residential cases with complete information on geographic, procedural, and legal variables.

3.2 *Validation*

To assess the accuracy of our case-level estimates, we randomly selected 300 cases from the full dataset and manually coded all outcomes for these cases. Each case required about 5 minutes of manual coding by research assistants, for an approximate total of 25 hours spent on manual review. We sampled 100 cases from each of the three filing years (2022-2024) to ensure temporal diversity and capture variations in document formats and court procedures as they evolved throughout the study period.

4. Results

4.1 *Validation*

To verify the accuracy of our analytic pipeline, we calculated accuracy, precision, recall, and F1 scores for all boolean variables by comparing LLM estimates to manually coded results for the 300 randomly sampled cases. For the continuous variable measuring alleged rental arrears, we also estimated mean absolute error (MAE), mean absolute percentage error (MAPE), and the proportion of estimates within $100 of the true value – an amount that was intended to reflect a relatively minor error level in the context of monthly rent and rental arrears.

Table 2: Accuracy metrics for LLM-based extraction of case-level variables (n = 300)

Variable	**Accuracy**	**Precision**	**Recall**	**F1**
Landlord had representation	99.7%	99.7%	100.0%	0.998
Tenant submitted written response	98.3%	100.0%	95.5%	0.977
Hearing was held	99.7%	100.0%	99.4%	0.997
Tenant attended a hearing	98.3%	98.5%	97.7%	0.981
Eviction judgment	99.0%	100.0%	98.4%	0.992
Case dismissed	98.3%	99.0%	96.1%	0.975
Record protected	99.0%	98.7%	97.5%	0.981
Tenant received representation	99.0%	97.8%	100.0%	0.989
Tenant displaced	94.0%	97.8%	94.4%	0.961
Total	98.4%	99.1%	97.7%	0.983

Overall case-level accuracy was 98.4%, with equally strong precision (99.1%), recall (97.7%), and F1 score (0.983), demonstrating consistently high performance across all variable types. For the continuous arrears variable, approximately 97% of estimated values fell within $100 of actual amounts, with mean absolute error of $50.50 and mean absolute percentage error of 0.89%.

Individual variable accuracy ranged from 94.0% to 99.7%. The most imprecise estimates were those of court-ordered displacement, which required documenting multiple outcomes across different documents in proper chronological order. Despite this complexity, accuracy remained above 94% for that outcome.

4.3 *Descriptive statistics*

Table 3 presents descriptive statistics for case characteristics of the full sample of 8,234 valid residential eviction cases filed in Pierce County between 2022 and 2024. Case filings increased substantially over the study period as the federal eviction moratorium and Washington state's additional protections lapsed (Demkovich 2024). The vast majority of cases involved landlords with legal representation (98.9%) and nearly 90% of cases included a monetary complaint against the tenant. Monetary complaint amounts demonstrated a significant right skew, with an average of $7,275 and a median of $4,395. Monetary complaint statistics exclude cases with no monetary basis (11.9%). Just over a quarter of cases were filed in neighborhoods below 50% of Pierce County's area median income (AMI).

Table 3: Case characteristics (n = 8,234)

Filing year	
2022	23.8%
2023	34.2%
2024	42.0%
Case characteristics at filing	
Contains non-monetary complaint	11.9%
Median monetary complaint ($)	$4,395
Mean monetary complaint (SD)	$7,275 ($13,997)
Landlord represented	98.9%
Filed in low-income neighborhood (below 50% Pierce County AMI)	29.3%
Estimated travel time to courthouse (mins)	
Median driving time	14.9
Mean driving time (SD)	16.4 (8.7)
Median public transit time	48.7

Mean public transit time (SD)	55.0 (32.6)

We then disaggregate case outcomes by whether the property that is the subject of the complaint is in Tacoma because of a local ordinance prohibiting default judgments against tenants without a hearing that could affect tenant or landlord decision-making during the eviction process (Tacoma Municipal Code 2018). Figure 1 demonstrates the stage of the legal process at which this ordinance disallows eviction judgments on the basis of tenants' non-participation.

Figure 1: The eviction process, right to counsel program, and intended effect of Tacoma Municipal Ordinance 1.95.090

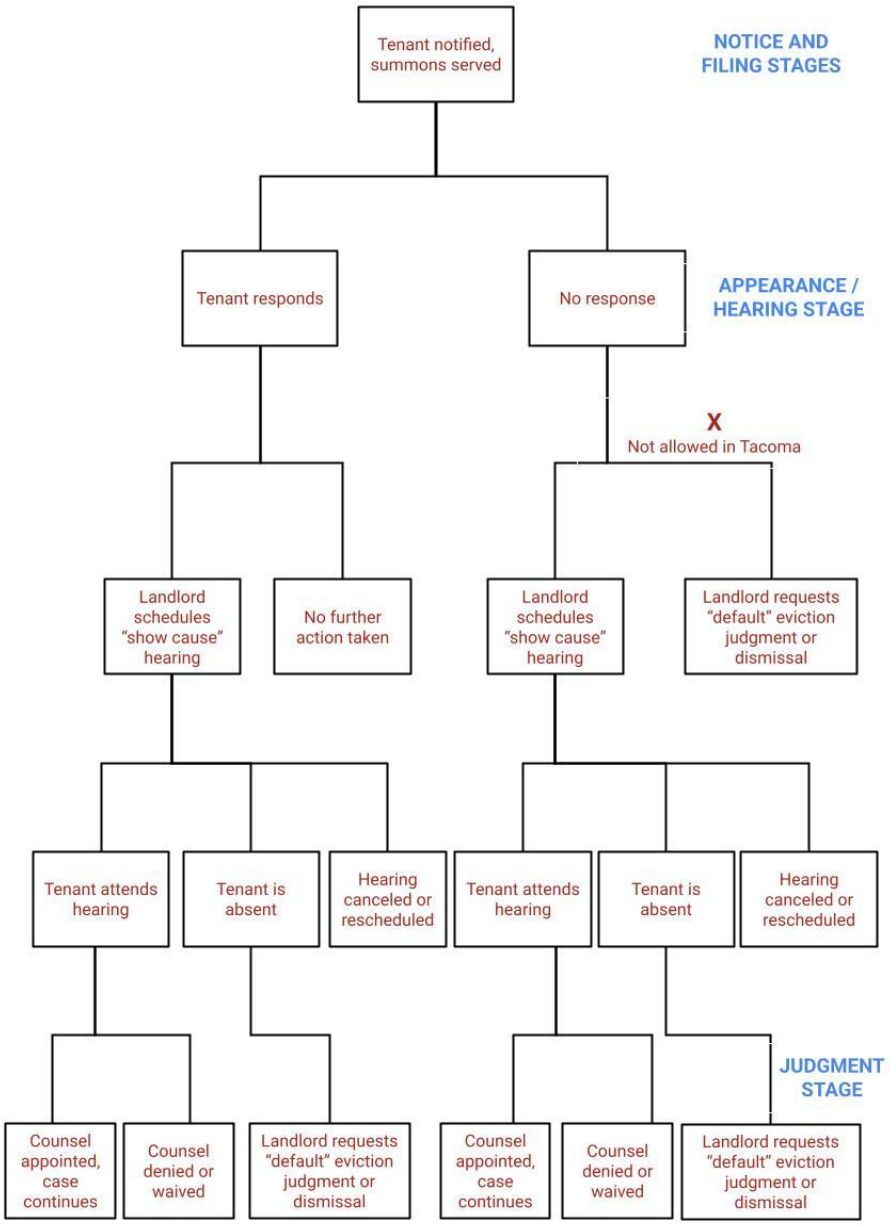

Table 4 shows the results of disaggregating case outcomes by location. Tenants submitted written responses in 39.7% of cases overall, with lower rates in Tacoma than elsewhere. Despite lower response rates, Tacoma cases were more likely to include a scheduled hearing in alignment with the intended effect of the local Tacoma ordinance on default judgments. Although case-level hearing attendance rates were slightly higher in Tacoma because hearings were more likely to be scheduled in those cases, tenants in Tacoma were actually less likely to attend hearings in

cases where one was scheduled. Representation rates were slightly higher in Tacoma, and case outcomes (eviction judgments, dismissals, record protection orders, and court displacement) were slightly more favorable in Tacoma when compared to other parts of Pierce County.

Table 4: Case-level outcomes (n = 8,234)

	Address outside of Tacoma (n = 4,029)	Address in Tacoma (n = 4,205)	Total (n = 8,234)
Submitted written response	42.2%	37.3%	39.7%
Hearing held	55.2%	65.0%	60.2%
Tenant attended hearing	43.9%	47.9%	45.9%
Tenant attended hearing (conditional on hearing held)	79.5%	73.7%	76.3%
Received representation	44.0%	47.6%	45.8%
Eviction judgment	67.8%	64.7%	66.2%
Dismissal	28.6%	29.8%	29.2%
Record protection order	27.4%	29.3%	28.4%
Court displacement	76.2%	73.4%	74.8%

4.4 *Logistic regression models*

To examine predictors of tenants' participation in the legal process, access to legal representation, and case outcomes, we estimated a series of year-stratified logistic regression models. Stratifying by year accounts for the rapidly changing policy and economic environment throughout the study period (2022–2024), allowing examination of how the relationships between participation, access to legal assistance, and case outcomes evolved throughout the course of this volatile period.

The first set of models examines case-level predictors of tenant behavior. We first estimate a model predicting tenants' submission of a written response (Table 5) and then two subsequent models predicting whether a hearing occurred and whether the defendant or their attorney attended a hearing at any point during the case (Table 6). To predict whether tenants filed a written response, we estimated the following logistic regression model:

$$Y_i = \beta_0 + \beta_1(\text{Tacoma})_i + \beta_2(\text{Monetary})_i + \beta_3(\text{Plaintiff Representation})_i + \beta_4(\text{Low-income})_i + \beta_5(\text{Drive Time})_i + \beta_6(\text{Transit Time})_i + \mathcal{E}_i$$

In this equation, i indexes the case and Y_i represents the log odds of a tenant submitting a written response in that case. β_1 through β_6 measure the associations between submission of a response and, respectively, whether the property was located in Tacoma, whether the complaint included a monetary claim, whether the landlord had legal representation, whether the property was in a low-income neighborhood, driving time to the courthouse, and public transit time to the courthouse. All models are estimated separately for each filing year (2022, 2023, 2024) to account for the changing policy environment.

We then estimated a series of year-stratified models for two other procedural outcomes: whether a hearing was held in case, and whether a tenant or their attorney attended a hearing in the case. The first predicts whether a hearing was held using the following equation:

$$Y_i = \beta_0 + \beta_1(\text{Tacoma})_i + \beta_2(\text{Monetary})_i + \beta_3(\text{Plaintiff Representation})_i + \beta_4(\text{Low-income})_i + \beta_5(\text{Response Filed})_i + \mathcal{E}_i$$

In this equation, i indexes the case and Y_i represents the log odds of a hearing being held in that case. β_1 through β_6 measure the associations between submission of a response and, respectively, whether the property was located in Tacoma, whether the complaint included a monetary claim, whether the landlord had legal representation, whether the property was in a low-income neighborhood, and whether response was filed in the case. The second model predicts hearing attendance conditional on a hearing being held:

$$Y_i = \beta_0 + \beta_1(\text{Tacoma})_i + \beta_2(\text{Monetary})_i + \beta_3(\text{Plaintiff Representation})_i + \beta_4(\text{Low-income})_i + \beta_5(\text{Response Filed})_i + \beta_6(\text{Drive Time})_i + \beta_7(\text{Transit Time})_i + \mathcal{E}_i$$

Next, we estimated year-stratified models to assess the associations between case characteristics and procedural outcomes and tenants' access to legal representation. The models used the following equation:

$$Y_i = \beta_0 + \beta_1(\text{Tacoma})_i + \beta_2(\text{Monetary})_i + \beta_3(\text{Plaintiff Representation})_i + \beta_4(\text{Low-income})_i + \beta_5(\text{Response Filed})_i + \beta_6(\text{Attended Hearing})_i + \mathcal{E}_i$$

Here, i indexes the case and Y_i represents the log odds of a tenant successfully accessing legal representation during the case. In addition to predictors that had been included in prior models, we included β_6 to measure the association between hearing attendance and access to legal representation.

Finally, we estimated a series of year-stratified models that examine associations between procedural engagement, legal representation, and case outcomes. All case outcome models use the same specification:

$$Y_i = \beta_0 + \beta_1(\text{Tacoma})_i + \beta_2(\text{Monetary})_i + \beta_3(\text{Plaintiff Representation})_i + \beta_4(\text{Low-income})_i + \beta_5(\text{Response Filed})_i + \beta_6(\text{Attended Hearing})_i + \beta_7(\text{Legal Representation})_i + \mathcal{E}_i$$

The outcome (Y_i) varies across models: whether an eviction judgment was issued, whether the case was dismissed, whether an Order for Limited Dissemination (record protection order) was issued, and whether the case resulted in the tenant's physical displacement. This model specification allows examination of how procedural engagement at the checkpoint stages and access to legal representation relate to final case outcomes, while controlling for case characteristics present at filing.

4.5 *Predictors of procedural outcomes*

There are two requirements for tenants to avoid a default eviction judgment in Pierce County Superior Court: submission of a written response and attendance at an initial hearing. A prior study of Washington's right to counsel program showed that about 40 percent of tenants submit a written response and about 20 percent of cases result in a default eviction judgment against tenants who do not submit one (von Geldern and Martin 2025). This study, however, did not examine statistical predictors of written responses or how these relationships changed over time. Figure 2 shows the variation in tenants' ability to overcome these checkpoints throughout the study period, and subsequent tables present year-stratified logistic regression models predicting tenant responses and hearing attendance.

Figure 2 depicts the percentage of cases in which a response was filed, a hearing was held, and a tenant or their attorney attended at least one hearing for each month throughout the study period. Approximately 40% of tenant households submitted a written response (n = 3,269), although the response rate declined slightly in 2024 (37.9%) relative to prior years. Despite this

decline in responses, hearings were held in a growing share of cases over the study period, rising from 54.2% in 2022 to 65.3% in 2024. Tenant hearing attendance followed a similar upward trend, increasing from 40.6% of all cases in 2022 to 49.6% in 2024.

Figure 2: Eviction filings and procedural outcomes

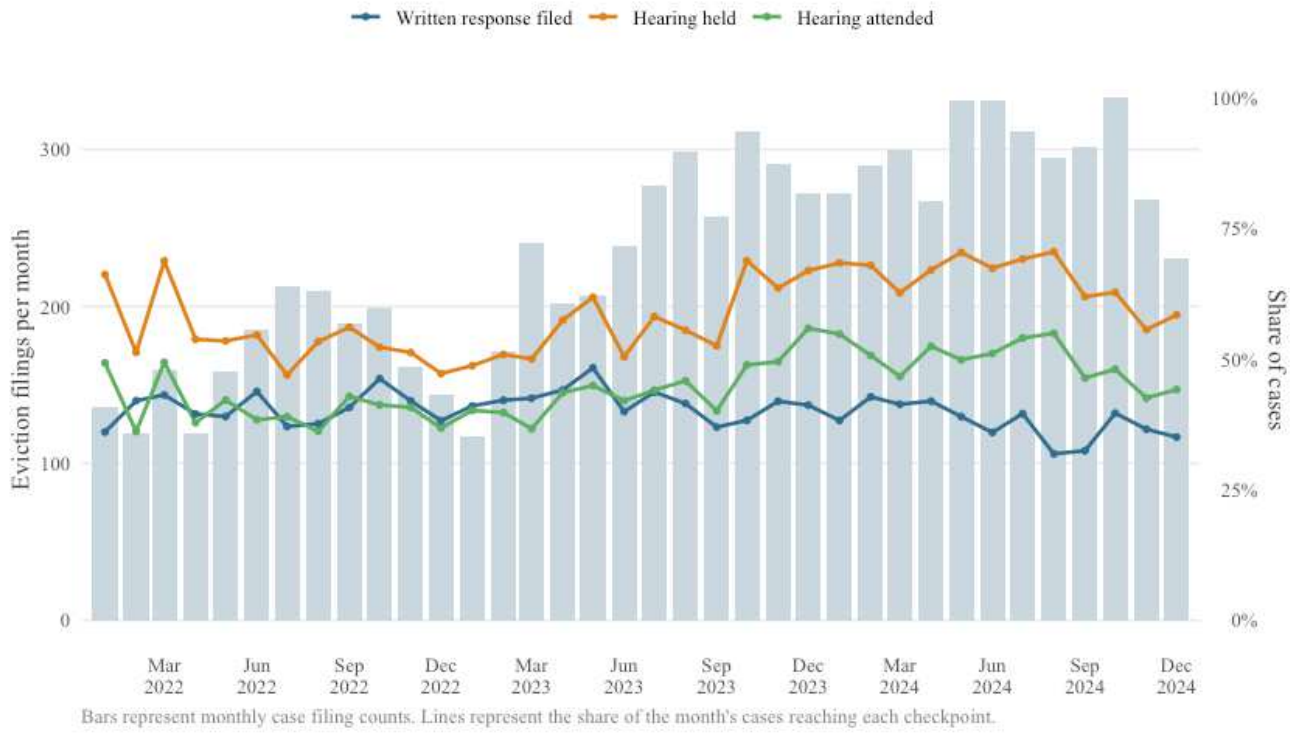

Disaggregation of results by location suggests that changes over time may be related to increased adherence to the Tacoma ordinance preventing default judgments against tenants who do not submit a written response. Figure 3 depicts this trend. In 2022, default eviction judgments without a hearing were issued at nearly identical rates in Tacoma (36.8%) and elsewhere (36.5%). By 2023 a modest gap in these cases had opened (28.5% vs 36.1%), and in 2024 only 10.2% of cases filed in Tacoma resulted in an eviction judgment without a hearing compared to 28.6% elsewhere. October 2023 appears to be a significant juncture for this trend. After this

point, tenants in Tacoma became less likely to submit a response, more likely to have the opportunity to attend a hearing, and less likely to attend hearings when they were held.

Figure 3: Default eviction judgments in cases with no hearing

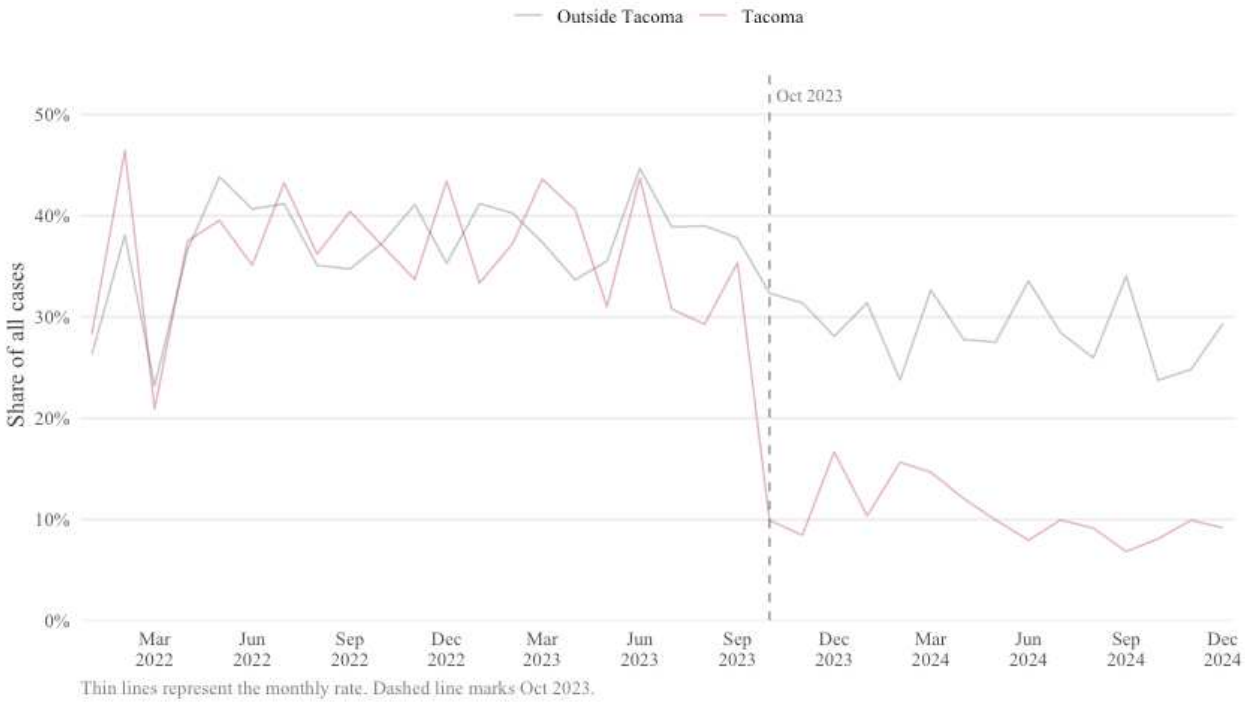

Model outputs also support the possibility of a changing policy environment over time. The Tacoma ordinance showed a weak negative association with response rates in 2022 and then a strong negative association in 2024. Higher alleged rent amounts and non-monetary complaints were associated with lower response rates in 2022 and 2024, while low-income neighborhood status showed a marginally negative association with response in 2022 but not in subsequent years. Monetary complaint status was not significantly associated with response in any year, suggesting that the type of complaint does not independently predict tenant engagement after controlling for other case characteristics.

Table 5: Predictors of tenants' written responses

	2022	2023	2024
Intercept	0.195	-1.645	-0.596
SE	(0.449)	(0.641)	(0.565)
Address in Tacoma	-0.261*	-0.133	-0.267**
SE	(0.128)	(0.105)	(0.094)
Monetary complaint	0.033	0.262	0.174
SE	(0.120)	(0.146)	(0.154)
Landlord represented	-0.178	1.400*	0.229
SE	(0.408)	(0.625)	(0.551)
Low-income neighborhood	-0.320**	-0.003	-0.117
SE	(0.112)	(0.090)	(0.082)
Drive time (mins)	-0.024	0.004	0.011
SE	(0.015)	(0.013)	(0.012)
Transit time (mins)	0.003	-0.001	-0.002
SE	(0.003)	(0.003)	(0.002)
N	1771	2550	3165
AIC	2378.1	3453.8	4172.5
BIC	2416.5	3494.7	4214.9
Log Likelihood	-1182.06	-1719.91	-2079.24

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Results of the models predicting hearing occurrence and attendance show a similar pattern. In 2022, filing in Tacoma was not predictive of whether a hearing was held, by 2023 it was statistically significant, and by 2024 it had grown further. Tenants' written response rates remained the strongest predictor of hearings being held in all three years, while the association between low-income neighborhood status and hearing occurrence fluctuated over time.

Conditional on a hearing being held, the predictors of tenant attendance were relatively stable across years, with responses being consistently the strongest predictor. Filing in Tacoma was not associated with conditional attendance in 2022 and 2023 but became significantly negative in 2024, consistent with the large increase in hearings being scheduled for Tacoma-based tenants who had not submitted a response. Travel time was not a significant predictor of conditional hearing attendance in any year, consistent with the fact that parties may attend hearings virtually in Pierce County and prior research showing that geographic proximity effects dissipate when virtual hearing attendance is allowed.

Table 6: Predictors of hearing outcomes

	2022	2023	2024	2022	2023	2024
	Hearing held			Attendance (conditional)		
Intercept	-0.255	-0.340	-1.228	1.346	0.115	1.206
SE	(0.415)	(0.406)	(0.440)	(0.838)	(0.623)	(1.131)
Address in Tacoma	0.061	0.616***	0.944***	-0.051	-0.148	-0.336*
SE	(0.105)	(0.092)	(0.082)	(0.201)	(0.174)	(0.142)
Monetary complaint	-0.108	-0.112	0.049	0.078	0.013	-0.169
SE	(0.124)	(0.147)	(0.153)	(0.186)	(0.217)	(0.218)
Landlord represented	-0.390	-0.518	0.704	-0.743	0.564	-0.242
SE	(0.412)	(0.405)	(0.440)	(0.780)	(0.578)	(1.118)
Low-income neighborhood	-0.236*	-0.007	0.272**	-0.222	-0.049	-0.164
SE	(0.118)	(0.099)	(0.088)	(0.178)	(0.143)	(0.112)
Drive time (mins)	—	—	—	0.029	0.010	0.017
SE				(0.023)	(0.019)	(0.017)
Transit time (mins)	—	—	—	-0.007	-0.002	-0.004
SE				(0.005)	(0.004)	(0.004)
Submitted response	2.302***	2.470***	2.086***	1.010***	1.212***	0.896***
SE	(0.115)	(0.103)	(0.099)	(0.155)	(0.130)	(0.109)
N	1771	2550	3165	959	1495	2104
AIC	2204.9	3051.8	3785.6	1037.8	1521.0	2256.9
BIC	2238.3	3087.5	3822.4	1076.8	1563.5	2302.1
Log Likelihood	-1096.44	-1519.91	-1886.78	-510.91	-752.50	-1120.43

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.6 *Predictors of access to legal representation*

Throughout the study period, tenants' ability to surmount procedural hurdles was highly predictive of access to legal counsel. Table 7 presents logistic regression models predicting tenants' access to representation. Overall, 77.3% of tenants who submitted a response received representation, compared to 25.1% of those who did not – and 89.8% of tenants who attended a hearing received representation, while only 8.4% of those who did not were able to access legal counsel. The share of tenants who received representation but did not submit a response grew substantially over the study period, from 19.3% in 2022 to 31.2% in 2024. Despite this change, however, logistic regression models show that both written responses and hearing attendance are enduring predictors of access to representation.

Figure 4: Representation rates by status of written responses

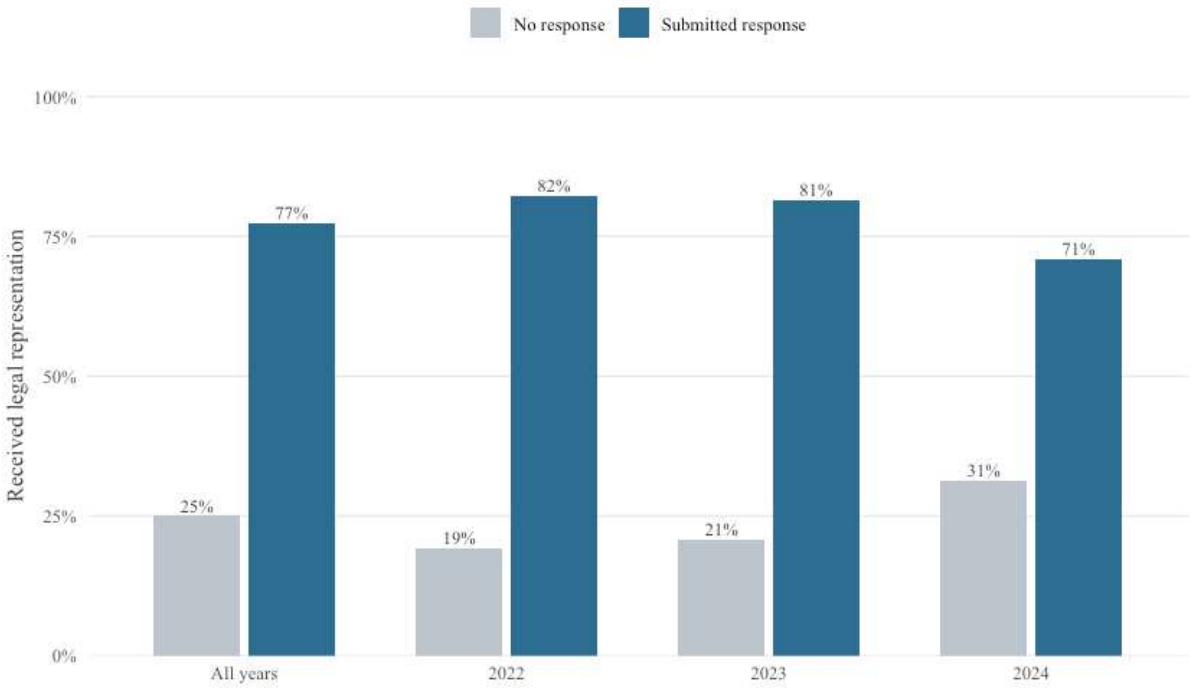

The year-stratified models also support the idea that the pathway to representation has changed over time. While the association between written responses and representation decreases with each year, the association between hearing attendance and representation grows – although all of these behavioral predictors are strongly associated with access to representation throughout the study period. Filing in Tacoma and in low-income neighborhoods was associated with higher rates of representation in 2023 but not in the other two years. Landlord representation appears as a positive predictor of tenant representation in 2024, but this result is difficult to interpret given that only 27 cases in that year (0.8%) involved an unrepresented landlord.

Table 7: Predictors of tenant legal representation

	2022	2023	2024
Intercept	-2.710***	-2.955***	-6.310***
SE	(0.587)	(0.662)	(0.738)
Address in Tacoma	0.005	0.259	0.169
SE	(0.156)	(0.137)	(0.147)
Monetary complaint	0.113	0.002	0.320
SE	(0.183)	(0.231)	(0.266)
Landlord represented	-0.234	-0.262	2.462***
SE	(0.572)	(0.656)	(0.717)
Low-income neighborhood	0.102	0.157	0.070
SE	(0.179)	(0.150)	(0.158)
Submitted response	2.813***	2.434***	0.917***
SE	(0.169)	(0.142)	(0.149)
Attended hearing	3.896***	4.025***	5.574***
SE	(0.173)	(0.140)	(0.167)
N	1958	2819	3457
AIC	1131.2	1566.2	1513.0
BIC	1170.2	1607.8	1556.1
Log Likelihood	-558.58	-776.12	-749.52

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

4.7 *Predictors of case outcomes*

The final set of models in Tables 8-10 predict four case-level outcomes using the same year-stratified approach: eviction judgments, dismissals, record protection orders, and court-ordered displacement. These models include tenants' written responses, hearing attendance, and legal representation as predictors along with other case characteristics at filing. These models allow us to examine the association between procedural outcomes and case outcomes after controlling for case characteristics.

Table 8: Predictors of eviction judgments and dismissals

	2022	2023	2024	2022	2023	2024
	Eviction judgment			Dismissal		
Intercept	-.0348	-0.127	-0.519	0.051	0.124	-0.244
SE	(0.442)	(0.402)	(0.418)	(0.414)	(0.392)	(0.401)
Address in Tacoma	-0.189	-0.052	-0.265***	0.162	-0.002	0.182*
SE	(0.109)	(0.087)	(0.073)	(0.109)	(0.088)	(0.077)
Monetary complaint	-0.220	-0.077	-0.181	0.111	0.040	0.294
SE	(0.130)	(0.145)	(0.145)	(0.129)	(0.147)	(0.158)
Landlord represented	2.172***	1.709***	1.317**	-1.886***	-1.658***	-0.651
SE	(0.441)	(0.402)	(0.419)	(0.411)	(0.391)	(0.402)
Low-income neighborhood	0.251*	-0.160	0.033	-0.178	0.120	-0.050
SE	(0.126)	(0.095)	(0.079)	(0.126)	(0.096)	(0.083)
Submitted response	-0.661***	-0.644***	-0.516***	0.497***	0.836***	0.688***
SE	(0.130)	(0.103)	(0.080)	(0.132)	(0.107)	(0.085)
Attended hearing	0.162	0.186	0.702***	-0.205	-0.308*	-0.738***
SE	(0.155)	(0.136)	(0.159)	(0.155)	(0.137)	(0.163)
Received representation	-1.451***	-1.058***	-0.582***	1.384***	0.654***	0.152
SE	(0.175)	(0.147)	(0.159)	(0.178)	(0.148)	(0.163)
N	1958	2819	3457	1958	2819	3457
AIC	2095.7	3268.7	4477.7	2096.2	3205.5	4149.2
BIC	2140.3	3316.2	4526.9	2140.8	3253.1	4198.4
Log Likelihood	-1039.83	-1626.33	-2230.87	-1040.09	-1594.77	-2066.59

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Tenants' written responses were negatively associated with eviction judgments and displacement and positively associated with dismissals across all three years. Monetary complaint type was not a significant predictor of eviction judgments or dismissals in any year after controlling for other factors. Hearing attendance was not significantly associated with eviction judgments or dismissals in 2022 or 2023, but in 2024 it counterintuitively became positively associated with eviction judgments and negatively associated with dismissals. This finding suggests that the implementation of the Tacoma ordinance led to more tenants attending a hearing, but that tenants who became more likely to attend did not necessarily have strong legal cases.

The association between legal representation and case outcomes changed significantly during the study period. Representation was associated with reduced eviction judgments and displacement, and increased dismissals and record protection orders, but the relationships with eviction judgments, displacement, and dismissals diminished with each successive year. This attenuation coincides with the persistent reduction in availability of COVID-era emergency rental assistance funding from 2022 through 2024 (Demkovich 2024).

Table 9: Predictors of court displacement

	2022	2023	2024
Intercept	0.437	-0.167	-0.587
SE	(0.383)	(0.389)	(0.419)
Address in Tacoma	-0.102	-0.076	-0.365***
SE	(0.111)	(0.095)	(0.079)
Monetary complaint	-0.271*	0.060	-0.170
SE	(0.135)	(0.156)	(0.157)
Landlord represented	1.344***	1.915***	1.583***
SE	(0.376)	(0.389)	(0.420)
Low-income neighborhood	0.297*	-0.149	-0.077
SE	(0.131)	(0.103)	(0.084)
Submitted response	-0.662***	-0.722***	-0.584***
SE	(0.137)	(0.118)	(0.088)
Attended hearing	0.109	0.383**	0.824***
SE	(0.157)	(0.146)	(0.165)
Received representation	-0.830***	-0.395*	-0.048
SE	(0.179)	(0.160)	(0.167)
N	1958	2819	3457
AIC	2019.5	2873.5	4027.7
BIC	2064.1	2921.0	4076.9
Log Likelihood	-1001.74	-1428.73	-2005.84

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Tenant responses and representation were negatively associated with displacement across all three years. Landlord representation and higher alleged arrears were positively associated with displacement throughout the study period. Filing in Tacoma showed a weak negative

association to displacement in 2022 and 2023 and became significantly negative in 2024. Non-monetary complaints were negatively associated with displacement in 2022 but not in other years.

Table 10: Predictors of record protection (Orders for Limited Dissemination)

	2022	2023	2024
Intercept	-3.310***	-4.001***	-4.430***
SE	(0.520)	(0.640)	(0.936)
Address in Tacoma	0.250	-0.016	-0.324**
SE	(0.133)	(0.112)	(0.113)
Monetary complaint	0.021	0.411*	0.183
SE	(0.155)	(0.185)	(0.215)
Landlord represented	-0.741	-0.116	-0.534
SE	(0.473)	(0.612)	(0.919)
Low-income neighborhood	0.006	-0.170	0.153
SE	(0.154)	(0.122)	(0.121)
Submitted response	0.117	0.363**	-0.005
SE	(0.154)	(0.124)	(0.113)
Attended hearing	0.108	0.692***	2.642***
SE	(0.175)	(0.162)	(0.284)
Received representation	3.434***	3.661***	3.479***
SE	(0.281)	(0.245)	(0.253)
N	1958	2819	3457
AIC	1413.8	2021.1	2071.6
BIC	1458.5	2068.7	2120.8
Log Likelihood	-698.91	-1002.57	-1027.79

$* p < 0.05, ** p < 0.01, *** p < 0.001$

Legal representation was by far the strongest and most stable predictor of record protection orders across all three years. The relationship between hearing attendance and record protection became stronger over time, while the relationship between submitting a response and record protection varied. Tacoma residency was not a significant predictor of record protection in 2022 or 2023 but became negatively associated in 2024. Monetary complaint type was positively associated with record protection orders in 2023 but not in other years.

5. Discussion

This study demonstrates that document-conditioned LLM analysis can be used to create highly accurate estimates of procedural and legal outcomes in eviction proceedings. By conditioning LLM inference on OCR-extracted document text, constraining outputs to structured JSON schemas, and algorithmically reconciling document-level extractions across multiple documents into case-level outcomes, we demonstrate that large language models can reliably support measurement of procedural and legal events in administrative records with accuracy comparable to manual coding processes.

5.1 *Methodological contributions*

Our results show that document-conditioned LLM analysis is a highly accurate but imperfect means of capturing case-level outcomes in this context. We evaluated the accuracy of our analytic pipeline against a sample of 300 manually coded cases spanning three years of eviction filings in Pierce County, Washington. Accuracy scores were high across a number of

procedural and legal outcomes, and results suggest that variables requiring inputs based on LLM analysis of multiple documents and chronological ordering were measured less accurately. This finding suggests that tasks combining multiple LLM responses and cross-document inference present the greatest challenges for LLM analysis of text extracted from court documents.

The 0.89% mean absolute percentage error on the continuous arrears variable demonstrates that LLMs can also extract and interpret numerical information from court documents with high fidelity. This finding has implications beyond eviction research: administrative records across different legal and policy domains contain numerical data that may be amenable to similar extraction approaches. This approach does not appear to be reliable, however, for estimation tasks that require perfection at the individual case level or where manual review is not possible.

Nonetheless, we show that the combined use of document identification from html metadata and OCR text verification via an LLM is an effective way to balance computational efficiency with accuracy, avoiding the expense of LLM-based classification for every document filed in court proceedings while ensuring correct estimates of variables from critical document types. The validation framework itself also provides a template for other researchers implementing LLM-based extraction on administrative records or documents.

5.2 *Limitations*

Document-conditioned LLM analysis has several important limitations in this context. Primarily, this study also shows that LLM analysis is unlikely to be sufficient for high-stakes legal analysis and is more suitable for empirical research. While accuracy scores are very high

overall at the case level, this analytic approach may be inappropriate in contexts requiring perfection.

The findings of our analysis are also limited by the fact that we examined a single county over a three-year period during a period of many significant policy reforms including a federal eviction moratorium, implementation of a groundbreaking tenant right to counsel law, and significant changes to local court procedures. Results may not generalize to jurisdictions with different legal frameworks, representation structures, or case volumes. Additionally, the validation sample of 300 cases, while stratified across years, may not adequately capture rare procedural scenarios or unusual document formats. Expanding the scope of similar analyses to multiple jurisdictions and creating larger validation sets would provide stronger evidence on the portability and accuracy of document-conditioned LLM extraction.

5.3 *Implications for access to justice*

Beyond the methodological contribution, this work enables measurement of tenants' participation and case outcomes in eviction proceedings at a previously infeasible scale. The pipeline's accuracy permits researchers to examine the barriers that tenants face in their pursuit of legal representation and favorable case outcomes. It also allows us to observe the extent to which specific procedural barriers – namely submission of a written response and hearing attendance – create meaningful barriers to accessing legal representation when it is available. Understanding and reducing the difficulty of overcoming these "administrative checkpoints" (Kim et al. 2025) may allow policymakers, legal aid providers, and court administrators to ensure more robust access to justice.

The temporal disaggregation possible with large-scale data also reveals suggestive evidence about how policy changes affect tenant behavior, access to legal proceedings, and case-level outcomes. For policymakers and legal advocates, the ability to measure procedural barriers at scale enables more precise targeting of interventions. If research using this type of analytic pipeline identifies specific procedural barriers that disproportionately affect particular tenant populations, interventions could be designed to address those specific barriers rather than assuming uniform challenges across all tenant households or communities.

This approach represents the next frontier of technologically-enabled access to justice research. Although the access to justice crisis has been thoroughly documented in theoretical research and limited empirical work, more granular evidence and larger datasets are required to design effective, targeted interventions that promote access to legal resources for underserved and historically marginalized populations at a broader scale. Document-conditioned LLM analysis represents a promising avenue for accelerating this type of empirical research.

6. Conclusion

This study shows both the potential benefits and limits of document-conditioned LLM analysis. Although our approach created highly accurate estimates of case-level outcomes, this analytic pipeline would be insufficient in contexts requiring perfection such as in individual case adjudication or high-stakes legal determinations. Rather, the value of LLM-based analysis lies in enabling population-level empirical research on legal systems and the access to justice crisis that can generate more detailed and robust empirical evidence than manual coding methods have previously produced. The pipeline described here represents one step toward that goal, and future work should aim to expand this approach across multiple jurisdictions and longer time horizons.

Substantively, our findings provide a detailed empirical portrait of eviction proceedings that prior research has been unable to produce at scale. We show that tenants face substantial procedural barriers to participation – fewer than half submit a written response, many do not attend hearings, and access to legal representation is strongly conditioned on these behaviors. The associations between representation and favorable case outcomes, while consistently present throughout the study period, change considerably from 2022 to 2024. Because we identify changing patterns that coincide with the drawdown of emergency rental assistance and the implementation of a local ordinance in the county seat of Tacoma, this study does not causally link policy developments during the study period to procedural or case outcomes.

Although we do not draw causal findings from our work, these findings nonetheless carry direct implications for the design of tenant right to counsel programs and related access to justice interventions. The strong association between procedural engagement and access to representation indicates that outreach to tenants in the early stages of the legal process may be an effective means of increasing participation. More broadly, the ability to document case characteristics and measure procedural barriers at scale enables more precise targeting of interventions than existing small-sample studies have previously enabled. Understanding which checkpoints matter most, and for which tenant populations, is a prerequisite for designing programs that meaningfully expand access to justice in the real world.

References

Breton, Julien, Mokhtar Mokhtar Billami, Max Chevalier, et al. 2025. "Leveraging LLMs for Legal Terms Extraction with Limited Annotated Data." *Artificial Intelligence and Law*, ahead of print, March 14. https://doi.org/10.1007/s10506-025-09448-8.

Cassidy, Mike, and Janet Currie. 2023. "The Effects of Legal Representation on Tenant Outcomes in Housing Court: Evidence from New York City's Universal Access Program." *Journal of Public Economics* 222 (June): 104844. https://doi.org/10.1016/j.jpubeco.2023.104844.

Chiappetta, Casey, and Darcy White. 2026. How Court Eviction Practices Affect Families, The Pew Charitable Trusts, June 4. https://pewtrsts.org/3RH1QcJ.

Demkovich, Laura. 2024. "Eviction Filings around Washington Soar to Record High Levels." *Washington State Standard*, December 16. https://washingtonstatestandard.com/2024/12/16/evictions-around-washington-soar-to-record-high-levels/.

Desmond, Matthew, Ashley Gromis, Lavar Edmonds, et al. 2018. "Eviction Lab National Database: Version 1.0." Princeton University. www.evictionlab.org.

Graetz, Nick, Carl Gershenson, Peter Hepburn, Sonya R. Porter, Danielle H. Sandler, and Matthew Desmond. 2023. "A Comprehensive Demographic Profile of the US Evicted Population." *Proceedings of the National Academy of Sciences* 120 (41): e2305860120. https://doi.org/10.1073/pnas.2305860120.

Kim, Jae Yeon, Pamela Herd, Sebastian Jilke, Donald Moynihan, and Kerry Rodden. 2025. "Administrative Checkpoints, Burdens, and Human-Centered Design: Increasing Interview Access to Raise SNAP Participation." *Journal of Policy Analysis and Management* n/a (n/a). https://doi.org/10.1002/pam.70007.

Ribeiro de Faria, Joana, Huiyuan Xie, and Felix Steffek. 2025. "Information Extraction from Employment Tribunal Judgments Using a Large Language Model | Artificial Intelligence and Law | Springer Nature Link." https://link.springer.com/article/10.1007/s10506-025-09443-z.

Shastri, Ishana, Shomik Jain, Barbara Engelhardt, and Ashia Wilson. 2024. "Automating Transparency Mechanisms in the Judicial System Using LLMs: Opportunities and Challenges." *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 7 (1): 1357–67. https://doi.org/10.1609/aies.v7i1.31729.

Siino, Marco, Mariana Falco, Daniele Croce, and Paolo Rosso. 2025. “Exploring LLMs Applications in Law: A Literature Review on Current Legal NLP Approaches.” *IEEE Access* 13: 18253–76. https://doi.org/10.1109/ACCESS.2025.3533217.

Singh, Amrita, H. Suhan Karaca, Aditya Joshi, Hye-young Paik, and Jiaojiao Jiang (arXiv.Org). 2025. “Evaluating Customized vs. Generalist Transformer-Based Models for Legal Contract Classification.” August 11. https://arxiv.org/abs/2508.07849v2.

Sudeall, Lauren, and Daniel Pasciuti. 2021. “Praxis and Paradox: Inside the Black Box of Eviction Court.” *Vanderbilt Law Review* 74: 1365.

Summers, Nicole, and Justin Steil. 2024. “Evicted by Default.” *Connecticut Law Review* 57: 1233.

Tacoma Municipal Code, 1.95.090 (2018). https://www.cityoftacoma.org/UserFiles/Servers/Server_6/File/CMO/OEHR/RHC/TMC1.95-20230724.pdf.

The Eviction Lab (Eviction Lab). 2025. “Eviction Tracking System.” https://evictionlab.org/eviction-tracking.

Thomas, Tim. 2026. *Washington Eviction Filings*. Eviction Research Network. https://evictionresearch.net/washington/

von Geldern, Will. 2025. “Evictions, Legal Counsel, and Population Health: A Mixed Methods Study.” *Social Science & Medicine* 377 (July): 118134. https://doi.org/10.1016/j.socscimed.2025.118134.

von Geldern, Will, and Karin D. Martin. 2025. “No Right to Counsel: Evictions, Administrative Burdens, and Access to Civil Justice.” *Stanford Law and Policy Review* 36 (2): 313–48.

Westermann, Hannes, and Jaromir Savelka. 2024. *Analyzing Images of Legal Documents: Toward Multi-Modal LLMs for Access to Justice*.