

READING LAW'S GREAT UNREAD

QUALITATIVE COMPUTATIONAL METHODS, ARTIFICIAL
INTELLIGENCE, AND THE NEW EMPIRICAL LEGAL
RESEARCH

SIMON WALLACE

A DISSERTATION SUBMITTED TO
THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF LAW

GRADUATE PROGRAM IN LAW
YORK UNIVERSITY
TORONTO, ONTARIO

December 2024

© Simon Wallace, 2024

Abstract

How will new computational technologies change legal research and our visions of what law is? Inspired by the work of digital humanists, Bourdieu, and sociologists of literature, this dissertation explores how the methods of “distant reading” can be used to develop new classes of critical insights about law. After situating the project theoretically, this dissertation reports on a series of new computational studies about Canadian law.

Chapter 1 measures Canadian statutory and regulatory law, showing that law has grown unevenly over the past decade and a half. Chapter 2 uses new artificial intelligence to transcribe and analyze Supreme Court of Canada hearings, revealing gendered and linguistic speaking patterns among justices. Chapter 3 shows how computational methods can be deployed to detect inconsistency and discord in a jurisprudence, in this case Canada’s law of terrorism. Chapter 4 uses machine learning to study refugee law jurisprudence, particularly showing how it has developed over the past decade. Chapter 5 leverages new computational techniques to analyze Social Security Tribunal of Canada decisions regarding employment insurance appeals and suggests that new computational analyses might usefully change legal education. It concludes by considering how some visions of computational legal analysis—despite the sweep and scope of their projects—are part of old and traditional visions of what law is.

Acknowledgements

Already I am starting to forget: when I began my doctorate in September 2020, life looked quite different. If the promise of graduate education is conviviality and exchange, seminars and forums, and shared drafts and hallway conversations, it was an open question whether the interstitial but important components of higher education could survive us being six feet apart or only meeting through a laptop screen. For reasons large and small, the return to the university made a tremendous amount of sense to me, but it was also a gamble that— given the world was entirely reconfigured— that learning could still work.

Wagers are easier to make when they are insured. But for the unwavering moral and emotional support of my parents (Mary and Donald), my sister (Megan), and Dana’s parents (Karen and Dave), it would not have been possible to even contemplate, let alone continue, doctoral research. Remember, schools and daycares were closed, restaurants were shuttered, and we kept to our bubbles. Without everyone helping with childcare, sharing dinners, and providing stability in unstable times, this document would not have been produced.

Institutional support was critical. From a practical perspective (“Leave a secure job? In this economy?”) the funding provided by the Social Sciences and Humanities Research Council, the Ontario Graduate Scholarship Program, Osgoode Hall Law School, and York University, made all the difference. Equally, I am indebted to the Refugee Law Lab and the Centre for Refugee Studies for extending intellectually enriching and exciting employment at York University. One of the great pleasures of the past four years was teaching at the Lincoln Alexander School of Law and I extend my deep and sincere thanks to Graham Hudson for

the steadfast support and guidance as I developed my professorial identity and pedagogical philosophy.

This is a dissertation that combines materialist literary theory with legal analysis with computational methods. You might call it strange, but you might also call it very “Trent University.” Some twenty years ago I was fortunate enough to find myself in Peterborough, Ontario surrounded by people who were also excited by strange things. There are too many people here to name, but suffice it to say that if I am, as all lawyers are, hemmed in by the conservatism and narrowness of legalism (“legal education sharpens the mind by narrowing it”), my instincts to look past and beyond are most explicitly cultivated and encouraged by friends and colleagues from this time.

Which brings me to Clifford McCarten. Years ago, as undergraduates we had an argument about environmentalism in the pages of Trent’s student newspaper. I cited some data from the Soviet Union and Clifford cited Nietzsche. This, of course, meant that our differences were substantial and irreconcilable (we were, after all, in our twenties). But then we met again on the first day of law school and recognized what should have been obvious: kindred commitments, humours, and values. This is a friendship with tremendous meaning for me. Cliff makes difficult things easy. Out of law school we started a firm together, but this is only the most obvious and outward expression of how Cliff inspires everyone to do things and to build things. I am glad that, in return for all that he has given me, I was able to give him running, though it is only appropriate that—here too—his accomplishments far exceed my own.

At Osgoode, I was fortunate to meet many new friends and colleagues, even if some I still have never met in person. I am particularly grateful to have met Catherine Le Guerrier, who is bound to be an outstanding researcher and teacher. I was equally fortunate to have met Ruth Buchanan and Susan Drummond—each of whom provided invaluable assistance to me as graduate program directors. I also must thank Stephanie Ben-Ishai and Carys Craig for introducing me to new areas of scholarship that I otherwise would not have encountered. Craig Scott has, even when he was overburdened with work as Associate Dean, been a tremendous support—thank you. I am indebted to Jon Penney for his comments and feedback

on the dissertation as it developed and to Amy Salyzyn and Özgün Topak for examining the final product.

Parts of this dissertation have already been published or are in the publication pipeline. Chapter 4 was published in the Dalhousie Law Journal, an abridged version of Chapter 5 will hopefully be published in an edited collection soon, and Chapter 3 was published in the Supreme Court Reports. Each of the editors and reviewers who worked on these essays made them significantly better. To Anonymous Reviewer #1 of my paper on terrorism deportations (Chapter 4), your detailed and thoughtful report made me a believer in peer review.

But unquestionably the best editor I have ever encountered is Sonia Lawrence. For this reason—and for many others—I am so pleased that Sonia agreed to chair my dissertation defence. The reason that the text in this dissertation is largely free of ornamentation is because of Sonia. Years ago, Sonia quoted Lucy Maude Montgomery to me, explaining that in *Emily's Quest* a dying teacher uses their last breath to issue a most important warning to their student: “Beware — of — italics.” Point taken: there are not *too* many italics here. But Sonia’s guidance is, as it has for every student who has encountered her, about much more than just style. Every time Sonia has read any of my work, she has returned with the most incisive and useful commentary. Her care and enthusiasm for good writing, good research, and lively thinking is simply infectious. I am just so grateful to have learned from, and been inspired by, her example.

Graduate school at Osgoode is a return of sorts. Just over a decade ago Dana and I met in Ben Berger’s criminal law class. Memory is a funny thing. We know that it would be just as factually accurate to say that we met in the hallways, or on the 196 bus, or in (horror) our property law class, but we are certain that we met in Ben’s criminal law class. How to account for this? When we both look back it is undeniable that Ben’s class was the most dynamic and exciting of our first-year experience: sometimes learning law can be rote, but sometimes it can reshape how you understand the world. Put differently, in our memories’ Ben’s class was the educational complement to our experience of falling in love. Perhaps this reads as a bit gratuitous, but Ben will understand my point: I am paying the greatest compliment I can to a teacher.

For autobiographical reasons, then, it was such a pleasure to reconnect with Ben when he agreed to serve on the committee. But it was just as much a pleasure to work again with Ben because he so thoughtfully and generously invited me to explore terrain that I would have otherwise failed to consider or failed to properly grapple with. At bottom, I see this dissertation as a humanist one that bears Ben's humane and generous imprint throughout.

Now that the dissertation is written and the defence is wrapped up, I can safely make this declaration: I am a Rehaagist. Let me say three things on this. First, Sean's contributions to knowledge are singular. I am uniquely positioned to observe how the study of law in Canada, particularly the study of refugee law and administrative law, turns around Sean's work. It is not an exaggeration to say that Sean has generated knowledge that has enabled whole new ways of understanding the relationships between law and society. There is a reason that every public law litigator and decision-maker in the country knows Sean's work: it is work of consequence.

Much of this is due, second, to Sean's commitments to methodological innovation and interdisciplinarity. To read Sean's work is to encounter pluralist accounts of law, informed by history, empiricism, steadfast commitments to human dignity, and large analyses of the law-in-action. To the extent that my own work in this field is possible, it is only because of the paths broken by Sean. Why could a graduate student with no coding experience conceive of the value of computational analyses? Because Sean showed the way. How could a graduate student orient to the "great unread" of the legal archive? Because Sean created, collected, and published the data. It is simply a fact that this entire project depends on the methodological, archival, and theoretical paths broken by Sean.

Third, to be supervised by Sean is to be supervised by a scholarly citizen of the highest order. To say that Sean is generous is to understate the point. At each stage of my graduate journey, Sean went well beyond what could ever be expected of a supervisor to ensure that the work was advancing, that there was funding available, and that I could take risks. Putting it together, whatever trajectory I am on, I am on it because Sean made it possible. And that is the bottom line: to be a Rehaagist is to do a category of work, and want to do this work, that could not be done but for Sean.

Graduate studies correlated with a great transformation in my life: having children. This text is therefore dedicated to my daughters, who have made each day just tremendously fun. Once a great philosopher's daughter asked her father what his favourite maxim was and he replied, "Nothing human is alien to me." Watching Poppy and Rosa begin to discover and explore books, school, friends, and the world has made me appreciate that sentiment anew. There is nothing quite like spending time with curious people to make one curious too.

And now, Dana. It is to you that my greatest debts are owed. This is a document about what measurement and quantification can illuminate. But you prove other illuminations matter much more. At the end of a long opera about Einstein, someone asks an important question. How much do I love you?

Count the stars in the sky;
Measure the waters of the ocean with a teaspoon;
Number the grains of sand on the seashore.
Impossible, you say?

Yes, and it is just as impossible for me to say how much I love you.

Statement on Generative Artificial Intelligence

At the time of writing, York University has not implemented a policy regarding the use of generative artificial intelligence for graduate work, but it is only a matter of time. It appears that through schools the world over there is a concern (perhaps even a majoritarian one) that generative artificial intelligence *necessarily* threatens the integrity of the scholarly project. Sometimes teachers worry that students will short-circuit the hard work of learning by asking generative artificial intelligence tools to produce their essays in seconds. Other times the concern is that generative artificial intelligence will make students lazy, preferring to rely on the outputs of a computer instead of doing independent research. To a degree, these concerns are well-founded. But only, I think, to the degree that these have always been concerns. It is not as if artificial intelligence has created cheating, nor is it that some students only now have started to search out short-cuts to learning. Plainly: the panic over pedagogy and artificial intelligence looks to me like a familiar one, just dressed up in new clothes for the technological moment.

Nonetheless, some of the ways that universities are responding to this old, made new again, concern seem positive. Chief among these innovations is an increasingly standard requirement that students who use generative artificial intelligence disclose the fact of the usage. This seems to me to be entirely in keeping with largely scholarly traditions: we are all supposed to disclose our methodologies, we are supposed to set out our approaches so others may replicate our results, we are supposed to be transparent about authorship and attribution.

While it is no easy task to simply define generative artificial intelligence so that its use can be disclosed (is spellcheck generative artificial intelligence? Are old machine learning algorithms artificial intelligence? Are older rule-based computer programs artificial intelligence?), the overall spirit of the rule is familiar and it is good.

So, here is the statement: this dissertation would have been impossible to produce without generative artificial intelligence. Let me explain how I used it. The first, and most important, disclosure will be repeated throughout the dissertation. Some of the research here depends, and depends explicitly, on artificial intelligence research tools. As you will soon see, I used advanced community detection algorithms, facial recognition algorithms, voice recognition, speaker diarization, speech recognition, sentence parsing, part-of-speech detection, summarization tools, embedding generators, and clustering algorithms to develop the insights I report in these pages. If generative artificial intelligence is really just developed forms of computational statistical learning, where computers learn from examples as opposed to depending pre-specified coded rules, then this is a generative artificial intelligence project. Truthfully, that does not strike me as a particularly helpful framing (just because popular culture has discovered the power of this form of machine learning does not mean that it is a new discovery, unknown and unused to previous generations of scholars). Be that as it may, what follows is an exploration of what the absolute state-of-the-art makes possible for legal research. So, yes, generative artificial intelligence is here because a major point of this dissertation is to think about what this new technology can do for legal scholars.

In a second way, it would have been extraordinarily difficult for me to undertake this work without generative artificial intelligence. I am a self-taught coder and, while I characterize my skills as advanced, they were not always advanced. To learn how to program I used a range of open-source materials and courses, but I also used tools like GitHub's Co-Pilot and ChatGPT to help me debug and develop my code. My view is that this is one of the enormous pedagogical opportunities of generative artificial intelligence. Before the artificial intelligence moment, it would have taken years of sustained study to develop my capacities to the point that they are at now. I am convinced—though I cannot empirically prove it—that generative artificial intelligence made me a more capable coder by helping me solve problems and develop programs in a fraction of the time it otherwise would have taken. We are often

promised that technology will introduce efficiencies into our lives and our learning, but these promises frequently fail to materialize. Here, I can report that there is a good news learning story: artificial intelligence can advance learning meaningfully. It also has helped me reduce some of the tedium of dissertation writing. For example, to build my bibliography I developed a simple artificial intelligence program that extracted every citation from my footnotes and classified it, so that I could simply export a final bibliography, properly formatted and organized. This tool's final product was not perfect, but it saved me hours of tedious work.

But, of course, these are not really the concerns that explain why universities are scrambling to regulate generative artificial intelligence. The concern that is most frequently cited appears to be about writing: did someone have artificial intelligence do their writing for them? In this case, no. I experimented with using artificial intelligence tools to improve my writing using different models and approaches. Over the years, I have asked programs to return my writing in syllogisms, to see if the program's reading matches my intention; I have asked the program to correct my grammar and spelling; and I have asked it to review my writing holistically to give feedback on how it might be improved. Here is my report: for what I wanted, generative artificial intelligence was not useful.

I do not claim to be a stylish writer, but I do claim that I would like to be one. Over and over, I have found that generative artificial intelligence systems produce bland, boring, and derivative outputs. The writerly advice of a generative artificial intelligence system strikes me as routine and repetitive, stressing a few rote notions: thesis statements, for example, should be clearly stated! Traditional spellcheck remains superior. And I have never received computational feedback that came even close to matching the quality of the feedback of an interested human editor. I acknowledge that others have found that there are forms of writing for which generative artificial intelligence can be useful, but what unites these forms? The professional email, the brief memo, the re-writing of meeting notes: in each case, excellence in the form is marked by consistency and predictability. No one is looking for creativity or stylistic originality in the office email that outlines some changes to some human resources policy. In those domains (where good writing is recognized for cleaving to a style, instead of developing it), generative artificial intelligence as a writing tool excels.

Research in this field is yet young, but I sense that this is a feature of the political economy of artificial intelligence. Does it surprise you that artificial intelligence developed with billions of dollars of capital from a few leading tech firms has, embedded deep within it, the stylistic preferences of the corporate culture? Ask your preferred generative artificial tool to write you an email on any topic and I will wager that, try as you might, it will return some draft that expresses a few mildly optimistic themes. The future will bring great and meaningful change. Balance is important. But we must act with intention and care, mindful of our responsibilities to each other. Now think of the runners listening to inspiring music, the smiling families, and the friends connected across oceans, each of whom is regularly featured in the advertisements of large technology firms. Ask yourself: how would these people write an email?

Against the major generative artificial intelligence products, open-source developers and artists point to a different vision of what generative artificial intelligence might do and suggest, perhaps, that stylistic and creative innovation may be aided by new machine learning technologies. In “Frontier,” the artist and musician Holly Herndon uses a generative artificial intelligence model, trained on thousands of different voices, to sing and reinterpret the Appalachian music she remembered from her childhood. The result is new and old. Quoting electronic music, choruses, and tradition, Herndon presents us with music that surfaces tensions between individual artistic production and our relationships with machines and each other. First, a single vocalist sings a few notes then, second, they are joined by a machine chorus heralding a new future: “Don’t listen to what the silvers say / Their world is long gone / A leap or push...” The end of the song is cacophonous, as human vocalists and the machine choruses sing back and forth, becoming indistinguishable from each other.

The point: our moment (like each one that preceded it) offers us new questions, new answers, and new lines of sight on social life. I worry that the anxieties of the professoriate, narrowly focused on the impact of old technologies made new and made shiny by large technology firms, focus our attention and our students’ attention on the most banal and uninteresting parts of generative artificial intelligence. Against these quotidian concerns, and against the suggestions of a tech leadership who (as they seek to raise billions in capital) make large promises about what artificial intelligence will do, it might be better to more closely

examine the paths broken by those artists who wonder how we might see the world differently, those computer scientists who have long toiled in the basements of university buildings, and those researchers motivated first by their curiosity. There are tools here that can be used to expand knowledge and those of us in the university ought to encourage experimentation and exploration. So, to end, my final statement on generative artificial intelligence: may this dissertation at least gesture to some of the ways that we might intentionally and usefully use this technology to open new perspectives on the ways we live and ways we might live.

Table of Contents

Abstract	ii
Acknowledgements	iii
Statement on Generative Artificial Intelligence	viii
Preface	xvii
1. Introduction: Reading from a Distance	1
1.1 Law’s “great unread”	2
1.2 Close reading and the discovery of form.....	10
1.3 Literature (and culture) as a symptom of material life.....	17
1.4 The sociology of literature and Bourdieu	22
1.5 Modelling text and the digital humanities.....	25
1.6 Reading legal text at scale	30
1.7 The plan for this dissertation.....	33
2. The Measurement of Federal Law	37
2.1 Scholarly Context.....	40
2.2 Generating a Corpus of Canadian Law.....	44
2.3 Taking Law’s Measurements	47
2.3.1 A Growing Law	47
2.3.2 Of big laws and small laws	49
2.3.3 Law grows and law churns	52
2.3.4 A new vantage: peeking at the micro	54
2.3.5 Summary: a dynamic and growing law	56
2.4 Mapping Canadian Law	57
2.4.1 Using network theory to spatialize law	57
2.4.2 Inferring a community structure in Canadian law	61
2.4.3 Looking back at the micro.....	66
2.5 Discussion and conclusions	69
2.6 Appendix A: Genres of law (2007).....	72
2.7 Appendix B: Genres of law (2023)	73
3. Speaking Like a Judge	74
3.1 Building an AI transcriber and transcribing the 2021-2022 SCC term	78
3.2 A gentle introduction to neural networks.....	79
3.3 Analyzing and validating machine-generated transcripts of 51 SCC oral hearings	83
3.4 Who speaks, when, and in which language?	86
3.5 Judges who are men speak more than judges who are women.....	90

3.6 Non-Québécois judges rarely speak French	93
3.7 The research opportunities and limitations of computational methodologies	95
4. Persistent Discord: Inconsistent Adjudication of Terrorism Cases in Canada	98
4.1 Canadian national security deportation law	101
4.1.1 Bases for security inadmissibility	104
4.1.1.1 Espionage against Canada or against Canadian interests.....	104
4.1.1.2 Subversion of any government by force or subversion of any democratic government by any means.....	105
4.1.1.3 Terrorism.....	106
4.1.1.4 Residual grounds	107
4.1.2 Adjudication	108
4.1.2.1 Adjudication before the Immigration Division	109
4.1.2.2 Appeals before the Immigration Appeal Division.....	110
4.1.2.3 Judicial review before the Federal Court of Canada	110
4.1.3 Relief available for inadmissible non-citizens.....	111
4.2 Empirical study of security deportation cases (2018–2020).....	113
4.2.1 Methodology	116
4.2.1.1 Access to information request	116
4.2.1.2 Development of a Python program to analyze Federal Court of Canada dockets.....	118
4.2.1.3 Reviewing each Federal Court of Canada case file	119
4.3 High-level analysis of the dataset	120
4.3.1 The Immigration Division is Canada’s busiest national security tribunal.....	121
4.3.2 Most national security cases were brought against asylum seekers from the Global South	121
4.4 Case study of deportation cases concerning members of the Bangladesh National Party	123
4.4.1 Evidentiary basis for the BNP terrorism allegation.....	124
4.4.2 Adjudication before the Immigration Division	126
4.4.2.1 Inconsistent outcomes and reasons between adjudicators.....	127
4.4.2.2 Canada led functionally identical evidence across cases	128
4.4.2.3 Montreal-based adjudicators recycled their reasons across cases.....	129
4.4.3 Adjudication of judicial reviews before the Federal Court of Canada.....	133
4.5 Discussion and Recommendations	139
4.5.1 Recommendation 1: Tribunal and Court administrators should help decision-makers look beyond the four corners of each case.....	140
4.5.2 Recommendation 2: Supplement traditional research methodologies with computational methodologies	142
4.5.3 Recommendation 3: Consolidate appeal routes	143
5. Modeling Canada’s Refugee Law Jurisprudence.....	145
5.1 Using computational methods to read from a distance	149
5.1.1 The study of the “great unread”	150
5.1.2 Using computational approaches to uncover patterns in text data.....	154
5.1.3 The value of validation.....	158
5.2 The study: measuring change in the Federal Court of Canada’s refugee law jurisprudence .	159
5.2.1 Collecting the dataset	161
5.2.2 Isolating for refugee related decisions	163

5.2.3 Partitioning the decisions into clusters	166
5.2.4 Visualizing multi-dimensional data.....	170
5.3 Findings.....	172
5.3.1 Canadian refugee law jurisprudence is increasingly claimant-negative	172
5.3.2 Visualizing and clustering refugee law jurisprudence shows that it is not topically diverse.....	178
5.3.2.1 Mapping Canadian refugee law.....	178
5.3.2.2 Clusters of refugee law	180
5.3.2.3 Summary	194
5.3.3 Jurisprudential change over time	195
5.3.4 Algorithmic insights: more than just a sum of different parts?	197
5.4 Discussion	198
5.4.1 An economic case for intermediate appellate review?.....	199
5.4.2 Learning the law by knowing the jurisprudence’s shape	200
5.4.3 What accounts for jurisprudential change? And what do lawyers do?	201
6. Law School Goes Back to the Laboratory.....	205
6.1 Employment Insurance (“EI”) and the Social Security Tribunal of Canada (“SST”).....	209
6.1.1 Origins of the EI program.....	209
6.1.2 The establishment of the SST	211
6.1.3 Little is known about the substance of the EI disputes brought to the SST	214
6.2 Designing a computational study to understand the SST’s EI jurisprudence	215
6.2.1 Data acquisition	216
6.2.2 Data transformation: using start-of-the-art artificial intelligence models	218
6.2.3 Clustering the data	221
6.2.4 Quantitative validation	223
6.2.5 Qualitative discovery and theory building	224
6.3 Report on the clustering.....	226
6.3.1 LLM text embeddings produce useful clusters for analysis, but can hallucinate relationships between decisions	227
6.3.2 Qualitative analysis of the clusters.....	230
6.3.2.1 Dark pink: Refusing extension requests for reconsideration requests.....	231
6.3.2.2 Light brown: Good cause for delay in making an application for benefits.....	233
6.3.2.3 Light pink: Making the wrong election for parental leave benefits	234
6.3.2.4 Light green: The non-reviewability of overpayments	237
6.3.2.5 Dark olive: Definition of earnings and the allocation of earnings.....	239
6.3.2.6 Light olive: Consequences for misrepresentations	241
6.3.2.7 Light purple: Disentitlement to regular benefits.....	243
6.3.2.8 Dark green: Disqualification from coverage because the claimant voluntarily left their employment.....	245
6.3.2.9 Dark brown: Disqualification from coverage because the claimant’s employment was terminated because of misconduct.....	248
6.3.2.10 Dark brown: Disqualification from coverage because the claimant’s employment was terminated because of misconduct (non-compliance with mandatory vaccination policy)	250
6.4 A theoretical model of EI law amenable for teaching	253
7. Conclusion: The Paths of Computational Law.....	256
8. Bibliography.....	282

Legislation.....	282
Jurisprudence.....	282
Secondary material	288
Monographs and collections	288
Scholarly articles and chapters	291
Other.....	299

Preface

This dissertation asks: how will law be changed by its encounter with new, sophisticated computational methodologies? It answers: the significance of precedent, rules, and canons will diminish; the meanings of habits, patterns, and deep trends will come into sharper focus. The study of law has always been suspended between idealism and experience. But artificial intelligence, machine learning, big data, networks: these are tools that give new power to empirical and theoretical descriptions of law's shape based on inductive reasoning.

But this question and answer are not this dissertation's subject. No, it is a frame that came to me only lately. When I returned to graduate school I was seized by a factual question. I had just finished a stint as a deportation and detention lawyer. With an office in Canada's largest immigration detention facility, I represented an untold number of people in their emergency, last-ditch efforts to stop their removals from this country or to get out of jail. This was edifying and, for a while, spiritually invigorating work. I met many extraordinary people who lived extraordinary lives in extraordinary circumstances, far beyond the rhythms and routines of my own life.

There was a class of case that came to deeply trouble me, in part because of its familiarity. Non-citizens, even those with some forms of status, can be deported from Canada if they are convicted of a sufficiently serious crime. And sometimes they will be kept in immigration detention after serving their criminal sentence while the government arranges for their removal. Sometimes this can take weeks, months, or even years. As I

worked with detainees facing criminality deportations, obvious similarities between us presented themselves. Usually we were the same age; more than a few times we shared birth dates. We were men and sometimes we went to the same high school, or at the very least we knew the same neighbourhoods and places in Toronto. Like me, we grew up as kids in Canada. But I had citizenship and was safe from deportation, while they did not and were not. And, almost invariably, they shared three common characteristics: they were brought to Canada as infants, they had some form of psychiatric diagnosis, and they were Black.

A confusing part of lawyering is that we often do not talk about those obvious things that matter most. In court and in tribunal rooms I parsed the reasonableness and coherence of legal decisions. I made arguments about risk and balance; public safety and personal liberty. As a lawyer, I was brought up to believe that these arguments might be useful and, yes, sometimes they were. Sometimes they stopped or delayed deportations and sometimes they helped get people out of jail. But I found them to be useless against that which I wanted to inveigh against: the obvious racist pattern in front of me. Remarkable: a pattern all of us could see, but the law—with all of its regularity—reinforced, while denying me even the language to talk about it.

The great promise and appeal of the university is that sustained study and inquiry will develop new theoretical models of the world, to let us better understand it and maybe shape it. This explains my return to school. I, like many others, saw something and wanted the chance to explore what I observed, believing that knowledge might beget something worth begetting. But exploration is not as simple as finding four years to reflect quietly. Research needs evidence and it needs a way of marshalling, organizing, categorizing, and reasoning about that evidence.

I have a quarrel with law schools: the tyranny of idealism and deduction. Especially in our teaching, but also in our research, law schools often encourage us to work with first principles downwards. Do not get me wrong: deduction is a critical and necessary mode of reasoning for lawyers and researchers. Indeed, for all of us. If part of the purpose of theoretical models is to help us predict the behaviour of the world, deduction is what gives

life to our descriptions of the world. But, and here is my claim, in law schools we do not spend enough time or effort building those theoretical descriptions.

So, considering my interests, motivations, and objectives, my first work in graduate school was decidedly inductive. History: working with archival materials, I sought to understand the circumstances that produced the legal technologies of deportation and immigration detention.¹ Sociology: working with dozens of detention transcripts, I set out to explore who gets described as a dangerous immigration detainee and who does not.² Doctrine: working with hundreds of terrorism and war crimes deportation decisions, I aimed to understand the on-the-ground shape of these grave areas of the law.³

And then a methodological nudge. My first two years of graduate school featured boxes and boxes of casefiles. Yet hopelessly incomplete. I can only read so much and, as a lawyer, I live in an era where the consequences of these physiological limitations are thrown into sharper relief. I might read half a dozen casefiles in a day, then glance at Canada's online jurisprudential database and note the tens of thousands of cases out there that neither I nor anyone else would ever read. To ensure the validity of my samples, I always selected narrow research questions that led me to read a narrow part of the field comprehensively. But the narrowness of my answers, measured against the largeness of all that law I could see out there, made me want for more.

How might this larger law be read? Sean Rehaag pointed to a solution: we are physiologically limited in ways that computers are not. Maybe, he and his work powerfully suggest, we might set the machine on that information out there that interests us and, just maybe, it will come back with something we can use to describe law in its largeness.

And here the dissertation properly begins. I write at the dawn of a new computational age. While, as readers who make it through to the conclusion will see, I do

¹ Simon Wallace, "'Police Authority is Necessary': The Canadian Origins of the Legal Powers to Detain and Deport, 1893-1902" (2023) 48:2 *Queens LJ* 101.

² Simon Wallace, "The next jailor: An empirical study of danger to the public immigration detentions in Canada (summer 2021)" (2024) 26:3 *Punishment & Society* 566.

³ Simon Wallace, "The new Canadian law of refugee exclusion: An empirical analysis of international criminal law deportation orders, January 2018 to July 2020" (2022) 22:4 *International Criminal Law Review* 721.

not understand this moment in cataclysmic terms (the singularity, no, is not upon us), I do understand that we can do things with computers today that we could not yesterday. It is too early to confidently periodize and describe this moment, but this much is clear: if computers once operated with reference to sturdy rules and code (deduction, once again), they now begin to operate with reference to bodies of knowledge and learned patterns (induction, once again). The full implications of this transition are not known to us, but I do know that they have made me a more competent programmer than I could have ever hoped to be.

The work that follows ought, I hope in the best possible way, to be described as grasping: what will this emergent computational paradigm allow us to see? Once I read, although the source is now lost to me, that computational methods give researchers a new optic. Many methodologies illuminate the micro (think here of interviews, archival study, focus groups, studies of samples) or the macro (think here of humanist studies of the grand sweep of human history and philosophy) but we struggle to see human affairs on a meso scale. This is the optic, I now believe, that computational methods unlock for legal researchers. Institutions, the evolution of legal forms: these rarely change overnight, but by the slow accretion of changes, habits, and patterns.

Let me show you something. The Supreme Court of Canada is this country's highest court, deciding appeals from across the country. Its judgements form the most authoritative statements of the law. The Court's judgements, however, do not just issue from Zeus' head. The Court makes a point of embedding, referencing, and citing to other cases. This makes law look like an incremental effort: each new case builds on some insights from past cases. But some moments are more incremental than others, some cite deeper into the past, while other moments locate their interest more in the present.

What about our moment? Figure 1 was simply made. I wrote a script that identified every case cited in every Supreme Court of Canada judgement. I then took the age of each case cited (the year of the original judgment minus the year that the cited case was decided) and took the mean for each case. Then I calculated the mean citation age for each case, and then calculated the mean citation age for each year. Finally, the visualization shows the five-year rolling average from 1950 to 2022.

What do we see? In the 1970s, a jurisprudential age that reached deepest into a legal past, during which the average citation was about sixteen years old. Then, with the repatriation of the Constitution in 1982, an age of relative novelty. Cases of influence became younger and younger before, in the new millennium, a reversal of that trend. Beginning around the year 2000, the Court began to reach back further for citations. To be sure, more study is required, but do we not see major (though profoundly intuitive) changes in the nature of law? In a moment of constitutional change, jurisprudence became




Figure 1: The average age of a citation in the caselaw of the Supreme Court of Canada (five-year rolling average).

younger as it metabolized new legal arrangements. But then, eventually, law settled, and zeroed in on a few, now older, authoritative precedents.

And the form of judgements has stabilized. Standard deviation measures how much variation there is from the mean. A low standard deviation means that most data points cluster closely around the average, while a high standard deviation means that data points are spread out more widely around the mean. When we look at Figure 2, we see

that, since the 1980s, judgements look—at least in this regard—more alike. While the 1970s were characterized by wild differences in the average age of decisions, now the Court has settled into a new stable rhythm. We can expect now, as opposed to in the 1970s and




Figure 2: The standard deviation of the mean citation age in the Supreme Court of Canada's jurisprudence (five-year rolling average).

1980s, judgements to feature the same average age of cited cases.

We can look at this in a different way. In Figure 3, on the left we see the distribution of cases by average age of citation in the 1980s and, on the right, we see the same but for the 2010s. Look: in both decades the most decisions had a very low average citation age (perhaps these are quick, one-off judgements discussing matters the Court deems to be of little import). But then, in the 2020s, a different form: there is a major bulge showing that




Figure 3: The distribution of average ages of citations for cases in the 1980s (left) and in the 2020s (right).

the Court's judgements turn on a more common average age of citation. Deep trends in the law that suddenly become apparent when we measure aspects of the text. Does it not look to you like, in the 2010s, there is a body of law of a more or less consistent age that the Court engages with? And does this not look different from the jurisprudence of the 1980s? Or, are we seeing the court cite to the same body of law (maybe early *Charter* jurisprudence) that is simply aging?

All of this may strike you as a long way off from where I began. Back in that office in the immigration holding centre we did not talk the changing averages of citation ages in the Supreme Court of Canada jurisprudence. But we did talk about the law, its stability and instability, and its deep patterns. This dissertation does not explicitly address the distress about systemic racism that inspired by return to school (though, in fairness to myself, this is a matter that continues to occupy my attention),⁴ but it tantalizes. Soon, this dissertation will ask us to consider together how our justice system tolerates the inconsistent and unpredictable application of the law of terrorism on citizens of the global south; how judges who are men speak more in Court than judges who are women; how Indigenous reconciliation has only lately become a feature of statutory and regulatory law; how refugee law jurisprudence has shifted focus in the past decade; and how employment insurance disputes are adjudicated in Canada.

Once we get our hands on more data, more decisions, more evidence of experience, the methods I work with here promise to illuminate other dark corners of the law, in turn enabling the construction of new theoretical perspectives and models. Bourdieu, whose insights haunt this project, once wrote that “[t]he mind is a metaphor of the world of objects.”⁵ A larger focus on legal texts, a focus we have not previously had access to, may well give us new vantages on some of our metaphors and some of our relations with each other and the world. But these are questions I take up another day. For now, the fruits of my grasping.

⁴ See “Police Authority is Necessary,” *supra* note 1; “The Next Jailor,” *supra* note 2; and “The New Law of Refugee Exclusion” *supra* note 3.

⁵ Pierre Bourdieu, *Outline of a Theory of Practice* (Cambridge: Cambridge University Press, 1977) at 91.

1. Introduction: Reading from a Distance

Not reading is the dirty open secret of all literary critics.¹

[S]ome people have read more, but the point is that there are thirty thousand nineteenth-century British novels out there, forty, fifty, sixty thousand—no one really knows, no one has read them, no one ever will. And then there are French novels, Chinese, Argentinian, American... Reading ‘more’ is always a good thing, but not the solution.²

If one dimension is lost, in the shift from hermeneutics to quantification, one dimension is gained: we still know almost nothing about how literary systems function, and that’s exactly what the logic of measurement shifts our gaze towards.³

Culture. Theory. Data. In that group of three terms, the obvious interloper is data.⁴

It would be nice if, one day, big data could lead us back to big questions.⁵

¹ Lisa Marie Rhody, “Beyond Darwinian Distance: Situating Distant Reading in a Feminist *Ut Pictura Poesis* Tradition” (2017) 132:3 PMLA 659.

² Franco Moretti, “Conjectures on World Literature” in Franco Moretti (ed.), *Distant reading* (London: Verso, 2013) at 47.

³ Franco Moretti, “The Roads to Rome” (2020) 124 New Left Review 125.

⁴ Laura McGrath, Richard Jean So & Chad Wellmon, “Culture, Theory, Data: An Introduction” (2022) 54:1 New Literary History 519 at 519.

⁵ Franco Moretti, *Literature, Measured*, Pamphlets of the Stanford Literary Lab (Stanford, 2016) at 7.

1.1 Law’s “great unread”

How will technology change the study and practice of law? I begin my answer by pointing to some cases: *McNally Contractors (2011) Ltd. v. Director, Alberta Environment and Parks*,⁶ *Blois v. Butcher*,⁷ *J.M.A. v. L.A.M.*⁸, *Appellant v Northern Regional Health Authority*,⁹ *Vemb v. The Samuel T*,¹⁰ *Gupta v. The King*,¹¹ *Haben Abrham Weldekidan v. His Majesty the King*,¹² *Layzell v. Ontario (Human Rights Commission)*,¹³ *EL (Re)*,¹⁴ *Algonquin College of Applied Arts and Technology (Re)*,¹⁵ and *The Highlands Residential Association v. Welland (City)*.^{16 17}

Each of these cases is unremarkable. When I say that I do not mean that each case is uninteresting, unimportant, or inconsequential. I mean that no one will remark on them. Even though there are more lawyers, judges, and law professors than ever before, the text from each of these cases will, the moment the case is closed, go unanalyzed. Every day, an uncounted but massive volume of legal text is produced, considered, reasoned with, and then filed in the archive—where it is never read again.

Why? It is no mystery: you and I can only read so much. You study criminal law, contract law, or refugee law? No, you study its “canonical fraction.” Beyond the law that ends up in the books, the law in action produces a massive quantity of data. Consider these numbers: in 2022, the Canadian Immigration and Refugee Board heard approximately

⁶ 2023 ABEAB 6.

⁷ 2023 NSSM 26.

⁸ 2023 MBKB 110.

⁹ 2023 CanLII 56246 (MB HAB).

¹⁰ 1971 CanLII 1976 (CA EXC), [1970] Ex CR 1036.

¹¹ 2023 TCC 82.

¹² 2023 CanLII 57199 (SCC).

¹³ 2003 CanLII 23954 (ON SCDC).

¹⁴ 2015 CanLII 94649 (ON CCB).

¹⁵ 2022 CanLII 122328 (ON IPC).

¹⁶ 2021 CanLII 48340 (ON LPAT).

¹⁷ Let me acknowledge early on an important intellectual debt. Franco Moretti, the literary scholar who is often credited with founding the digital humanities, opened an important essay on literature’s unread the same way: by listing books that were forgotten but were, in their own times and in their own ways, important: Franco Moretti, “The Slaughterhouse of World Literature” in Franco Moretti, *Distant reading* (London: Verso, 2013) at 65.

25,000 applications for refugee protection, issuing explanatory written reasons in each case.¹⁸ If we assume that each decision was 2,000 words long, that year's first instance refugee law jurisprudence equalled 50 million words. If we assume an average novel is 60,000 words, a researcher who set out to study the entirety of that year's case law would have to read the equivalent of 833 novels, and tens of thousands more if they decided to also read each transcript, each filed exhibit, each appeal record, and so on. An impossible task.

So instead of reading at scale, legal scholars read selectively or read around. Scan the tables of contents of the law journals and you will find thoughtful engagements with doctrinal and policy questions, close readings of leading cases, ethnographies of institutions, and studies based on interviews with system actors, but you will rarely find a sustained engagement with legal text at scale. The methods of traditional legal research, as illuminating and important as they are, have no approach on the "great unread."¹⁹

But perhaps law's unread is a feature, not a bug, and something that should not concern us. The great strength of legal systems, particularly common law legal systems, is their ability to pull out a few important signals from the noise of all legal disputes. The mechanisms of judicial review and appeal require judges to synthesize and justify their readings of first instance law. These readings get recycled back into the legal system and prescribe patterns of thought and behaviour for legal actors. For this reason, the texts from "law's capitals," staffed by princely judges, are rightly influential and revelatory because appellate law offers the best access to the distilled rules of argumentation that all system actors are expected to follow.²⁰ Understand law's canon, understand the law. Change the canon, change the law.

¹⁸ The statistics show that 45,555 individual applications were closed. Some matters, usually families who claim together, are heard together. For this reason, I estimate that 25,000 cases were decided that year. See Immigration and Refugee Board of Canada, "Claims by Country of Alleged Persecution - 2022", online: <<https://irb.gc.ca:443/en/statistics/protection/Pages/RPDStat2022.aspx>>, Last Modified: 2023-03-01.

¹⁹ Margaret Cohen, *The Sentimental Education of the Novel* (Princeton: Princeton University Press, 2002) at 23.

²⁰ Ronald Dworkin, *Law's Empire* (Cambridge: Belknap Press, 1997).

Or maybe it is the other way around. Beneath the judges, a whole social universe of legal actors compete to define legality. The words of apex courts may get used this way or that way, or they may be ignored altogether. Different adjudicators will consider the same facts, but arrive at different conclusions, even though their analyses may cite and use the same legal ideas. Legal tropes develop, rise, adapt, and fade. “Give me the facts,” Bourdieu says, “and I will give you the law.”²¹ Not the idiosyncratic facts of each case, but the facts about how the juridical field—the space where legal competitions occur—is arranged. If we understood the relative power of actors, their unconscious ways of thinking, their “feel for the game,” their disputes and motivations, we could explain the deep conflicts and competitions of law’s lived, as opposed to bookish, reality. Better, Roscoe Pound said, to know about the law in action than the law in the books.²²

For this latter perspective, legal text at scale is an invaluable evidentiary resource. In amongst those 25,000 refugee cases are instances where adjudicators disagree but reference the same doctrinal concepts; there are claimants advancing new sorts of claims; there are official and non-official actors adapting their circumstances to their perceptions about jurisprudential reality; there are people who ignore precedent, who misunderstand it, who apply it inconsistently; there are ideas that will one day find some expression in some appellate judgement; there are ideas that will, once some adjudicator retires, disappear. Put differently, the unread is an extraordinary sociological resource. There is no better evidence available to us regarding the legal present and recent legal past than the textual residue each legal matter leaves behind.

But the trick is in the reading. Legal scholars are accustomed to paying close attention to each text before them, but this type of reading just will not work with text at scale. Just beyond our discipline, however, literary scholars have developed new ways of reading, new ways of mining literary texts for insights, that legal researchers could learn from. In many ways, our concerns are similar: what is the role of the canon? How is the canon produced? What visions and voices does the canon privilege? Which does it ignore?

²¹ Pierre Bourdieu, “The Force of Law: Toward a Sociology of the Juridical Field” 38 *Hastings LJ* 805 at 814.

²² Roscoe Pound, “Law in Books and Law in Action” (1910) 44:1 *Am L Rev* 12.

How do systems function to produce forms of art or law? The *Law and Literature* movement (particularly the subfield of *Law as Literature*) already authorized legal scholars to use the tools of literary analysis to explore law, though it is not an approach that ever really caught on. But now, I argue, there are literary tools there that are uniquely useful to legal researchers. We ought to use them.

So, what are these new literary methods of analysis? Careful attention, not to each individual text, but to the common forms across texts. The goal: to model how forms—those repeating patterns, devices, genres, rules, categories, modes of analysis—develop and change across texts. Instead of close reading, distant reading:

[I]f you want to look beyond the canon... close reading will not do it. It's not designed to do it, it's designed to do the opposite. At bottom, it's a theological exercise—very solemn treatment of very few texts taken very seriously—whereas what we really need is a little pact with the devil: we know how to read texts, now let's learn how not to read them. Distant reading: where distance, let me repeat it, is a condition of knowledge: it allows you to focus on units that are much smaller or much larger than the text: devices, themes, tropes—or genres and systems. And if, between the very small and the very large, the text itself disappears, well, it is one of those cases when one can justifiably say, Less is more.²³

Just as geographers can look at terrain and look at a map, and see the same geological phenomena described different ways, literary and legal scholars can look at texts from different perspectives. Represent Hamlet as a social network and you see that every character who speaks to both the Prince and the King dies.²⁴ Represent a genre as a tree, and you see how dozens of early detective fiction writers iteratively developed the concept of the clue—making way for Sherlock Holmes.²⁵ Represent a century's worth of country music on a graph and you see that the industry allows women artists a single hit.²⁶

²³ Franco Moretti, "Conjectures on World Literature" in Franco Moretti, *Distant Reading* (London: Verso, 2013) at 48-49.

²⁴ "Network Theory, Plot Analysis" in *ibid* at 211.

²⁵ "The Slaughterhouse of Literature" in *ibid* at 73.

²⁶ Juda Watson, "Gender on the Billboard Hot Country Songs Chart: 1996-2016" (2019) 42 *Popular Music and Society* 528.

Represent 23,000 applications for judicial review as a table and you see that an applicant's success turns on the identity of the judge assigned to the case.²⁷

Modeling culture, inspired by natural and social scientific practice of representing systems and data, unite a loose coalition of humanities scholars under the disciplinary heading of the digital humanities. Digital, because constructing models with large amounts of data requires the assistance of computational methodologies; humanities, to affirm that, yes, even though computers are about, this is humanities work. But this big tent disciplinary title is, nonetheless, misleading because it “tends to obscure the deep roots of this project in literary study itself”²⁸ by unduly highlighting the technological. For this reason, some scholars say that their work should be classified differently: quantitative formalism, the material study of form, or (my preferred label) sociological formalism.

Why these terms? First, formalism. Literary formalism overlaps with, but is much broader, than legal formalism. For most legal scholars, legal formalism invokes ideas about rules and categorical reasoning. Often pilloried as “mechanical jurisprudence,” legal formalism is often cast as overly idealistic and resistant to change. This is, at least as it describes more modern legal formalism, a mischaracterization. As Ernest Weinrib explains, legal formalism is interested in the “ensemble of concepts, principles, and processes that come into play when a legal claim is asserted.”²⁹ Put differently, when a person makes a claim based on breach of contract, certain legal forms (ideas, devices, genres) get invoked: ideas about, for example, unconscionability, consideration, and distress. An administrative law case, in contrast, will invoke different legal forms: reasonableness, purpose, fairness, and so on. Together, the concepts, idioms, rules, categories, approaches, and textual patterns come together to give a legal text its expression.

²⁷ Sean Rehaag, “Judicial Review of Refugee Determinations: The Luck of the Draw?” (2012) 38:1 Queen's LJ 1 at 52.

²⁸ Ted Underwood, *Distant Horizons: Digital evidence and literary change* (Chicago: The University of Chicago Press, 2019) at 164.

²⁹ Ernest Weinrib, “Legal Formalism” in Dennis M Patterson, ed, *A Companion to Philosophy of Law and Legal Theory*, 2nd ed (Chichester: Wiley-Blackwell, 2010) at 327.

Seen this way, there is a correspondence between legal formalism and literary or artistic formalist modes of analysis. As the art historian Amanda Wasielewski explains:

Form is the essence of a work of art, but it is also highly ambiguous. It is an artwork's visual and material properties or its compositional elements (i.e., its superficial appearance), but it also contains an artwork's effervescent quality, its expressive power. ... Form on its own, however, is merely a collection of interrelated features.³⁰ [emphasis added]

Formalism, then, is the study of these features (these forms) and their interaction with each other. A literary formalism, for example, might attend to the grammars, the imagery, the genre of a poem to study how it works as a literary cultural object. More than, but inclusive of, content, a study of form pays attention to how a text is constructed using the raw literary material available to the producer.

This points to an important distinction between the literary formalism under consideration here and legal formalism. For most legal formalists, there is a second object to formalist analysis: elaboration, coherence, and clarity. Once repeating legal forms are identified by the legal formalist, they often then ask whether the forms are principled, consistent, moral, and correct. This normative project distinguishes the critic of law from the critic of literature, whose interest is descriptive. The legal question: Are these invocations of legal forms principled and justified, with reference to morality and jurisprudence? The literary question: Why do these forms work together to express this?

Second, what does it mean to think of form materially or sociologically? For some literary scholars, especially scholars interested in the sociology of literature (that is, the study of the social relations that produce some types of literature over others), form is an index to the social world from which it emerges. For example, art produced in a feudal society for the benefit of the nobility will look, express, and qualitatively differ from that funded by, say, a modern national arts council. In this tradition, critics are less interested in the genius that produces great works of literature and art, but in the social conditions and facts of life that invited and allowed people to produce cultural work in certain ways.

³⁰ Amanda Wasielewski, *Computational Formalism: Art History and Machine Learning* (Cambridge: MIT Press, 2023) at 30.

Each piece of art, the theory goes, is an answer to some social problematic encountered by the artist and, viewed individually or in aggregate, art and literature are a perspective on those problematics.

This is not, I think, a difficult concept for lawyers to grasp and appreciate as relevant for their own field. Each day lawyers see how different types of claims—that is, different types of social conflicts—get grouped together and analyzed with reference to genre-specific rules and principles. If you were to study all contract law disputes from one decade to the next, you would be able to track how different social dynamics brought people to contract law, how contract law adapted to new social facts, and how legal forms rose, fell, or changed in response to social developments. Lawyers know that the law is not static. A sociological formalism invites us to think about how legal forms are related to social life.

Finally, quantitative. If each text is a collection of features (its forms) and if each form—genre, devices, rules, categories, word choice—tells us about some material facts of social life, one way these sociological facts can be monitored is through measurement of form. Once a form is isolated, we can describe its shape, growth, and decline numerically. Imagine two different types of refugee claim, each from the same country. In the first class of claim, applicants say that the state is actively persecuting them because of their political views. This claim will invite certain forms of assessment: does the person actually have these political views? Is the state actually seeking out people with these political views? In the second class of claim, applicants say that a local police force is targeting people of a particular ethnicity. This will invite other forms of assessment: are there parts of the country beyond the reach of the local police force where the claimant can relocate to? Having isolated these forms of inquiry, we can then measure them: how do different adjudicators answer these questions? How have answers developed together with time? And so on. With enough aggregate cases, average features will emerge that might be able to usefully model trends.

So, putting the pieces together: a quantitative sociological formalism. From legal formalism, an interest in the features that appear in a text when a legal claim is made. From sociology, a recognition that forms relate to social facts and provide a line of sight on how

legal ideas develop in response to social conflicts. From the natural and social sciences, a recognition that social phenomena can be quantitatively measured and modeled. This is how I say we might approach the great mass of unread legal texts. No, not in aid of a search for better doctrine or principled law, but to better understand how law operates as a social system. In the age of big data, examining legal texts in their aggregate might help us understand the location of individual influential texts, the stability or instability of legal forms, the rise and fall of legal genres, and the ways ideas about legality develop.

If my argument is correct, that this literary way of reading texts at scale can usefully support new legal research programs, it is useful to understand how and why literary scholars began to think this way. This is a methodological essay about the digital humanities that focuses, not on the technology available to us, but on why humanists reached for digital, model-building tools in the first place. In so doing, it points to a critical path forward for law and technology. For many law and technology scholars, new digital innovations promise to help us perfect law by using big data techniques to increase its coherence and predictability. This is a view from which I dissent. When we step back, when we look at law from a distance, we will not see anything that makes sense on its own terms: we will see a temporary and changing social system that is a product of its own moment.

My argument develops chronologically. I begin by exploring the moment when the discipline of literary studies elbowed its way into the university. In this early heady moment, new scholars of literature organized themselves around a shared method (close reading), a commitment to an area of study (the formal attributes of a text), and a general communitarian impulse (the idea that the study of art would improve and liberate the liberal student). As we will see, at different moments each of these touchstones was complicated, allowing for new visions of the literary to develop. In the 1960s, for example, literary studies' interaction with the counterculture, Marxism, the Frankfurt School, anti-colonialism, and psychoanalysis invited faculties to rethink the role of form and begin to think about how literary analyses could be enriched, not by more close readings, but with the lessons of the social sciences. More, cultural studies invited scholars to abandon the idea that there was such a distinct thing as the literary to recognize, instead, that all forms

of cultural production were amenable to literary analysis. Out with close readings of poems, in with analyses about why culture is the way it is.

Later, as the canon came under withering attack, scholars began to wonder why the canon was made the way it was in the first place. This sociological question invited scholars to think about all those non-canonical texts and the reasons for their exclusion from the canon. Increasingly, literary scholars began to draw on sociological modes of seeing to map the literary field and think through how institutions and competitions shaped literature. And so, the great unread came into view. Perhaps no scholar has done more to explore how the unread can be explored than the Marxist literary critic Franco Moretti. In the early and mid 2000s, Moretti published a series of interventions in the *New Left Review* sketching out the case for modelling literature using (though it took some time for him to accept this) computational methods. I therefore spend some time analyzing Moretti's work and then think about how his approach might usefully inform legal scholarship.

Throughout this essay, another figure looms large: Bourdieu. In the 1970s, 1980s, and 1990s, Bourdieu devoted a significant portion of his scholarly agenda to the sociological analysis of both art and law. Though his analyses of art and law were topically quite different, they shared a common empirical outlook. It is one that I have found inspiring but, more than that, useful because it allows me to see how ideas about literature might also be ideas about law, and *vice versa*. I do not propose to rehearse Bourdieu's sociology here. However, I do pause at several moments to use his work to further explain why legal scholars might look to literary methods and literary research programs.

One final point. This essay began with a list of unread cases. No, I do not suggest that we down tools, rush to the archive, and begin reading. But I do suggest that we make these unread cases count.

1.2 Close reading and the discovery of form

Literary studies is a young discipline. Before the First World War, the study of literature did not exist “in the sense of a distinct and widely legitimate academic discipline in which

one may pursue higher as well as lower degrees and make a career as a practising scholar.”³¹ What did exist was the study of language. Students interested in rhetoric, the relationship between morality and aesthetics, and understanding what distinguished good writing from bad, studied “Belles Lettres.”³² Students who wanted to understand the relationships between linguistic change and phonetics, myth, and ethnology (a forerunner to anthropology), studied the now dead discipline of philology.³³

Against these recognized scholarly disciplines, the study of literature was decidedly provincial. Scottish pupils, women, working men, and colonial subjects studied the great English writers in the lower-class educational institutions, peripheral colleges, or colonial schools. This literary study was, Terry Eagleton explains, “the poor man’s Classics — a way of providing a cheapish ‘liberal’ education for those beyond the charged circles of public school and Oxbridge.”³⁴ This is not to suggest that elite students did not read Shakespeare or Wordsworth, but it is to say that this reading was not thought of as scholarly. The exception demonstrates the point. On those rare occasions when literature was taught, it was not taught on literary terms. Each year, hundreds of Harvard undergraduate students would gather to listen to the eminent scholar G. L. Kittredge lecture on Shakespeare. But his lectures were not about poetry, values, or drama. Instead, as a philologist, he would lecture on “etymologies and meanings of words, allusions, and references.”³⁵

The first university students to study literature as literature were women. When, in the late nineteenth and early twentieth century, universities “grudgingly” admitted a few women students, literary pursuits seemed suitable because of literature’s supposed

³¹ John F English, “Literary Studies” in Tony Bennett & John Frow, eds, *The SAGE Handbook of Cultural Analysis* (London: SAGE Publications, 2008) 126 at 126. [English]

³² John Guillory, “Literary Study and the Modern System of the Disciplines” in Amanda Anderson & Joseph Valente, eds, *Disciplinary at the Fin de Siècle* (Princeton: Princeton University Press, 2021) at 19.

³³ Chris Baldick, “Literature in the Academy” in Patricia Waugh, ed. *Literary Theory and Criticism: An Oxford Guide* (Oxford: Oxford University Press, 2006) at 88.

³⁴ Terry Eagleton, *Literary Theory: an introduction*, (Minneapolis: University of Minnesota Press, 2008) at 23.

³⁵ William E Cain, “British and American New Criticism” in David Richter, ed, *A Companion to Literary Theory* (Hoboken: John Wiley & Sons, 2018) at 13.

“softening” and “humanizing” tendencies. Sending women to read English tracts then, in turn, protected the research university’s “virile” classical disciplines from the new students.³⁶ But this first, semi-democratic beachhead, expanded during and after the war. Philology was always associated with the Germanic scholarly tradition and as Anglo nationalism rose, so did the feeling that English and American universities ought to be interested in English and American traditions.³⁷

But market forces alone will not explain why a university suddenly establishes a literary department, staffed by professors teaching previously untaught texts. To become a discipline, proponents of literary studies needed to meet certain disciplinary prerequisites. This was especially so as universities transformed from places of teaching to places of research:

With the development in the nineteenth century of the new research university with specialised and professionalised academic disciplines, based on a general model provided by natural science, the main function of the university changed from being the custodian and transmitter of bodies of knowledge to being a producer of new knowledge. For the humanities, this meant that they had to ‘conform to a widely accepted set of criteria and conditions which distinguish disciplines of knowledge from bodies of knowledge.’³⁸

What were these criteria? A discipline needed to speak to a discrete and identifiable part of social reality, have a shared methodological commitment to studying that area of social reality, with a community of partisan practitioners who believed that study was worthwhile behind it.³⁹

The first step towards disciplinary status was taken by teachers who compiled lists of great literary works, evoking the “theological traditions of scriptural collation,

³⁶ *Ibid* at 24.

³⁷ *Ibid* at 26.

³⁸ Stein Haugom Olsen, “The Discipline of Literary Studies” in Andrea Selleri & Philip Gaydon, eds, *Literary Studies and the Philosophy of Literature: New Interdisciplinary Directions* (New York: Springer, 2016) at 38.

³⁹ *Ibid*.

evaluative winnowing, and hermeneutic scrutiny.”⁴⁰ With a canon, the first criterion was met: here was a discrete area of study. At the same time, early practitioners made a lofty promise about the returns study on this part of social life promised. Literary study was cast as a site for deep thinking about the human condition, because it was the study of an art form that provided the opportunity to think about what it would mean “to live from the vital centre of the most essential values.”⁴¹ A discrete area of study, then, that promised to help produce an engaged and thoughtful citizenry.

But the real breakthrough was, James English argues, methodological. To warrant disciplinary status, the study of literature had to be more than just reading. Instead, literary proponents needed an approach to text that distinguished their study from that which any literate person could do. To be sure, the innovations that came were not *sui generis*: literary did not invent its methodology, but cut from the cloth readily available to them. Just as the early students of literature were inspired by the study of religion when they set out to develop a canon, the notion of close reading was a familiar one to students of religion. The close study of scripture is a method that stretches back to antiquity, with different traditions developing different modes (and different reasons) for the study of religious text.⁴² Exegesis: the study of scripture to better understand one’s relationship with the sacred. From the study of the Talmud to the catechism, readers have paid close attention to the scripture to understand the world, religion, and their relationships with each other and with the divine. It is no coincidence that the Gutenberg press gave the world the Gutenberg Bible. Indeed, leading into the twentieth century scholars of the Old Testament began to use new critical intellectual tools to reinterpret religious texts, in ways that anticipated the study of literature to come.⁴³ What was novel was the conviction that

⁴⁰ Simone Murray, “Varieties of Digital Literary Studies: Micro, Macro, Meso” (2022) 16:2 DHQ at para 23.

⁴¹ Eagleton, *supra* note 34 at 27.

⁴² James L. Kugel, “The emergence of biblical interpretation in antiquity” in Mordechai Cohen and Adele Berlin (eds) *Interpreting Scriptures in Judaism, Christianity, and Islam: Overlapping inquiries* (Cambridge: Cambridge University Press, 2016).

⁴³ John Rogerson, *Old Testament Criticism in the Nineteenth Century: England and Germany* (Philadelphia: Fortress, 1985).

literature, those texts just lying around, were worthy of the sort of close interpretive study that had once been reserved for religious texts.

Here, the first wave of key innovators came from Moscow. The Russian formalists, a group of early literary scholars, worked to substitute a “hard science” for the emotive or mythic analyses that had, up to that point, dominated criticism of Russian literature.⁴⁴ Their investigations focused on the rhythm, style, and structures of a text as opposed to, for example, their moral or emotive impacts. This extremely close attention to the text (Russian formalists might, for example, count the beats in a stanza) seemed to resemble philology for its focus on language, but there was a key difference. Where the philologists sought to understand how language developed scientifically, the Russian formalists wanted to understand how great artists used and manipulated text to make artistic masterpieces. What, they wondered, distinguished normal text from “estranged” literary text?

Philology was interested in the words and grammar, all the raw material; formalism was interested in the artistic uses of that material, the uses whose deviance from ordinary language practices might be said to define the ‘art’ or the ‘poetic function’ of a text, its special qualities or properties as literature.⁴⁵

Their literary movement was short-lived. Though they earned Trotsky’s grudging respect (the formalists, he said, “may open a path—one of the paths—to the artist’s feeling for the world.”⁴⁶), he developed an early critique of their work: too esoteric, too disconnected from the raw experience of working-class life. This intellectual critique assumed, in Stalin’s Russia, a coercive form of repression during which, by the late 1920s and early 1930s, formalisms of all sorts (in literary, artistic, and musical circles) became targets of the purges.⁴⁷

⁴⁴ James F English, “Literary Studies” in Tony Bennett & John Frow, eds, *The SAGE Handbook of Cultural Analysis* (London: SAGE Publications, 2008) 126 at 128.

⁴⁵ *Ibid.*

⁴⁶ Leon Trotsky, *Literature and Revolution* (Chicago: Haymarket Books, 2005) at 139.

⁴⁷ David Gorman, “Russian Formalism” in Richter, David, ed, *A Companion to Literary Theory* (Hoboken: John Wiley & Sons, 2018) 36 at 36.

But the movement influenced and heralded similar ways of thinking in Britain and America where the parallel schools of ‘practical criticism’ and ‘new criticism’ each set out to set the study of literature in more precise and scientific terms. Like the Russian formalists, and in tandem with efforts to develop a canon, they set out to “defend the specificity of the literary object” and its “unique and irreducible value as art.”⁴⁸ Literature, they thought, ought to be distinguished from the normal cultural artefacts of everyday life.⁴⁹ To explore how literary objects functioned and were distinct types of cultural productions, they developed a new methodological concept: the idea of close reading.⁵⁰ This method, neither reading itself nor the type of attention a linguist might pay to words, asked students to attend to the ways forms and features came together in the text. To read closely meant to focus on a text’s:

constitutive devices, patterns, or elements of style, including such effects as rhythmic regularity or irregularity, such tropes as irony or apostrophe, such narrative techniques as free indirect style or stream of consciousness, such thematic features as doppelgangers or Manichean binaries, and such structuring devices as foreshadowing, misdirection or flashback. From the formalist vantage, literariness consisted of problems or complications, ‘impediments’,... to a reader’s rapid and virtually automatic comprehension. Close reading was a way of highlighting these textual intersections or roadblocks where the mind is forced to slow down and scout for detours.⁵¹

Armed with a coherent object of study, a unique way of seeing it, and a group of committed practitioners, literary studies became a recognized form of study. And a popular one at that. If once it was, Eagleton says, “desperately unclear why English was worth studying at all; by the early 1920s it had become a question of why it was worth wasting your time on anything else.”⁵²

⁴⁸ James F English, “Literary Studies” in Tony Bennett & John Frow, eds, *The SAGE Handbook of Cultural Analysis* (London: SAGE Publications, 2008) 126 at 129.

⁴⁹ *Ibid.*

⁵⁰ *Ibid.*

⁵¹ *Ibid.*

⁵² Eagleton, *supra* note 34 at 27.

Methodological decisions are consequential. For the first generation of literary scholars, their understanding of what close reading was (why is this word here? Why this meter? Why this rhythm?) led them to study short poems. It was difficult to sustain the energy and momentum of close reading for more than a page or two, making close analyses of novels or plays functionally impossible. Moreover, if a lecture was only two or three hours long, poetry was an effective form to present to students the method of close reading: everyone, together, could see and explore the same text. And scholars worried about what would happen to the discipline if they tolerated methodological departures from these close readings. This sentiment found clear expression, for example, in Cleanth Brooks' famous essay "The Heresy of Paraphrase." For Brooks, it was a mistake to work at any level abstracted from the immediacy of the words on the page because the critic would lose access to the text's literary qualities. Yes, he says, the lines "So wore night" and "Thus night passed" share a common denominator of meaning but to say that both lines simply concern the passage of time at night, would prevent the critic from exploring why different artists approached the concept differently. Form and content, he thought, are indivisible and each construction, each small decision, is a matter of artistic significance.⁵³

By the 1950s, new analytic interventions, however, made it possible to practically analyze different textual mediums. Northrop Frye, for example, pointed out that genre—something that manifested with macro, not micro, analyses of text—was a form that could be studied. For Frye, literature was not "just a random collection of writings strewn throughout history,"⁵⁴ it was a self-contained system of texts that drew from a shared pool of myths, metaphors, narratives, and symbols. This suggestion opened new opportunities for analysis by inviting scholars not to just closely read the forms as they manifested line-by-line, but to analyze how larger, mid-level, or text-level, forms expressed meaning.

Put together, the first decades of literary studies' life represent a startling and exciting achievement. From a place where literary texts were not studied, it moved to one

⁵³ Cleanth Brooks, *The Well Wrought Urn: Studies in the structure of poetry* (New York: Harcourt, Brace, and World, 1947) at 201.

⁵⁴ Eagleton *supra* note 34 at 79.

where a large and growing trans-Atlantic group of scholars drew the interest of a large body of students. Indeed, sacred texts themselves became the subject of literary study. But though they had proven capable of methodological innovation, their shared commitments to studying only the literary object (the canon) and to the idea that the literary was an autonomous sphere of artistic genius, meant that questions about art's relationship, and location within, social life were "bracketed out."⁵⁵ With the 1960s just around the corner, this soon would change.

1.3 Literature (and culture) as a symptom of material life

A centre of the counterculture was the university. By the late 1960s, campuses the world over were in revolt. Students, protesting their professors' conservatism, racism, the persistence of colonial apparatuses, patriarchy, and imperial wars, looked for inspiration from both new thinkers and old neglected ones. It was just as much the age of Angela Davis and Frantz Fanon as it was of Karl Marx and Sigmund Freud.

No discipline escaped contact with the spirit of the season and no discipline left the decade fundamentally the same. Literary studies, for example, was particularly touched by Marxism. Against an older generation's conviction that art was an autonomous zone of creativity, students who saw literary traditions as vectors of a particular antiquarian way of thinking were inspired by Marx's suggestion that cultural thought was a product of unequal social relations. Why do people think the way they do? Because the economic base, the class organization of society, predicts the form of its cultural, superstructural organization:

In the social production of their existence, men inevitably enter into definite relations, which are independent of their will, namely relations of production appropriate to a given stage in the development of their material forces of production. The totality of these relations of production constitutes the economic structure of society, the real foundation, on which

⁵⁵ *Ibid* at 130.

arises a legal and political superstructure and to which correspond definite forms of social consciousness. The mode of production of material life conditions the general process of social, political and intellectual life. It is not the consciousness of men that determines their existence, but their social existence that determines their consciousness.⁵⁶

A famous formulation, but one prone to vulgar generalization. What distinguishes, for example, economic zones from cultural zones? Is thought totally determined by the class composition and arrangement of society? Are superstructures just mildly autonomous zones that, in the final analysis, do nothing other than justify and support the dominant as they dominate the dominated? No, Marx himself said that people had a wide ambit to develop their own interpretations of social life and own visions of, for example, art. To paraphrase the key lines from *The Eighteenth Brumaire of Louis Bonaparte*: authors produce literature, but they do not do so as they please, but from circumstances existing already, given and transmitted from the past. Social and material facts come first, posing questions that can be answered in any number of different cultural ways.

A more nuanced and complicated Marxist analysis of culture had, since the 1930s, developed in Europe.⁵⁷ The Stalinist purges largely cut short innovative Soviet cultural scholarship, but elsewhere the Frankfurt school, Lukács, Marcuse, and Gramsci had all explored the complicated ways that economic and cultural life were integrated together. Walter Benjamin was an early innovator. His idea of correspondences said that there was a relationship between how people occupied and moved through space and the patterns of their thought. As the geography and architecture of Paris changed—as arcades were built and Haussmann rearranged the city—corresponding new forms and conceptions about the natures of the crowd developed, which in turn allowed for new conceptions of the individual to emerge. This historical moment, Benjamin says, is what provoked

⁵⁶ “Preface” in Karl Marx, *A Contribution to the Critique of Political Economy* (Progress Publishers: Moscow, 1977) online: < <https://www.marxists.org/archive/marx/works/1859/critique-pol-economy/preface.htm>>.

⁵⁷ James F English, “Literary Studies” in Tony Bennett & John Frow, eds, *The SAGE Handbook of Cultural Analysis* (London: SAGE Publications, 2008) 126 at 137.

Baudelaire's new forms of poetry (a correspondence) because new conceptions of individuality allowed for new poetics.⁵⁸

To these Marxist critics, formalism itself was evidence of a shift in social relations and realities. Medvedev and Bakhtin argued in *The Formal Method in Literary Scholarship: a critical introduction to sociological poetics*, formalism was related to futurism because, as a method, it combined "extreme modernism and radical negation of the past."⁵⁹ Insisting that art was an autonomous zone of production, decipherable by close study that ignored social context, at once rejected the idea that art was a mere reflection of bourgeois morality but also "those forms of thinking about art which connected it with that or indeed any other social order."⁶⁰ This was an advantage for showing that culture could not be deciphered with reference to an economic or class base alone, but it denied the formalist access to the conditions (the correspondences) that produced ways of thinking. For this reason, they said that formalism needed to be celebrated for teaching critics how to pay a new type of detailed attention to the text, but they also urged critics to ground their analyses of poetics in a sociological analysis.

This tradition was largely inaccessible to Anglo scholars in the post-war period, both because there were few Marxist scholars working in literary departments and because the work had never been translated into English.⁶¹ This changed in the 1960s and 1970s as Marxist scholars began to enter official literary departments and old texts were introduced to the Anglo reader for the first time. Writing without an immediate set of English-language scholars to draw on, young Marxist literary scholars embarked on an ambitious project to craft a sweeping, sociologically minded, approach to cultural texts. No scholar better represents this critical season better than Fredric Jameson, one of the great American literary critics. It is telling that Jameson's second book, *Marxism and Form*, sought to combine traditional formalism, psychoanalytic thought, and old Western

⁵⁸ See Raymond Williams, "Marxism, Structuralism and Literary Analysis" (1981) 129 *New Left Review* 51 at 58.

⁵⁹ Pavel Medvedev & Bakhtin Mikhail, *The Formal Method in Literary Scholarship: a critical introduction to sociological poetics* (Cambridge: Harvard University Press, 1985).

⁶⁰ Raymond Williams, "The Uses of Cultural Theory" (1986) 158 *New Left Review* 19 at 23.

⁶¹ *Ibid* at 137.

Marxism texts. Following that older tradition, he said, close reading, a “precise analytical instrument,” presumed that all literary texts could be examined the same way and on the same terms, meaning that the same sort of analysis always was “rigidly applied to every kind of text.”⁶² This denied the critic access to an important truth: that all cultural works:

come to us as signs in all-but-forgotten code, as symptoms of diseases no longer even recognized as such, as fragments of a totality we have long since lost the organs to see.⁶³

Symptomatic analysis: the idea that art and its forms are the “final articulation of the deeper logic of the content itself.”⁶⁴ For Jameson, analyzing the literary was a method of sociological analysis. Study the art of a moment, and you gain access to a moment’s deep anxieties and troubles.

This theoretical approach authorized new readings of literary texts. Nancy Armstrong, writing about the domestic novel, showed how this genre of fiction, helped create the idea of a feminized private space and inculcate new ideas of subjectivity and domesticity.⁶⁵ Franco Moretti, writing on the coming-of-age novel (the ‘bildungsroman’), argued that the form, this genre of literature, helped the emerging bourgeois class reconcile its commitment to individuality with its belief in the value of social order.⁶⁶ Jameson showed that the *Heart of Darkness* expressed the contradictions and anxieties associated with imperialism.⁶⁷

But perhaps more importantly, to say that literary texts were a product and symptom of their historical circumstances also implied that any cultural text was equally related to its moment. The founders of cultural studies, Raymond Williams and Stuart Hall, for example, were proponents of a method of social formalism. Every text, not just

⁶² Fredric Jameson, *Marxism and Form: Twentieth-century dialectical theories of literature* (Princeton: Princeton University Press, 1972) at 93.

⁶³ *Ibid* at 419.

⁶⁴ *Ibid* at 329.

⁶⁵ Nancy Armstrong, *Desire and Domestic Fiction: A Political History of the Novel* (Oxford: Oxford University Press, 1987).

⁶⁶ Franco Moretti, *The Way of the World: The Bildungsroman in European Culture* (London: Verso: 2000).

⁶⁷ Fredric Jameson, *The Political Unconscious: Narrative as a socially symbolic act* (Ithaca: Cornell University Press, 1981).

great literary texts, was for them an index to the time and moment of its production. In Williams' *Keywords*, for example, he traced how the meaning of specific words developed and changed through time to adapt to changing social conditions and conflicts. The immediate effect was to dramatically widen the field of study available to literary scholars. Where once students of literature were restricted to the study of canonical works, the turn to culture authorized the study of the full range of cultural productions, from music, to newspapers, to poetry, to advertisements.

Equally, it invited scholars to attend less to the immediacy of a specific cultural production and more to its position within a cultural field and within long temporalities. Against the day-to-day events that invite people to express themselves, it is better to see each text as "little more than an individual *parole* or utterance" that combines to form collective discourses.⁶⁸ For related reasons, Williams thought it best to conceive of cultural studies on three levels: that of lived culture, that of recorded culture, and that of selected (or canonical) culture. Everyone who lives in a moment, he thought, will inevitably have access to cultural features that will escape the notice of any commentator, be they a contemporary sociologist or a historian. And canonical culture will inevitably be a partial and partisan representation of reality. For this reason, he thought that literary and cultural critics should attend primarily to recorded culture as best they could. But, in a comment that anticipates what would come, he acknowledged how difficult this would be: "One can say with confidence, for example, that nobody really knows the nineteenth-century novel; nobody has read, or could have read, all its examples, over the range from printed volumes to penny serials."⁶⁹

This sort of critical spirit was infectious. Soon, a *Law and Literature* movement would emerge that, amongst its several agendas, included a sub-field of study: law as literature. But even before then, legal scholars were examining how social life spoke to the form of law. In his famous essay on the rise of rule-based patterns of legal reasoning, Morton Horowitz asked how and why in the 1850s new and unyielding legal rules and

⁶⁸ Michael Hardt & Kathi Weeks, eds, *The Jameson Reader* (Oxford: Blackwell, 2000) at 37.

⁶⁹ Raymond Williams, *The Long Revolution* (Peterborough: Broadview Press, 2001) at 66.

categories replaced 18th century legal forms that were “protective, regulative, paternalistic and, above all, a paramount expression of the moral sense of the community.” His answer: as American society oriented increasingly towards the market, law shifted to protect, over the worker and the farmer, the interests of the business owner and man of commerce. And,

once successful, those groups could only benefit if both the recent origins and the foundations in policy and group self-interest of all newly established legal doctrines could be disguised. There were, in short, major advantages in creating an intellectual system which gave common law rules the appearance of being self-contained, apolitical, and inexorable...⁷⁰

This analysis, which owed as much of its origins to social and intellectual history, nonetheless spoke to the tempo of the critical season: ideas and principles, even those that seemed value-neutral, were symptomatic of social realities.

Leaving the 1980s, then, the literary’s encounter with the spirit of the 1960s left the discipline changed in two important ways and, perhaps ironically, strengthened. Now, the methods of formal analysis promised not only to disclose the ways in which art functioned, but to help explain the social realities from which they emerged. Second, where once literary departments studied a narrow range of texts, now their methods were understood as valid to explore a range of subjects previously thought to be within the exclusive remit of the social sciences.

1.4 The sociology of literature and Bourdieu

From the micro analyses of text at formalism’s beginnings, to its macro sweep after its encounter with countercultural politics. Somehow lost in between, as literary scholars suddenly traversed from autonomous readings of art to questions about society’s fundamental and basic organization, were all the smaller institutions, tendencies, competitions, and traditions that influence the shape of culture. Put differently, the study of literature is haunted by the specters of two related errors: an excessive belief in either automatism or functionalism. In one account, art is pure because it is entirely disconnected

⁷⁰ Morton J Horwitz, “The Rise of Legal Formalism” (1975) 19:4 Am J Leg His 251 at 251–252.

from social life (subjectivism); in the other, art is pure because it is the functional expression of some class or group interest (objectivism).

At least, these were risks identified by French scholar Pierre Bourdieu. To Bourdieu, literary scholarship paid insufficient attention to the social universe of literature that mediated between social facts and the thought of cultural producers. To put it simply, books are not published because they are symptomatic of social relations, they are published because there is a market for them, there is a publisher who believes a particular work by that author will meet market demand, there is an author who wishes to write for that publisher, and because they are asked to write over another writer who might have been just as willing. Bourdieu proposed that each area of social life, or “field,” is a site of competition among individuals and institutions. In these fields position and success are dictated by the different types of capital (economic, cultural, social, and symbolic) and how they interact with other capital-bearing entities.

As these competitions unfold, they give rise to field-specific grammars, behaviors, and expectations, forming what Bourdieu refers to as *habitus*. This habitus, a congealed expression of history, produces a set of dispositions that influence how individuals perceive and act within the field. Consequently, a cultural producer produces within a complex web of aspirations, expectations, and conflicts.⁷¹ Against Marxism, Bourdieu insists on the vibrancy and independence of a literary field where field-specific positions are taken and field-specific competitions occur; against formalists (or, in his time, against Foucault), Bourdieu insists that scholars take note of how that literary field is, in turn, situated within a wider field of power that is itself riven by competitions between classes and over material resources.

This sociological approach to the study of culture appealed, in the 1980s and 1990s, to many literary scholars, especially those who were interested in how canonical works

⁷¹ Richard Harker, Cheleen Mahar & Chris Wilkes, “The Basic Theoretical Position” in Richard Harker, Cheleen Mahar & Chris Wilkes, eds, *An Introduction to the Work of Pierre Bourdieu: The Practice of Theory* (New York: Springer, 2016).

became canonical.⁷² Scholars of the sociology of literature, who sought to understand why some forms became ascendant over others and why others later fell, used Bourdieu to sketch out empirically informed analyses of the changing literary field. How, they wondered, did the competitions between funders, editors, writers, markets, and governments play out so that different types of literature sometimes rose or fell?

Critics here consider how different fields interact, such as a literary field for the mass market and another for the avant-garde, how meso-level institutions (publishers, funding agencies, professional editors) influence literary production, and the role of the different reading publics. This sort of sociology of literature depended upon access to new data. For these scholars, the canon was not so much an ethical or political problem in and of itself, but one that needed to be understood: how did these texts, against all others, become significant? To answer these questions, they needed evidence to reconstruct the literary field, to understand its actors, and identify the sorts of competitions they engaged in, and then reinsert the production back into the field to make sense of it.

It was exactly these sorts of problems that motivated Margaret Cohen as she wrote *The Sentimental Education of the Novel*. In this study, she sought to understand why, in the 1830s, the novel (as a form of art) “became an authoritative form of social and cultural analysis imbued with the highest literary prestige.”⁷³ Before, novels were pleasant entertainment; after, it was the mode *par excellence* of exploring issues related to social class and conflict. Literary history has singled out two realist authors, Balzac and Stendhal, as the leading innovators of “brilliant poetics” who engaged in a “heroic struggle to make sense of the economic, social, and political upheavals resulting from the French Revolution.”⁷⁴ But in Cohen’s view, a visit to the archive disturbs this view. Instead of seeing the works of Balzac and Stendhal as the pure products of genius, she shows how—when held against the forgotten texts of the “great unread”⁷⁵—they developed their work

⁷² James F English, “Everywhere and Nowhere: The Sociology of Literature After ‘the Sociology of Literature’” (2010) 41:2 *New Literary History* v at ix.

⁷³ Margaret Cohen, *The Sentimental Education of the Novel* (Princeton: Princeton University Press, 2002) at 1.

⁷⁴ *Ibid.*

⁷⁵ *Ibid* at 23.

in competition with, and against, a genre of sentimental literature written by women. Their efforts, in Bourdieu's sense, represented a "hostile takeover" of the literary field.⁷⁶ This helps explain the absence of women in our memory of French novels: it is not that they did not write, it is that their works fell into obscurity after the literary field to accommodate new genres.

But this study highlights the real practical methodological challenges of this approach. Cohen could read the unread only because, studying the early stages of novelistic production in France, the corpus was relatively small. As the number of texts, and the amount of associated metadata that helps us contextualize each text, grows, the reading becomes impossible. As scholars shifted their interests and began to ask more about the sociology undergirding literary production and literary value, the final original pillar of literary studies would be threatened. First the discrete and irreducible value of the literary object as art, then the canon, then the narrow field of study, soon the method of reading.

1.5 Modelling text and the digital humanities

No scholar has contributed more to the origins of the digital humanities than Franco Moretti, the Italian Marxist scholar who co-founded the Stanford Literary Lab. In a series of essays in the *New Left Review*, Moretti first described what would become the main research agenda of the digital humanities. But what is important to note for us now is that many of these essays were decidedly not about digital methods. Instead, in the early 2000s, Moretti sought to sketch out a new approach for the study of literature, responsive to the questions raised by sociologists of literature, that could understand how and why literary forms evolved and changed by drawing explicitly on the methods of social scientists.

This interest can first be clearly detected in his *Atlas of the European Novel*. In this text, Moretti aimed to make explicit the connection between geography and literature by representing, in a series of maps, how, first, space figured in literature and, second, of how

⁷⁶ *Ibid* at 6.

literature functioned in space. In the former, how do characters move through fictional space? In the latter, how do the points where literature is printed, shared, and sold impact the rise and fall of novelistic forms? His object was to use maps as analytical tools that would bring “to light relations that otherwise remain hidden.”⁷⁷ Midway through his text, Moretti paused for a “theoretical interlude.” Spending so much time in the libraries and mapping all these unknown novels highlighted how unexceptional most literature was. Not that it was unimportant, but from the perspective of the modern reader, all this text just looked so mundane and repetitive:

But then, are we so sure that boredom is boring? Once we learn to confront it, the flatness of literary conventions will appear for the genuine enigma it is. How does a new narrative form crystallize out of a collection of haphazard, half-baked, often horrendous attempts? How does a convention change, or, better: *does* it ever change? Or does it remain stable in a thousand disguised — until the day it suddenly disintegrates? And why does it remain stable so long? And why does it then collapse? And how on earth can the *same* convention work in such *different* places — Scotland and Italy, Denmark and Hungary?⁷⁸

But these, he said, were “[a]ll questions for another study.”⁷⁹

In 2000, one year after the publication of this *Atlas*, in the *New Left Review* Moretti published *Conjectures on World Literature*. The study of comparative literature, he said, was largely a disappointment because scholars usually just reverted to type and closely studied the texts of Western European countries. Yet, especially in the new millennium, it was clear that literature was (and always had been) a planetary system. While it was possible to speak of distinct national literatures, it was an error not to notice how different national literatures influences and drew from each other. After all, the novel as a form began in Western Europe in the nineteenth century, but by the twentieth century authors the world over were writing novels. How to account for the ways in which the forms in the great mass of literature, and here he cited Cohen’s for the expression “the great unread,” evolved, changed, and faded?

⁷⁷ Franco Moretti, *Atlas of the European Novel, 1800-1900* (London: Verso, 1999) at 3.

⁷⁸ *Ibid* at 150.

⁷⁹ *Ibid*.

Here he proposed the concept of distant reading: the attention to common features of texts in aggregate. But Moretti had not yet conceived that digital tools could be used in service of this type of analysis. Instead, in this early intervention his suggestion was that critics ought to attend more to the work of their peers, to aim to develop global level syntheses of secondary work on local literatures. At first, this essay was neither intended nor received as a call to embrace computational methodologies. Instead, “it seemed to be the culmination of a long argument over the canon” with Moretti’s detractors hearing “a forceful claim that critics’ bookshelves had to be replaced by the whole library.”⁸⁰ This initially provoked a sharp response from scholars who saw Moretti’s intervention as, not just another critique of the canon, but a call to abolish the canon altogether. Meanwhile critics interested in subjugated literatures worried that their efforts to develop new counter-canons would be overwhelmed by the scale of a world-literature analysis.⁸¹

But as the canon wars faded, so too did the initial controversy around *Conjectures*. Instead, it and Moretti’s related work were understood as calls for “maximum methodological boldness” for literary studies.⁸² Still, he remained committed to the study of form partly because “sociological formalism has always been my interpretive method.”⁸³ Continuing the work he began with the *Atlas*, Moretti began to publish essays exploring how other models from the social sciences could be used to model literature. Using graphs from quantitative history, he showed how the number of novels published each year in different literary markets followed the same patterns of take-off, plateau, and explosive growth. Using maps from geography, he showed how the physical space was arranged differently in novels, depending on whether they described a place before or after industrialization. Using trees from evolutionary biology, he explored how Darwinian

⁸⁰ Ted Underwood, *Distant Horizons: Digital evidence and literary change* (Chicago: The University of Chicago Press, 2019) at 174–175.

⁸¹ *Ibid.*

⁸² Franco Moretti, “The Slaughterhouse of World Literature” in Franco Moretti, *Distant Reading* (London: Verso, 2013) at 89.

⁸³ Franco Moretti, “Conjectures on World Literature” in Franco Moretti, *Distant Reading* (London: Verso, 2013) at 49

accounts of evolution could describe the increasing diversity of literary forms. His studies focused, he said, on “Shapes, relations, structures. Forms. Models.”⁸⁴

Moretti’s objective was to find common features in and around texts and find ways to isolate and present that commonality. Models are therefore a grossly simplified version of reality. In place of the rich complexity of each text, a model works by stripping from each unique instance those features that make it unique and distinguish it from each other instance. They are “simulations, a combination of technique and technology that aims to represent the world in some way but also reduce its complexity.”⁸⁵ Or, as Andrew Potter put it, attention to quantity “reveals the grooves and channels of cultural expression, the deep connections among words, ideas, and forms.”⁸⁶

And as the years passed, technology’s role in producing these models became increasingly evident. Soon, scholars like Moretti would develop a disciplinary identity: the digital humanities. Two events crystallized the movement. First, in 2004, Blackwell published an edited collection, the *Companion to Digital Humanities*. Originally titled the *Companion to Humanities Computing*, the editors wanted a new title to signal that their vision for the future of computational humanities work was more than “mere digitization.” Up to that point, the bulk of computational humanities work was focused on the work of making computationally-accessible versions of humanist and archaeological artefacts or—once again echoing the religious roots of close reading—in large computational studies of scripture.⁸⁷ But, by the time of the publication, the editors sensed that the discipline was reaching “a turning point” and, “for the first time, a wide range of theorists and practitioners” could be brought together to discuss what the field was and might become. In 2005, following this logic, the Association for Computers in the Humanities and the Association for Literary and Linguistic Computing merged into a new

⁸⁴ Franco Moretti, *Graphs, Maps, Trees: Abstract models for literary history* (London: Verso, 2007) at 1.

⁸⁵ Andrew Piper, *Can We Be Wrong? The Problem of Textual Evidence in a Time of Data*, (Cambridge: Cambridge University Press, 2020) at 18.

⁸⁶ Andrew Piper, *Enumerations* (Chicago: University of Chicago Press, 2019) at 3.

⁸⁷ Here Robert Busa’s work is emblematic, see Dolores Burton, “Review of Index Thomisticus: *Sancti Thomae Aquinatis operum omnium indices et concordantiae; Sancti Thomae Aquinatis opera omnia*” (1984) 18:2 *Computers and the Humanities* 109.

organization, the Alliance of Digital Humanities Organizations. Second, the digital humanities was “big news at the 2009 MLA Annual Convention.” Heralded as the next big thing, panels featuring digital humanists that year were plagued by “overflow crowds to too-small conference rooms.”⁸⁸ Soon, research funding dollars began to flow and digital humanities labs started to pop up all over North America.

This interest gave Moretti’s *Conjectures* a second life, becoming both a manifesto for the new discipline and a target for “even more violent [critique]—this time, equanimously from the left and the right—aimed at the idea of ‘distant reading.’”⁸⁹ As summarized by Martin Eve:

Sapping all of the conventional funding out of traditional humanistic pursuits, the so-called ‘digital humanities’ (or ‘DH’ to those in the know) brings a grim entrepreneurialism and technocratic mindset to English, history, classics, archaeology—and any other disciplinary space on which it can lay its hands. Seemingly charged with perverting the humanistic foundations of critical thinking and replacing them with techno-solutionist mindsets, the digital humanities are growing and thriving beneath our noses and many seem not even to have noticed the danger.⁹⁰

But the new discipline continued its growth. Soon scholars were publishing in the *Journal of Cultural Analytics*, the *Digital Humanities Quarterly*, the *New Left Review*, *Digital Scholarship in the Humanities*, and *Digital Studies*. As both industry and the academy became increasingly interested in uncovering methods to analyze text at scale, a new study of literature “developed tools, models, and software that facilitated the analysis and organization of texts at scale.”⁹¹

The discipline is now too large to be simply described: there are as many diverse research agendas as there are scholars. That said, Bourdieu’s presence remains

⁸⁸ Matthew G Kirschenbaum, “What Is Digital Humanities and What’s It Doing in English Departments?” (2010) 150 *ADE Bulletin*.

⁸⁹ Franco Moretti, “Conjectures on World Literature” in Franco Moretti, *Distant Reading* (London: Verso, 2013) at 44.

⁹⁰ Martin Paul Eve, *The Digital Humanities and Literary Studies* (Oxford: Oxford University Press, 2022) at 1.

⁹¹ Justin Grimmer, Margaret E Roberts & Brandon M Stewart, *Text as Data: A new framework for machine learning and the social sciences* (Princeton: Princeton University Press, 2022) at 1.

undeniable. In a retrospective piece, Moretti noted how Bourdieu (who “stands for a literary study that is empirical and sociological at once”⁹²) was a major presence in digital humanities work. Or, as Simone Murray observes, for sociologically minded digital humanists their goal is to harness “the enumerative power and analytical scale of digital data to ask new questions about how Bourdieu’s literary field functions.”⁹³ As scholars collect and release larger and larger datasets, a quantity of empirical evidence about the literary field, which would have been unimaginable to Bourdieu, comes into view. Through the analysis of texts, digital humanists can observe how literary forms develop over years and decades. Outside of texts, it is possible to track the identity of publishers, editors, agents, and writers. Digital humanists therefore can track deep changes in the habitus of literary agents by monitoring text at scale and text over time.

1.6 Reading legal text at scale

Concurrently, some legal scholars began to explore how the same technological tools used by the digital humanists could be usefully used to study law. “Law as data” scholarship hypothesizes, like digital humanists, that legal texts are amenable to computational analysis. Work here has not yet been organized institutionally, at least to the same degree as it has been in the digital humanities. There are not yet, for example, any computational legal studies journals. This is no doubt in part because there are so few legal researchers who have the skillsets required to undertake computational projects, but it is also a feature of the fact that, until quite recently, researchers had only limited access to large datasets of law. A recent survey of computational legal research found only thirty-nine studies based on textual legal data published since 2010 and found, in general, that a “majority of existing studies using computational methods... are concerned with shedding light on factors influencing the legal process itself.” Put differently, to date the trend in

⁹² Franco Moretti, *Literature, Measured*, Pamphlets of the Stanford Literary Lab (Stanford, 2016) at 7.

⁹³ Simone Murray, “Varieties of Digital Literary Studies: Micro, Macro, Meso” (2022) 16:2 DHQ at para 17.

computational legal research is to use methods to pull out variables of interest from text for empirical analysis.

There is, to be sure, considerable overlap between the digital humanities and computational legal work, but each field has arrived at computational methods from a different direction. For legal scholars, computational tools usually promise to continue and advance the tradition of empirical legal work that stretches back, through the realists, to Oliver Wendell Holmes Jr. Computers promise to help scholars figure out how outcomes are related to facts and to cut through the mystique of doctrine. The text is primarily a source of data, and less an interesting artefact in and of itself. This, of course, contrasts with the literary scholar's preoccupation: the text, its shape, and the reasons for its production. Perhaps there is a simple reason for this: unlike literary studies, and its modern interest in the sociology of literature, there is not an analogous pre-existing sub-field of legal study interested in understanding, in a systemic way, the relationship between law's canonical read and its ignored unread. Law and Society scholarship, that area of legal research most touched by sociological research methods, for example, is largely uninterested in Bourdieu-type empiricism, preferring to focus on different sets of questions.⁹⁴

This strikes me as somewhat odd because, from the perspective of most practicing lawyers, empirical accounts of law make much more intuitive sense than doctrinal accounts. Lawyers see that each case, each legal text, is a response to a felt social experience. Reasons justifying or denying an award of social assistance speak to a level of want. Negligence claims speak to injury. Refugee claims to a desire to move. Yes, each experience is mediated through the legal system, but the shape of text speaks to the social positions of all the impacted actors. And, in aggregate, cases speak to common social relations and common social disputes.

In tribunal rooms around the world, lawyers do a bit of spontaneous sociology every time the adjudicator walks into the room: given this decision maker, given this case,

⁹⁴ Yves Dezalay & Mikael Rask Madsen, "The Force of Law and Lawyers: Pierre Bourdieu and the Reflexive Sociology of Law" (2012) 8:1 *Annu Rev Law Soc Sci* 433.

given what I understand about the law, here is what I will say. Compared to the author, the lawyer is constrained. Law, Bourdieu explains, “achieves its effectiveness at the cost of a limitation in its autonomy.” Though different ideas regarding a legal text may circulate, all legal actors operate “within a body capable of resolving conflicts between interpreters and interpretations.”⁹⁵ As these conflicts resolve, the law obtains “the appearance of a transcendental basis”⁹⁶ that applies universally. But for all the stability, law remains dynamic. As social conflicts develop or change, a “tension [develops] between the available judicial norms... and the necessarily diverse, even conflicting and contradictory, social demand.”⁹⁷ To understand law’s functions and operations, it becomes necessary to “recover the profound logic of juridical work in its most specific locus, in the activity of formalization.”⁹⁸

And here is the promise of the methods of the digital humanist: instead of mining a text for variables, watch its forms, watch how the deep structures of law change and adapt. In practice, what does this look like? In *Bankspeak*, Moretti and his co-author, Dominique Pestre, computationally explore the annual reports of a major international legal institution: the World Bank. Their analysis shows a clear break between the styles of writing used by the Bank before and after the 1970s. In the 1940s, 1950s, and 1960s, the nouns of the reports often referred to infrastructure (dams, irrigation, steel, highways) and the actors were states, industry, and experts.

After, the language shifted to focus on finance. Key adjectives now described returns and investments (fair value, portfolio, derivative, accrual). Now concepts related to management dominated (opportunities, challenges, strategies). Style too changes:

Aside from individual words, it’s the nature of the Bank’s language that is changing: becoming more abstract, more distant from concrete social life; a technical code detached from everyday communication, and pared down to the economic factors crucial to the repayment of the debt. Solutions are

⁹⁵ Pierre Bourdieu, “The Force of Law: Toward a Sociology of the Juridical Field” (1987) 38 *Hastings LJ* 805 at 816.

⁹⁶ *Ibid.*

⁹⁷ *Ibid* at 841.

⁹⁸ *Ibid* at 842.

disengaged from any specificity: they are the same for everybody, everywhere.⁹⁹

Strange things, too, happen to the grammar. A nominalization is a derived abstract noun or a verb that is transformed into a noun: cooperation, development, implementation, prioritization. The literary effect, says Moretti, is to change the focus of the reports from, “not economic agents, but principles.” Instead of social forces or social conflicts, general ideas move to the centre of the Bank’s reports. But the biggest surprise concerned the word most frequently used by the Bank. In most forms of English text, the most common word is the word “the.” Yet in the Bank’s reports, especially after the 1970s, the most common word is “and.” In normal scholarly writing, the word “and” accounts for approximately 2.6% of all words; by the 1970s, the word “and” accounted for over 7% of the Bank’s words. Instead of the definite article (grammatically used to indicate an object), the ascendance of the word “and” indicates the “multiplying words, and most particularly nouns.”¹⁰⁰ Consider:

promote corporate governance *and* competition policies *and* reform *and* privatize state-owned enterprises *and* labor market/social protection reform
There is greater emphasis on quality, responsiveness, and partnerships; on knowledge-sharing *and* client orientation; *and* on poverty reduction.¹⁰¹

The result of the analysis shows how, in the 1970s, a new style emerged that spoke to the Bank’s new place in the world. Facts were out and broad, sweeping, generalizable policy statements were in. Right there in the grammar, evidence of a fundamental change in international law.

1.7 The plan for this dissertation

The sort of work I propose to embark upon is, especially in Canada, rare. This is for two main reasons: first, it is only recently that it has become technologically possible to

⁹⁹ Franco Moretti & Dominique Pestre, “Bankspeak: The Language of World Bank Reports” (2015) 92 *New Left Review* 75 at 87.

¹⁰⁰ *Ibid* at 94.

¹⁰¹ *Ibid*.

meaningfully analyze text computationally. Most of the techniques used by researchers are less than a decade old. Second, it was only in the past six months the researchers began to make large datasets of Canadian law available. In part, this is a feature of the fact that there are few computationally informed legal scholars in Canada who have the skillsets required to scrape and collate law from the internet. The Department of Justice has, for example, made copies of all Canadian federal statutes and regulations available for download for about a decade and a half—but no scholar has yet, to my knowledge, worked with this dataset. Similarly, the Supreme Court of Canada, the Ontario Court of Appeal, the Federal Court of Canada, and the British Columbia Supreme Court all allow for researchers to scrape their websites, but no researchers did so until quite recently.

This, then, is an exciting moment to engage in computational legal work in Canada: with the right skillset, there are opportunities to apply new research methods to as yet never computationally analyzed textual data. For this reason, I aim to produce several ‘first-pass’ studies with different datasets. If there is a constant, it is this: the literary study of form. Throughout my attention is on the repeating features of texts because I believe that the study of form can yield sociological insights. The studies come in two groups.

The first focuses on the legal change over long periods of time. Each of these studies shares the same overall objective: to quantitatively describe the size and shape of Canadian law. This essay began with a discussion about the archive’s size and each of these studies aims to help us better understand how much legal text, and its shapes, there is. Or, put another way, I make the case that there really is too much out there to read. I then use topic modeling and network analysis algorithms to map forms and text to show how they develop over time. Chapter 1 will be the first such analysis of Canadian federal law and statute. Chapter 2, asks a different question about words: how many words get uttered in the Supreme Court of Canada, by whom, and along what patterns? To develop this study, I used cutting edge generative artificial intelligence technology to transcribe and then analyze hours of Supreme Court of Canada video hearings.

The next section swaps out this focus on the long temporal view for a closer examination of law. Chapter 3 is a study of national security deportation orders. It computationally connects the decisions of first instance adjudicators to related judicial

review decisions and shows how, even when different adjudicators consider records that are word-for-word identical, they come to different reasons. Put differently, it is a meso-level study of how different institutional actors compete, despite what doctrine might suggest, to define what legality is.

Then, my most innovative studies. In Chapter 4, I report on a novel study where I analyzed ten years of Federal Court of Canada refugee law jurisprudence. It is millions of words long, well beyond what anyone can read, here I show how this text can be technologically sorted and organized. This study, which is in many ways an original approach to legal scholarship, provoked me to think how I could improve my approach and what the uses of this approach might be. In Chapter 5, I repeat the study but this time focusing on employment insurance law and using more impressive technology. The results, I believe, are even more compelling. As I worked on this project, I also began thinking about this artificial intelligence—particularly this approach—might be of use for legal teaching. These questions bookend this section.

To conclude, I step back from these empirical studies of the legal form. In a sense, this project is a rebuttal to a current strain of thinking about law and technology. For some scholars, new technologies and the emergence of new sorts of artificial intelligence promise to fundamentally transform law's forms and relationships. In *The Legal Singularity*,¹⁰² Abdi Aidid and Benjamin Alarie describe a futurist vision of law in which technology eliminates much of law's uncertainty and incoherence. Building on an old Holmesian vision of what law could and should become, they argue that we can and should work to develop a version of law that is complete, consistent, and immediate. This is a vision from which I dissent. As Bourdieu reminds us, law is always the subject of contest and its meaning can never be fixed. To claim that technology can help achieve a complete and completely coherent law is only another type of claim about what law should be. To my mind, the true gift of this technological moment is critical: we will not see a singularity,

¹⁰² Abdi Aidid & Benjamin Alarie, *The Legal Singularity: How Artificial Intelligence Can Make Law Radically Better* (Toronto: University of Toronto Press, 2023).

but we may be able to illuminate, in new and unexpected ways, the moments and contests in the past that gave rise to the present.

2. The Measurement of Federal Law

How many words long is Canadian federal law? Is it growing? Is it shrinking? Or is it staying just about the same length? If its length is changing, is it changing across all law or in specific zones? Do law's forms—the shapes of its sentences—look the same as they did a decade ago? Put together, if we tried to measure law, what would we see? What would federal law, contemplated from a distance, look like and what does it tell us?

Counting: this is not a normal method of qualitative method analysis, but I hypothesize here is that it is the right method to detect deep, tectonic drifts in law. Yes, statutes come and go as new laws are passed and old laws are repealed to respond to the peculiarities of moments and events. But take those moments and events and sum them up, might a larger pattern of interest appear? As lawyers, we are so used to just reading the laws and thinking about how they could or should be applied that we do not, nearly as often as we should, pay attention to larger critical questions.

Let me put it differently. The frog does not notice when the water is slowly boiled. Well, is the same true for lawyers? So focused on discrete changes here and there, it may be that legal researchers and the legal profession miss those deep and slow changes—those measured in years or decades—that they are a part of and, sometimes, causing. This is measurement's promise: to help us better understand some of the changing, or unchanging, features of legality that, because of their size, are beyond our vision.

This study followed a four-part methodology. First, I built annual consolidations of statutes and regulations as they existed on January 1st of each year, from 2007 to 2023.

After that, I counted. How many laws are there? How many words within each law? How many sections are there? What major differences in word usage can we detect when we compare the law in 2007 to the law in 2023? Third, to understand the structure of Canadian law I drew four maps: two for 2007 and two for 2023. Relying on the explicit citations between statutes and regulations (when, for example, the *Canada Evidence Act* references the *Criminal Code*), I produced visualizations of law as a network. Then, fourth, to better understand these maps, I used algorithmic tools to partition and group laws into different zones, which are illustrated in two new maps, which show the different zones or genres of Canadian federal law.

The findings I present here strike me as intuitive. Yes, law is growing. Yes, there are more words of regulations than of statutes. Yes, both regulations and statutes are growing at roughly the same steady rate. Yes, there are more explicit interconnections between laws today than there were in 2007. Yet in the process of re-representing familiar legal texts as a corpus for quantitative analysis, surprising perspectives on law emerge. There may be more words of regulation, but that is largely because regulations often feature long schedules. Remove all schedules, and there are more statutory words than regulatory words. Even though statutes and regulations grow at roughly the same rate, there is considerably more regulatory churn than statutory churn because regulations get repealed and replaced more often. Put differently, Canadian law is not only expanding, it is renewing: only 49.7% of sections that are present in the 2023 consolidation were also present in identical form sixteen years earlier. And this is a striking observation: compare the law in 2007 to the law in 2023 and you will find that it is half new.

But, then again, not really. For all of this obvious evidence of dynamism (all these new words and connections), my ultimate conclusions point to the basic stability of Canadian law. New sections and sentences are not written into a void, they are written into and next to pre-existing structures. Add new sections on, for instance, medical assistance in dying to the criminal law, and you build and graft on new ideas to pre-existing patterns of thought. Bourdieu says that our deep and unconscious habits and predispositions form a habitus, a common sense that helps social actions understand their social world. The maps I show here imply that Canadian federal law has such a common

sense: there are new words all over the place, but beneath them there is a deep common sense to law's shape. Over here, revenue raising laws group together; over there, we find fishing regulations.

I do not claim to have discovered the fact that there are genres in Canadian law. Instead, I show that we algorithmic tools prove to be startlingly good at empirically confirming what we might sense. But they also open up our senses. By the end of this paper, I hope to have convinced you of things you have neither thought of nor that you would have guessed at. For example: (1) law concerned with the raising of revenue is stable, meaning that it is growing at a slower rate than law overall; (2) criminal law's character is changing: once it was a zone unto itself, but recently it has absorbed aspects of administrative law; and (3) most exciting, we will see extraordinary innovation: climate law and Indigenous self-government law, absent in 2007, emerge as two distinct new fields by 2023.

And, tantalizingly, some of these macro trends and observations find expressions in the grammatical micro. Law is not only largely renewed, by 2023 it is decidedly more conditional, descriptive, and object oriented than law in 2007. When we examine the sentences, we find major growth in a few key words over this time. When the law of 2023 is held against the law of 2007, we clearly see the significance of new words. Contemporary law features more "if"s, more "is"s, and more "that"s. What is more, different areas of law feature their own sentence structure. Look closely (or, more properly, look from afar) and we will see that some zones of law feature more nouns and others more verbs.

One point about style. This is a report on an exploratory study, and I spend a considerable amount of time discussing method. For algorithmic conclusions about text to be interesting, it is important to understand what the methods are working with and how they are treating them. Patterns abound in data, but many of those patterns will have little to say about the concerns of a particular researcher. Throughout I explain how and why I explored the subject matter the way I did so that each reader can judge for themselves whether my approach can appropriately support the conclusions I draw. In the sciences, researchers are often just as concerned about replicability and method as they are with conclusions. While these are not often the priorities of legal researchers, perhaps they

should be, and I make an effort to speak clearly to any other researcher who is interested in assessing my approach or who decides to undertake parallel research.

But first I set out my priors and acknowledge the sources of my inspiration. And that is the matter to which I now turn.

2.1 Scholarly Context

Most computational legal research can be grouped into one of two categories: law as code or law as data.¹ Law as code scholarship presumes that there is a basic affinity between legal codes and computer code. Once this affinity is noticed, it becomes possible to import insights developed in the computer sciences into the legal sciences and *vice versa*. In different moments, this scholarship has obtained different forms. In the 1960s, for example, legal scholars were interested in using mathematical and scientific tools of representation to design clearer and more effective statutes.² Indeed, today's leading journal of science and the law (*Jurimetrics: The journal of law, science, and technology*) began in 1955 as *MULL: Modern Uses of Logic in Law*.

By the 1980s and early 1990s, during an upsurge of interest in artificial intelligence and so-called expert systems, scholars were working to represent statutory codes as computer code.³ But quickly funders across the sciences and social sciences became disappointed with the “observed gulf between the promises and reality” of these early rules-based systems and this sort of computational research entered a period now referred to as an “AI winter.”⁴ In some legal fields, most notably tax law, governments and

¹ Jens Frankenreiter & Michael A. Livermore, “Computational Methods in Legal Analysis” (2020) 16 Annual Review of Law and Social Science 39 at 41.

² See, as an example, Laymen Allen, “Some uses of Symbolic Logic in Legal Practice” (1962) 3:2 *MULL: Modern Uses of Law in Logic* 119.

³ See as examples: MJ Sergot et al, “The British Nationality Act as a logic program” (1986) 29 *Comm ACM* 370; TJM Bench-Capon, ed., *Knowledge-Based Systems and Legal Applications* (San Diego: Academic Press, 1991); Edwina Rissland & David Skalak, “CABARET: rule interpretation in a hybrid architecture” (1991) 34 *Int J Man-Machine Studies* 839.

⁴ Chris Wiggins & Matthew Jones, *How Data Happened: A history from the age of reason to the age of algorithms* (Norton and Company, New York, 2023) at 182.

accountants remained interested in continuing this work, leading even to the development of a specific computer programming language for law.⁵

Law as data scholarship, in contrast, is less interested in thinking about how code can make or be law, but in using computationally enhanced methods to advance legal research. Inspired by the “text as data” work of social scientists⁶ and the “distant reading”⁷ methodology of the digital humanists, legal researchers started to explore how legal texts and sources could be mined for insights about the law.⁸ In this sense, law as data scholarship is an elaboration of traditional empirical legal research but, because of the increasing sophistication of computational methods, is able to engage with more types of data and much greater volumes of data. The range of law as data research is large, but a few recognizable sub-fields are now coming into view. For many scholars, an important preliminary step is to count, measure, and statistically summarize legal data. Relatively straightforward work can yield novel classes of insights. For example, a researcher might measure the length of judgements and search out the variables that lead to shorter or lengthier decisions⁹ or measure the readability of judicial texts.¹⁰

Because of how young the field is, little is known about some basic questions. For example, it was not until 2010 that researchers first reported that the US Code was 22 million words long.¹¹ And then it was not until 2020 that researchers sought to measure

⁵ Liane Huttner & Denis Merigoux, “Catala: moving towards the future of legal expert systems” (2022) 25 AI & L 1.

⁶ Justin Grimmer, Margaret Roberts & Brandon Stewart, *Text as Data: A new framework for machine learning and the social sciences* (Princeton: Princeton University Press, 2022).

⁷ See Franco Moretti, *Distant Reading* (New York and London: Verso, 2013) and Matthew Jockers, *Macroanalysis: Digital Methods and Literary History* (Irbana: University of Illinois Press, 2013).

⁸ Michael Livermore & Daniel Rockmore, eds., *Law as Data: Computation, text, and the future of legal analysis* (Santa Fe: Sante Fe Institute, 2019).

⁹ Ryan Black & James Spriggs, “An Empirical Analysis of the Length of U.S. Supreme Court Opinions” (2008) 45 Hous L Rev 622.

¹⁰ Ryan Whalen, “Judicial Gobbledygook: the readability of Supreme Court writing” (2015) 125 Yale LJ 200.

¹¹ Michael J Bommarito II & Daniel Martin Katz, “A Mathematical Approach to the Study of the United States Code” (2010) 389 Physica A: Statistical Mechanics and its Applications 4195.

the extent to which the Code was growing and changing over time.¹² As the same researchers explain, basic quantification has only recently become an interest for legal researchers because:

scholars have traditionally regarded the law as hardly quantifiable, and although there is no dearth of empirical legal studies it is only recently that researchers have begun to apply data science methods to law. To date, there have been relatively few quantitative works that explicitly address legal change and almost no scholarship exists that analyses the time-evolving outputs of the legislative and executive branches of national governments at scale.¹³

In a sense, new computational methods offer to swap out familiar views for new, wider, ones. If professors, judges, and lawyers are students of close reading, computational and quantitative methods suggest that texts can be “read” and understood in different terms.

But it remains early days. What empirical hints we have about law’s largest shapes are generally American or European, where statutory text is growing and growing in measurable complexity, and are drawn primarily from one recent study comparing Germany to the United States of America.¹⁴ That study found that law was growing in length and complexity in both countries, leading the researchers to conclude that “extensive growth in legal complexity as a function of volume, interconnectivity, and hierarchical structure of the legislation in both countries—evidence that the highly industrialized countries we study seek to manage behaviour by building increasingly complex bodies of legal rules.”¹⁵ Though, of greater interest, the researchers also found that growth in length and complexity was uneven. Social security and environmental law, for example, grew at faster rates in both countries.

Often lawyers, policy makers, and legal scholars consider legal growth with suspicion. In an essay promoting a book he co-wrote with one of his clerks, Neil Gorsuch—after quantitatively describing the scale of American law—declares that “[t]he explosion

¹² Daniel Katz et al. “Complex societies and the growth of the law” (2020) 10 *Scientific Reports*. [Complex societies]

¹³ *Ibid* at 1.

¹⁴ *Ibid*.

¹⁵ *Ibid*.

of law has taken a very real toll on the lives of everyday Americans.”¹⁶ When America was founded, he claims, constitutional drafters were believers in the rule of law, but nonetheless worried about what would happen if “lawmaking becomes too easy, when it is a task too far removed from the people, and when laws become too hard to find and too difficult to understand.”¹⁷ Perhaps there is something here. It may be that the texts we are about to engage with, in combination with the fact of their size, are evidence of a dangerous trend, where no one really knows or understands (or can hope to know or understand) the nature of the legal relationships that they are a part of.

But equally, scale may be an indicator of overall system health. A new field of study, complexity science or complexity studies, holds that the growth of information in a society is basic evidence of human growth and flourishing.¹⁸ Perhaps statutory and regulatory sections are not evidence of information, but perhaps they show how our society is becoming increasingly interconnected, is increasingly looking to coordinate new forms of social interaction, is increasingly preoccupied with subjects that were once unknown to us. But, again, this is a developing field. It was only in February of 2024 that a journal first published a special issue considering what complexity science might have to say about law and legal relations.¹⁹ And while there is certainly scholarly interest in developing metrics and theoretical perspectives that could help researchers make claims about the beneficence of some forms of complexity over others, studies continue to be largely descriptive.²⁰

So, what follows is in keeping with the wider scholarship: it is a Canadian effort to describe Canadian law, just as other bodies of law have been described. I do not claim to

¹⁶ Neil Gorsuch & Janie Nitze, “America Has Too Many Laws” (5 August 2024) The Atlantic, online: < <https://www.theatlantic.com/ideas/archive/2024/08/america-has-too-many-laws-neil-gorsuch/679237/>>.

¹⁷ *Ibid.*

¹⁸ César Hidalgo, *Why Information Grows: The Evolution of Order, from Atoms to Economics* (New York: Basic Books, 2015).

¹⁹ Pierpaolo Vivo, Daniel Katz & JB Ruhl, “A complexity science approach to law and governance” (2024) 382:2270 *Phil Trans R Soc A* 1.

²⁰ JB Ruhl & Daniel Katz, “Measuring, Monitoring, and Managing Legal Complexity” (2015) 101 *Iowa L Rev* 191.

validate Gorsuch's views, nor do I claim to prove that Canadian law exhibits characteristics of healthy growth. But I do claim there is something here. Law is larger and just looks different today than it did over a decade ago. Let us now figure out why that might be.

2.2 Generating a Corpus of Canadian Law

How many laws are there? How many sections are there? How many words long is Canadian law? These first research questions, methodologically the simplest, accounted for this study's largest difficulty: how could I get a hold of all Canadian law so that I could write a program to measure it? The challenge here is legal. Canadian legal researchers and professionals benefit from a comparatively rich information research environment. Major private corporations (WestLaw and QuickLaw) each maintain large databases of Canadian primary and secondary law that researchers can access for a fee. The Federation of Canadian Law Societies also operates a non-profit corporation known as the Canadian Legal Information Institute, which provides free access to Canadian law and has a mandate to "provide efficient and open online access to judicial decisions and legislative documents."²¹

The challenge for researchers interested in engaging in computational research is that each of these organizations forbids bulk access to its data.²² In a basic sense, this means that researchers (and all users for that matter) are not allowed to download large quantities of the data held in each data repository. To my knowledge, no corporation has gone on the record to justify these policies, though several possible explanations seem clear. First, to allow bulk access would be to undermine the value of each corporation's asset: the data. It is only by keeping each database close that the company that it can justify charging lawyers and researchers for the data.

²¹ "What is CanLii?", online: <<https://www.canlii.org/en/info/about.html>>.

²² The terms of use for CanLii, for example, prohibit "[b]ulk or systematic downloading of documents, including via programmatic means or, for greater certainty, the hiring of human resources used to manually download documents: "Terms of Use," online <<https://www.canlii.org/en/info/terms.html>>.

Second, and this explanation might go some way to explain why even a non-profit corporation charged with improving access to data might decide to limit access to it, some Canadian tribunals have indicated some measure of discomfort with the idea of bulk access. In a paper published by the Canadian Judicial Council, for example, it was warned that openness needed to be balanced against risks that general access could inadvertently increase misinformation about the justice system and lead to the inadvertent disclosure of private information.²³ To the extent that these concerns hold water with respect to information produced for, within, or because of judicial proceedings, they seem empty with respect to statutes and regulations: those laws made generally by legislatures and the executive ought, it seems to me, be made generally available.

Fortunately, while I was researching this matter, I learned that the Canadian federal government has a more permissive and open policy. The *Legislation Revision and Consolidation Act* authorizes the Minister of Justice and Attorney General of Canada to publish an electronic consolidation of all Canadian law and regulation.²⁴ This is a relatively recent departure from past practice, in which an expert commission would every few decades or so collect, collate, and edit all Canadian federal law to present an up-to-date version of law.²⁵ Unlike the American example, where all laws are merged under distinct titles of the US Code, Canadian laws retain their discrete identities but amendments are merged into one final and authoritative document of laws. Digitization has allowed the Canadian government to maintain a much more responsive and complete consolidation. If once it took several decades to re-consolidate law, the Department of Justice has since 2006 published updated consolidations approximately every two weeks. In turn, it shares these consolidations with publishers to disseminate Canadian law. Unbeknownst to me, the Department of Justice made programmatic access to these consolidations generally

²³ Jo Sherman, *Guidelines For Canadian Courts: Management of Requests for Bulk Access to Court Information by Commercial Entities* (Canadian Judicial Council: Ottawa, 2021).

²⁴ *Legislation Revision and Consolidation Act*, RSC 1985, c S-20, s 28.

²⁵ Norman Larsen, “Statute Revision and Consolidation: History, Process, and Problems” (1987) 19:2 Ottawa L Rev 321.

available in 2020.²⁶ This means that anyone can download a local copy of a complete set of consolidations from 2006 to present day.

That said, the data is relatively raw and must be manipulated and rearranged to suit each particular use case. As of the time of writing, the complete consolidation is a 9.56 gigabyte dataset that contains English and French versions of Canadian law, sub-divided into statutes and regulations, with dated folders that contain files that are, for that particular date, up-to-date versions of each statute and regulation. Each statute and regulation is drafted using the extensible markup language (XML). XML is a type of computer language used, not to code programs, but to store and represent complex data structures.²⁷ As users of statutes and regulations will appreciate, providing a simple text file (for example) of statutes would make law quite difficult to work with: laws are made up of many different components (sections and subsections, paragraphs and subparagraphs), contain cross-references, and can be divided into different parts. While, yes, the text is the most important part of a law, that text's structures and relationships matter and XML drafting allows for this structure to be maintained and communicated, along with the textual content of each law.

Or, to restate all this in simpler terms, the consolidations provided by the federal government are comprehensive and extremely rich, and to make them workable for this research project, I needed to re-arrange them to suit my purposes. After some consideration, I decided to write a program that would produce consolidations of law, as it existed on January 1st at 12:01 a.m. for each year from 2006 (the first year for which a complete digital consolidation was published) to 2023. This program transformed each law into four data frames (or spreadsheets): one for metadata about the law, one for introductory text, one for body text, and one for schedules. For the body text data frame, I further subdivided all the text into its lowest common denominator and located that unit of text in its own row (if a section contained four subsections, each subsection was the final

²⁶ Department of Justice, “Consolidated federal Acts and regulations –Bulk XML (PIT (Point in time)),” online: <<https://open.canada.ca/data/en/dataset/4d277b82-c4af-4431-b31d-48fcece6dd8a/resource/dd28f08b-0b48-485a-a9db-b6176ff3d1f7>>.

²⁷ W3C, “Extensible Markup Language (XML),” online: <<https://www.w3.org/XML/>>.

textual unit; if a paragraph has two subparagraphs, each subparagraph was the final textual unit) with companion cells noting which sections, subsections, paragraphs, or subparagraphs each textual unit was a part of, and noting whether that section contained cross-references to other laws.

With the data restructured, I then wrote code to count the number of statutes, the number of textual units, and the number of words. I pause here to note a methodological problem. It was difficult to settle on a programmatic definition of a word. Statutes and regulations can contain all sorts of data: words, numbers, geographic coordinates, and so on. Because my primary interest concerned the linguistically expressive components of law, I endeavoured to draft code that would isolate English language words. To do this, the program separated units of data every time it encountered a whitespace character (i.e. a space between words), and then excluded all units that contained numbers and that were within parentheses. Recognizing that this method inevitably captured some units of data that are not words, I refer to each of these units as tokens, not as words.

Finally, while I had access to equivalent English and French data, I only examined English language representations of statutes. In the future, it may be of interest to explore how the same laws are represented in different languages. I note in passing that it appears that the French language versions of laws may be wordier than English language versions. The English version of all acts studied here was 3.42 gigabytes, while the same French language version was 3.51 gigabytes.

2.3 Taking Law's Measurements

2.3.1 A Growing Law

Canadian law is growing. In 2007 there were 3,346 laws and regulations, 71,387 sections of law and regulation, and 14,077,824 tokens within those laws. By 2023 there were 3,943 distinct laws and regulations, 89,236 sections, and 18,053,570 tokens. The scale of Canadian law and its growth is remarkable. In 2007, Canadian law equaled just over 25

editions of War and Peace; in 2023, it equaled just over 32. The median sizes of sections increased only slightly from 78 words in 2007 to 80 words in 2023.

Figure 1 shows that from 2007 to 2013 the number of statutes and regulations, the total number of sections, and the total number of words, each grew approximately 1.1 times. After 2013, the total number of words and sections each continued to grow at roughly the same rate (and have each grown approximately 1.3 times by 2023), but the growth in the total number of statutes and regulations slowed. The primary reason for this slowed growth is related to the relative number of statutes to regulations. In 2007, for example, there were 677 different statutes and 2,669 different regulations. By 2023, the number of statutes grew 1.25 times to 849; while the number of different regulatory instruments grew 1.16 times to 3,094. Put differently, in 2007 there were almost 4 regulatory instruments per statutory instrument; by 2023, this ratio dropped to 3.6 regulatory instruments per statutory instrument.




Figure 1: The rates at which Canadian federal law has grown, relative to 2007.

Even though there were many more regulatory instruments than there were statutory instruments, acts were on average much larger. As the left panel of figure 2 shows, each year the relative proportion of regulatory to statutory words remained steady, with 1.1 regulatory tokens to each statutory token. Yet, as the right panel shows, the

primary structural reason that there are more regulatory tokens than statutory tokens is that regulations often feature large schedules. If schedules are discounted from the totals,




Figure 2: On the left, the number of tokens of statutes and regulations including schedules; on the right, the number of tokens in statutes and regulations, excluding schedules.

then there are more statutory tokens than regulatory tokens. Excluding schedules, this means statutes grew from 5,452,440 tokens to 7,099,250 tokens (difference: 1.64 million tokens) and regulations grew from 5,286,537 to 6,936,615 tokens (difference: 1.65 million tokens).

This raises a question: what are schedules used for in statutes and regulations, and why might there be such a discrepancy in terms of their relative lengths? Reviewing the ten statutes with the longest schedules in 2023, shows that in eight cases the schedules were used to incorporate international legal instruments into Canadian law, in one case to describe the boundaries of Canada’s national parks, and in a final case to set out rules for the calculation of excise taxes. In contrast, the ten regulations with the longest schedules described geographic boundaries six times, set out rules and identified specific matters of concern for various regulator regimes three times, and listed regulatory offences once. Or, to put it simply, long schedules in acts tended to indicate an engagement with international law while long regulatory schedules were more often long lists of facts.

2.3.2 Of big laws and small laws

We all appreciate that there are some extraordinarily long acts and regulations and some that are extraordinarily short. Beyond that basic sense, however, we all will struggle to

answer second-order questions: which are the longest? What is the distribution of law sizes? Do large laws grow more or less than smaller acts? As Table 1 shows, comparing 2007 and 2023, there is a basic stability in terms of the ten largest statutes. In both years, the five largest statutes are the same, each representing roughly the same proportion of all law. The standout here concerns tax law, with the two largest statutes accounting for almost 15% of all statutory words. In contrast (see Table 2), there is more obvious change with respect to regulations and the regulatory space is not so dominated, at least when compared to the statutory space, by a few extraordinarily large laws.

Longest Acts	Body tokens	Schedule tokens	Total tokens	% of all tokens
2007				
<i>Income Tax Act</i>	701,862	0	701,862	10.8
<i>Excise Tax Act</i>	239,199	34,269	273,468	4.2
<i>Criminal Code</i>	243,514	14,023	257,537	4
<i>Bank Act</i>	189,421	574	190,016	2.9
<i>Insurance Companies Act</i>	186,262	1,261	187,544	2.9
<i>Canada Shipping Act</i>	97,074	9,157	106,231	1.6
<i>Trust and Loan Companies Act</i>	92,436	0	92,457	1.4
<i>Cooperative Credit Associations Act</i>	84,343	0	84,364	1.3
<i>Canada Elections Act</i>	83,348	868	84,237	1.3
<i>Geneva Conventions Act</i>	1,238	80,764	82,002	1.3
2023				
<i>Income Tax Act</i>	963,492	235	963,727	11.5
<i>Excise Tax Act</i>	322,771	38,693	361,464	4.3
<i>Criminal Code</i>	287,316	19,085	306,401	3.6
<i>Bank Act</i>	225,464	829	226,412	2.7
<i>Insurance Companies Act</i>	197,930	1,122	199,073	2.4
<i>Canada Elections Act</i>	126,313	3,366	129,700	1.5
<i>Trust and Loan Companies Act</i>	100,486	0	100,507	1.2
<i>Cooperative Credit Associations Act</i>	89,314	0	89,335	1.1
<i>Geneva Conventions Act</i>	1,297	82,620	83,917	1
<i>Canada Labour Code</i>	81,655	0	81,655	1

Table 1: The ten largest statutes (2007 and 2023)

On average, statutes were 9,589 tokens long in 2007 and 9,895 tokens long in 2023, but the median sizes were 2,503 and 2217 tokens respectively. Indeed, 75% of all statutes were shorter than 7,776 tokens in 2007 and 8,240 tokens in 2023. Put differently, the overwhelming majority of statutes were only a few thousand words long, while those long statutes discussed above were extreme outliers. The same pattern replicated itself with

respect to regulations, just that regulations tended to be even shorter. In 2007, 75% of all regulations were 1,752 tokens or less; in 2023, the number dipped slightly to 2,742 tokens.

Longest Regulations	Body tokens	Schedule tokens	Total tokens	% of all tokens
2007				
<i>Income Tax Regulations</i>	262,474	35,750	298,224	3.9
<i>Food and Drug Regulations</i>	160,966	10,613	171,579	2.3
<i>Canadian Aviation Regulations</i>	135,279	0	135,368	1.8
<i>NAFTA Rules of Origin Regulations</i>	36,862	70,044	107,011	1.4
<i>CCFTA Rules of Origin Regulations</i>	20,445	66,900	87,427	1.2
<i>Ontario Fishery Regulations, 1989</i>	7,709	70,689	78,456	1
<i>Motor Vehicle Safety Regulations</i>	16,084	57,482	73,566	1
<i>Processed Products Regulations</i>	10,812	55,701	66,513	0.9
<i>Immigration and Refugee Protection Regulations</i>	61,827	295	62,252	0.8
<i>Life Saving Equipment Regulations</i>	24,373	28,584	52,957	0.7
2023				
<i>Income Tax Regulations</i>	304,177	35,258	339,451	3.5
<i>Food and Drug Regulations</i>	216,185	18,735	234,920	2.4
<i>Canadian Aviation Regulations</i>	214,947	0	215,036	2.2
<i>Transportation of Dangerous Goods Regulations</i>	50,836	81,537	132,468	1.4
<i>Immigration and Refugee Protection Regulations</i>	90,476	1,319	91,925	1
<i>CCFTA Rules of Origin Regulations</i>	20,445	66,900	87,427	0.9
<i>Cannabis Regulations</i>	76,663	0	76,731	0.8
<i>Motor Vehicle Safety Regulations</i>	20,672	47,290	67,962	0.7
<i>Emergency Order for the Protection of the Greater Sage-Grouse</i>	1,638	62,131	63,837	0.7
<i>Selected Listed Financial Institutions Attribution Method (GST/HST) Regulations</i>	61,793	0	61,923	0.6

Table 2: The ten largest regulations (2007 and 2023)

Recall that statutes and regulations each grew about 1.3 times from 2007 to 2023. This general growth trend is not immediately apparent when examining only the largest acts and regulations. In some cases, individual growth exceeded overall growth (for

example, the *Income Tax Act* grew 1.37 times) but in other cases large acts shrank relative to the whole of law (for example, the *Criminal Code of Canada* only grew 1.19 times).

For the most part, the length of sections of law were the same in 2007 and 2023. In 2007 the average number of tokens for statutes was 171 and 126 for regulations; in 2023, these figures rose slightly to 175 and 134. There was, however, one outlier: income tax law. On average, a section of the *Income Tax Act* was almost 8.5 times longer than the average section of law. This likely shows that the drafters of tax law are not the same as drafters of all other legislation, and may suggest that income tax law may, if it is indeed located in a different bureaucracy, obey a different common sense than the rest of law.

2.3.3 Law grows and law churns

It is somewhat misleading, however, to only measure growth if the goal is to measure the extent of change within a corpus or within a law. This is because some of the most major changes a legislator or executive can make might not register as growth. When whole sections of parts of a law are repealed and replaced, wholesale and major changes are made to a law. But if those wholesale changes are more or less equal from a token count perspective, they will not register if the researcher only measures growth.

One way of appreciating the extent of change is to consider all those sections that were present in Canadian law in 2007. That year, as discussed above, there were 71,397 sections of statute and regulation. Of those, 31,691 sections were in statutes. By 2023, 3,918 of those sections had been repealed (12%), 7,332 (23%) were amended in some way, and 20,441 (65%) were left untouched. By 2023, there were 44,002 sections of statutory law. This means that only 46% of statutory law sections in 2023 looked the exact same way that they did in 2007. For regulations, this proportion of law that remained unchanged was even less. Of the 49,153 sections of regulation that existed in 2023, only 46% were also present in identical form in 2007.




Figure 3: How tokens entered and left law between 2007 and 2023 (left: statutes; right: regulations)

We can better appreciate the extent of these changes with a visualization. Recall that both states and regulations grew, when schedules are discounted, by approximately the same figure: 1.6 million tokens. Figure 3 shows how, when the 2007 and 2023 consolidations are compared, tokens entered law or left law. Tokens can only enter or exit a law (when considering sections) in one of three ways: they may be included in a whole new section, they may be removed when a whole section is removed, or they can be added or subtracted when a section is amended. As we can see, even though the final figure ends up being the same, there was considerably more regulatory churn than statutory churn because significantly more regulatory tokens were repealed and replaced. This is not necessarily surprising, but it shows that the regulatory space is, at least from a quantitative perspective, considerably more dynamic than the statutory space.

2.3.4 A new vantage: peeking at the micro

I was curious to see whether all this churn made for linguistically different sets of laws. In one sense, it obviously does: law has changed considerably over these 16 years and its meanings have inevitably changed with it. But more than that I wondered if deep changes in the grammar and vocabulary of law could be detected, beyond the individual changes in meanings of each prompted by additions and amendments of individual sections. To assess this, I undertook a series of increasingly complex methods.

Using the scikit-learn Python library (a toolkit of open-source machine learning methods), I calculated the word frequencies (occurrences of a word for every one thousand words) for the 2007 and 2023 consolidations. In both cases, the top words were the same and occurred at approximately the same frequency. As we would expect, the most common words were determiners and auxiliaries, such as “the,” “of,” “to,” and “or.” Yet interesting and subtle differences presented themselves between regulations and statutes. As Table 3 shows, the words “or” and “to” were more common in statutory text than in regulatory text, while the word “and” was more common in regulatory text.

Top statutory words	Occurrences per 1000 words (2023)	Top regulatory words	Occurrences per 1000 words (2023)
the	97	the	96
of	56	of	54
or	34	in	30
to	34	to	29
in	30	or	26
and	21	and	23

Table 3: Top word frequencies, statutes and regulations (2023)

Given that these discrepancies manifest in multiple years at the same rates, they imply a meaningful difference between statutes and regulations. Settling on what that meaningful difference might be is, however, challenging. My hypothesis is that each of these most common words serves, in the context of a sentence, a descriptive function. The word “the” declares the presence of an object, “of” describes a relationship between a noun and a verb and another noun, “or” and “and” describe different circumstances or conditions, while “to”—the most versatile word on the list—can, *inter alia*, announce the presence of a range (from here to there), a relationship (reports to the Minister), or be used to transform a verb into its infinitive form (to protect). That there are more of these

words (except for the word “and”) in statutes and regulations may indicate that there is more high-level description occurring in statute than regulation.

What is clearer, however, is that there are significant linguistic differences in the character of statutes and regulations over time. Figure 4 shows the top ten and bottom ten word frequency changes, for both statutes and regulations, between the 2007 consolidation and the 2023 consolidation. While at first glance, these figures may not appear significant (a one- or two-word swing out of a thousand instances), when aggregated they become quite significant. Over a 10-million-word jurisprudence, a one-word higher word frequency (per one thousand words) would equal 10,000 more total occurrences.

A few observations present themselves. First, and perhaps the clearest, is that language that may be more antiquated is being removed from Canadian law. “Shall” is steadily being replaced by “must” in both regulations and statutes, just as “pursuant” is

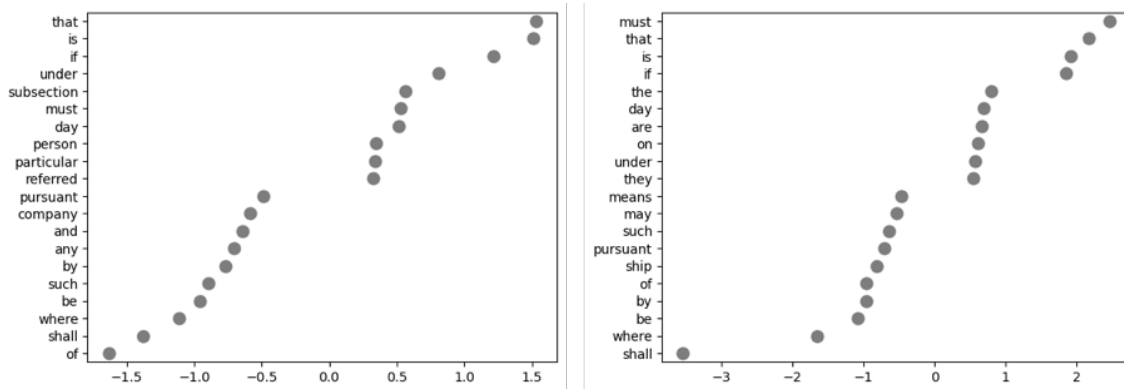


Figure 4: Word frequency changes (per 1000 words) between 2007 and 2023 (left: statutes; right: regulations)

being replaced by “under.” This observation points to the influence of the decisions and practices of legislative drafters. I suspect that sometime in this period an explicit decision was made by professional drafters to swap out some words for others.

Second, it appears that Canadian law is becoming more declarative and definitional. In both states and regulations, there are major increases in the frequency of the “is.” “Is,” which is the third-person present-tense form of the verb “to be,” describes a noun: something is this or it is that. As the frequency of the word “is” goes up, that suggests a general drift towards definition in law. The same logic may explain the

increasing use of the words “day” and “person” as there may, in an increasingly definitional and declarative law, be a tendency to refer more to discrete subjects, such as persons or days when things may, will, or have happened.

If this is the case, it would also help us understand the increase in the use of the word “that.” “That” is a versatile word that is often used to refer to a subject in a sentence. As the usage of the word “that” increases, it suggests that legal texts focus more on describing subjects or describing what subjects may do, or may have done to them. Alternatively, the explanation here may be more quotidian: perhaps a style decision was made to use “that” more often so that drafters would not repeat the names and identities of the subjects and objects in each sentence.

I doubt, however, that is the case. To understand how “that” is used in sentences I used a pre-trained machine learning model, released by explosion.ai as part of the spacy natural language processing toolkit, to identify the grammatical function of each use of the word “that” in both the 2007 and the 2023 consolidations. Across each year, the proportions of the function remained relatively constant. In 2007, in 77% of the usages of the word “that” referred to a subject or object, while in 21% of the usages it was used to mark the beginning of a marker of a subordinate conjunction (“an officer may seize anything that they believe is stolen goods.”), proportions that were fundamentally repeated in 2023 (78% and 20%). To summarize, it appears that the word “that” is being used in the same grammatical way, but those usages appear to have increased.

Third, it appears law is becoming more conditional. In statutes, but especially in regulations, there is a marked increase in the use of the word “if.”

2.3.5 Summary: a dynamic and growing law

These high-level measurements point to the dynamism of Canadian law. Law’s length has grown 1.3 times in just over a decade and a half, with the addition of 4 million tokens. More than that, there has been significant renewal as many sections have either been repealed completely or amended. From a purely textual perspective, Canadian legal texts are substantively different and larger than they were sixteen years ago. Yet there are also markers of basic stability. The laws that were large in 2007 are the laws that are large today.

The words that were common 16 years ago are also common today. The median sizes of each section have barely budged. What might explain this state of affairs? Perhaps there is a general inflationary tendency across all law? Perhaps some areas of law are growing more than others, or perhaps some areas are subject to more revision and renewal, while others are only added to? Better understanding the nature of this growth is the task of the next section.

2.4 Mapping Canadian Law

2.4.1 Using network theory to spatialize law

Every lawyer and legal researcher recognizes that there are distinct, though sometimes overlapping, genres of law. Tax law, for example, has an identity that is distinct from fisheries law, which in turn is distinct from veterans' law. What explains the development of these distinct categories? No doubt there is no single explanation. In part, distinct legal identities will flow from the teaching choices of law schools, the editorial decisions of legal publishers, the formal associations within the law profession, and so on. The point is that recognizing genres of law is somewhat idiosyncratic, a feature of different debates, overlaps, and ways of seeing particular to a moment and a position.

Nonetheless, I hypothesize that one explanation—and perhaps the dominant explanation—for the development of distinct genres of law is material and structural. Law can be understood as an answer to social problematics. The legislator or the executive produces or changes law as responses to social occurrence. As developments in the world, or conflicts in social life, present themselves, law, like any number of other cultural interventions, answers with a proposed resolution to those occurrences. When concerns about over-fishing develop, for example, fisheries regulations might be introduced to acknowledge and address those concerns. This is not to state a vulgar Marxist position that the legal superstructure is determined by the economic base, rather it is to say that social life poses questions that invite institutions, legal and otherwise, to develop and formulate their own answers. As social facts develop and change, and as legal answers

adapt and respond, legal texts themselves will develop their own shapes. Fisheries law, to continue the example, will, with the passage of time and with the continuing coherence of the underlying problematic, come to be a field of legality in its own right.

Thinking about genre is useful because the concept of the genre has analytic value. As genres of law rise, fall, and change, they say something large about the relationship between law and society. And for this study, which focuses on law as a complete corpus, genre is a more useful and practical optic than focusing on the hundreds of discrete laws or the tens of thousands of discrete sections.

How might these genres or legal fields be detected? Many computational legal scholars propose that the shape of law can be uncovered by mapping and examining the interconnections between legal texts. Much of this work focuses on the citation patterns of adjudicative decisions, but scholars have also explored how the links between primary legal documents can be used to uncover structure in law. Laws that develop around the problem of overfishing will, we might hypothesize, acknowledge each other because they each answer overlapping concerns. The study of these connections is referred to as graph theory or network theory. A decidedly twenty-first century research approach, the goal is to “model the complex interdependencies between entities... by emphasizing their shared contexts, relationships, and interactions.”²⁸ Here the classic example is a social network, which models the relationships between different people.

²⁸ John McLevy, *Doing Computational Social Science: A practical introduction* (London: Sage, 2022) at 231.



Figure 5: Networked representations of Canadian law (top: 2007; bottom: 2023).

To build my networks of law, I returned to the consolidations I already constructed for each year. One of the categories of data I extracted from the XML files published by the Department of Justice was the explicit citations within a law of another law. Using the `networkx` and `igraph` libraries, I created networked representations of Canadian law corresponding with each annual consolidation I constructed. Interesting observations immediately presented themselves. Statutes and regulations, for example, had different (but certainly intuitive) citation patterns. For each year under study between 96% and 98% of citations from acts targeted other acts while 74% to 77% of citations from regulations targeted acts. Each annual network was also clearly “scale free.” A “scale free” network is one in which there are a few nodes that are connected to many other nodes (i.e. hubs), but in which most nodes are sparsely connected. This contrasts with a network where nodes are equally connected to each other. For each year under study, the top four most-cited laws were the same (the *Criminal Code*, the *Income Tax Act*, the *Financial Administration Act*, and the *Customs Tariff*) and each of these were connected to over two hundred laws. In contrast, most laws and regulations (approximately 80%) were connected to five or fewer laws.

It was also striking to note the increasing interconnection of Canadian federal law. Recall that laws grew in length 1.3 times during the period under study. In contrast, my method of searching out connections between laws found that, over these 16 years, the number of explicit citations between laws increased 1.4 times: in 2007, there were 20,721 connections between regulations and statutes; in 2023, there were 29,520.

To visualize the networks, I used the ForceAtlas 2 algorithm to spatialize law. The algorithm works by simulating “a physical system in order to spatialize a network” where nodes “repulse each other like charged particles, while edges [connections] attract their nodes.”²⁹ It runs iteratively, moving each node, until the network reaches an equilibrium in which the nodes will not move any further. The appeal of this algorithm for me is that it corresponded with my hypothesis about how laws responding to similar concerns might

²⁹ Mathieu Jacomy et al, “ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software.” (2014) 9:6 PloS one.

recognize each other and group themselves: those with no connections between each other, or between other common nodes, would end up far apart, while those with many connections would end up closer together. Figure 5 shows the resulting visualizations for 2007 and 2023.

Comparing these two maps, several notable features come into view. First, and most importantly, the resulting maps validate the methodological choice to visualize law using a network approach. While both maps are distinct, they clearly represent similar morphological features. In both instances, at the right side of the maps there is a similar-looking outcropping. Moving left, there is a major spine that connects, from the top to the middle, several large major nodes. Below that, in the bottom right, a similar community of laws can be observed. In the middle there is a densely connected chaotic zone and, finally, the leftwards extension is more diffuse. At the same time, differences are apparent. The 2023 map is considerably more connected, noticeably wider, and more of a “blob” than the 2007 map, that better resembles a starfish. Though different, the 2023 map is clearly an extension or elaboration of the 2007 map.

2.4.2 Inferring a community structure in Canadian law

For present purposes, perhaps the most interesting feature of the maps is that they appear to show distinct zones, neighbourhoods, or communities of law. In some places, laws are densely connected. In others, loosely connected areas still appear distinct from other regions.

Network scientists have developed several algorithms that can be used to describe the community structure of a network. After reviewing a few of the options, I elected to use the Leiden algorithm.³⁰ The primary reason for this choice is simple: this algorithm identifies communities in a manner that approximates my own view about how legal scholars might recognize genres of law. In short, the Leiden algorithm attempts to maximize a quality function, which is some variable that describes the coherence of a

³⁰ Vincent Traag et al, “From Louvain to Leiden: guaranteeing well-connected communities” (2019) 9:1 Scientific Reports 5233.

community. I elected to maximize modularity, which refers to the extent to which the number of connections within a community exceeds the number which we would expect to be present if connections were randomly assigned. Or, to put it simply, Leiden assesses each node and its connections with its neighbours and, if those connections are stronger than expected, groups the nodes together. It repeats this process iteratively until modularity cannot be improved any further.

One feature of the Leiden algorithm is that it is non-deterministic. This means that if the algorithm is run on the same data multiple times it will produce multiple different results. This is because the algorithm, when it first begins assessing the strength of the connections between nodes, begins in a different random starting place each time. This variability promises to yield better ultimate conclusions because it means that, on repeated runs, the algorithm will return optimized results approached from a slightly different starting point, reducing the chance that a particular optimization is an isolated and non-representative result. To take advantage of this fact, I wrote code, following an approach already developed by the legal researchers whose original work inspired this chapter,³¹ that ran the community detection algorithm 1,000 times and returned those communities and connections that were present in at least 95% (or 950) different runs.

Even when stable communities are detected, some algorithmic decisions will necessarily be incomplete. Assigning a law to one community is a necessarily general decision that will sacrifice some of the particularity of each law. This is especially so for laws that arguably span multiple communities or that deal with multiple concerns. The decision to deploy an algorithm that focused on interconnections necessarily means other features of a legal text get downplayed or ignored. The *Judges Act*, which the algorithm grouped with laws concerned with public service pay, is a good example of the limits of this methodology. As the algorithm correctly noted, the majority of this law deals with questions concerning judicial salaries, benefits, and pensions. For this reason, and this is the data that the algorithm worked with, there are explicit citations to other laws concerning public sector pay. The *Judges Act*, however, also does more: it sets out

³¹ Complex Societies, *supra* note 12.

procedures for judicial discipline, it describes the role of the registrar of the Supreme Court of Canada, and so on. These provisions exist on their own and neither cite nor are cited by other laws, making the algorithm blind to this aspect of the law. This is neither a feature nor a bug: any method of community or genre detection will at some point be called upon to make a decision, the point is that it is important to understand what the tool is producing so that its insights can be properly balanced against its blind spots.

With this caveat in mind, however, I was pleased to find that this method yielded surprisingly interpretable and coherent categories of Canadian law. In Appendix A, I report the size of each stable community for which there were at least thirty members and, by word count, the five largest laws or regulations in that community for 2007. I report the same results, but for 2023, in Appendix B. In Figure 6, I present maps for both the 2007 and 2023 consolidations, but with the different regions or genres of law manually labeled.

This new perspective on Canadian law is both familiar and novel. In the centre of the map we find those genres of law where the interface between the individual and the state is the greatest: federal public law, administrative law, and criminal law. To the left, we find regulatory regimes that focus on public safety and protection, environmental protection, and resource management. At the bottom, there are genres of law focused on land and, more recently, Indigenous self-government. Finally, on the right, genres of law that deal with monetary and fiscal matters: tax, banking, insurance, benefits provision, the provision of veteran benefits, and federal service employment benefits.

Comparing the 2007 map to the 2023 map, we now can better understand some of the changes in structure. Along the right, import/export and excise tax law, income tax law, banking and insurance, federal public service employment entitlements, and veterans' law have each maintained their basic shapes and relationships with each other. These are zones of basic stability that relative to the rest of law, have kept their basic shape.

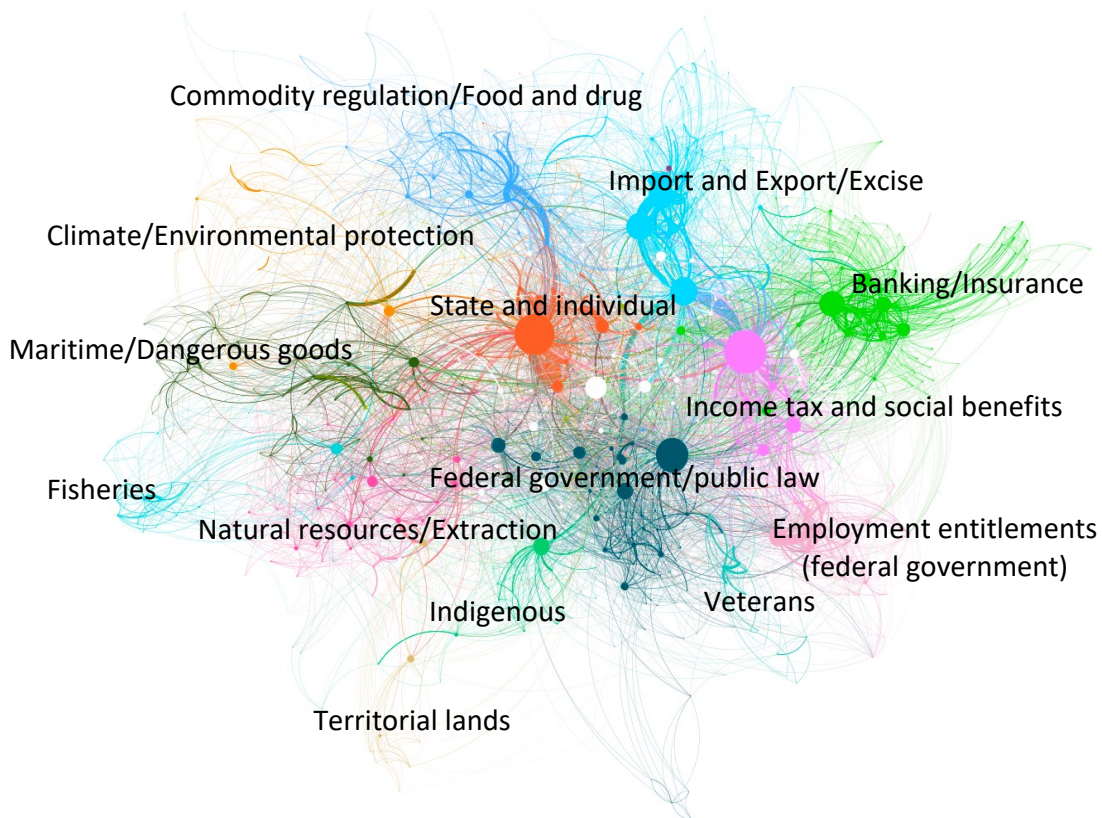
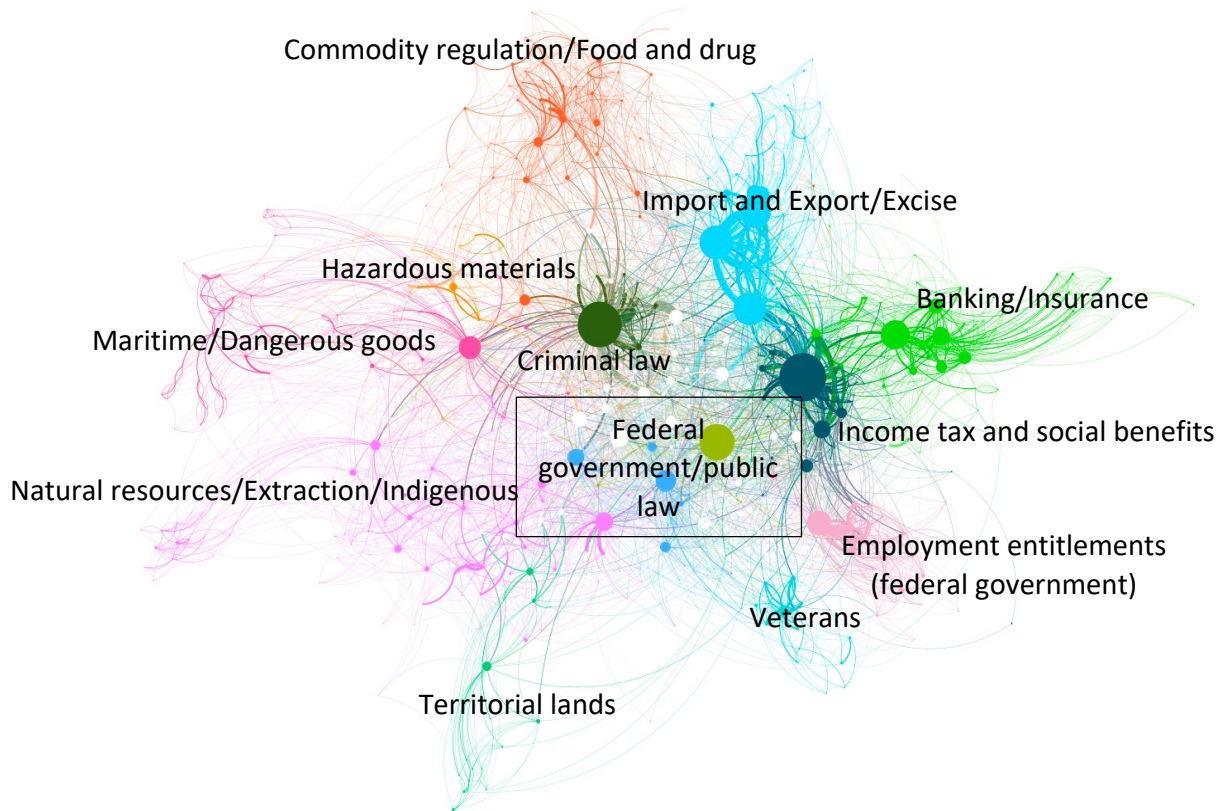


Figure 6: Community structure of Canadian law (top: 2007; bottom: 2023).

In the middle of the map, those areas that deal with the federal government, federal government policy, and the criminal law, we can observe clear consolidation. Those laws that deal with the establishment of ministries, the setting of government policy, and broad regulation all, by 2023, merged into a single large community. It is interesting that veterans' law, a genre of law that primarily concerns the provisions of benefits to veterans, has during this time drawn closer to the rest of standard federal government administrative law. Above the federal government, the community that in 2007 was concerned primarily with criminal law, expanded to include other, quasi-penal, law. By 2023, immigration law, citizenship law, and quarantine law were all grouped with criminal law. In both maps, we sense how important penal law is to other genres given that, more than other regions, there are linkages across the map—especially to those zones that are more regulatory. Along the bottom and to the left, the more regulatory spaces, there is a trend towards differentiation. Fisheries law, part of natural resources law in 2007, looks like a distinct genre in 2023.

We also see by 2023 two new genres of law, neither of which were present on the 2007 map. In the upper left, nestled between commodity regulation and hazardous goods regulation, we can observe new environmental and climate regulations. Where, in 2007, motor vehicle safety, occupational health and safety, and hazardous products regulation occupied a small space, by 2023 they are joined by a whole set of regulations and acts concerned with airborne pollutants, greenhouse gases, and environmental protection. Questions of proximity and distance become relevant because, this map suggests, that modern environmental law is more closely related to prescriptive regulatory regimes than, for example, older planning and project approval natural resource law, located further below. The second new genre, Indigenous self-government, found in between the federal government community, territorial lands, and natural resources, speaks to the proliferation of new laws and regulations that implicate natural resource developments that implicate Indigenous land and Indigenous self-government laws.

Different regions of the map are either more statutory or more regulatory. In the centre and to the right, genres of law feature more statutory words; while for each genre beginning at fisheries regulation and circling up to food and drug and commodity regulation is, when the words are counted, at least 80% regulatory. With a division of

Canadian law, it became possible to estimate the level of change within each community. Because communities sometimes divided and merged during the period under study, it was not possible to directly compare communities between years. To estimate the extent of change and growth within communities, I calculated the proportion of words in each 2023 community that were also present in 2007 law. To do this, I compared every section of law (excluding schedules) between the 2023 and 2007 consolidations and measured the change in the number of words between these two years. I then calculated, for each genre of law, the proportion of words in the 2023 cluster that were also present in the 2007 cluster. This method (Table 4) illuminates the comparative stasis and dynamism of different fields. Veterans' and public service employment law, for example, each maintained a core of 85% of the same words over the period under study, while the territorial lands, maritime and dangerous goods, and climate protection genres are today comprised more with words that were not a part of law in 2007.

2023 Genre	Estimated proportion of words present in 2007 consolidation	Genre's proportion of statutory words
Veterans	88.3%	65.3%
Public Service Employment	85.6%	49.6%
Fisheries	74.0%	70.5%
Banks, Insurance, and Credit	72.0%	68.1%
Income Tax and Benefits Provision	70.2%	8.8%
State and the Individual	67.7%	63.0%
Indigenous	65.3%	55.6%
Copyright and Broadcasting	63.1%	52.5%
Import/Export and Excise	56.0%	70.8%
Federal government/Public policy	55.7%	50.7%
Commodity and Food and Drug Regulation	52.6%	18.2%
Natural Resources/Extraction	52.1%	16.4%
Territorial lands	49.2%	56.6%
Maritime and Dangerous Goods	39.5%	14.0%
Climate/Environmental Protection	33.9%	11.9%

Table 4: The proportion of words in 2023 law that were also present unchanged in 2007 law and the proportion of each legal genre's words that were statutory.

2.4.3 Looking back at the micro

As I worked on this project, I developed a sense that the different genres of law followed somewhat different grammatical constructions. In large part, this was due to my own

training as a lawyer. Without too much difficulty, I can distinguish between a section of regulation and a section of statute taken out of context. Similarly, a section of criminal law feels different than a section of tax law, which feels different than a section of banking law. Given that the method of mapping and then partitioning law based on networks proved effective, I wondered if my intuition about the different grammatical feel of different genres of law would be borne out.

To investigate this matter, I used the open-source library spaCy³² to examine each section of law and count the number of determiners, proper nouns, nouns, verbs, adjectives, adverbs, auxiliary words, coordinating conjunctions, subordinating conjunctions, and particles. Then I calculated the word frequency, per genre, for each of these types of word. In essence, I stripped the content from all the sentences in each genre, preserving only the most basic grammatical information.

The resulting data were too complex to be interpreted. To make it more manageable, I used a machine learning technique known as principal component analysis³³ to reduce the dimensionality of the text. In simple terms, this means that I took a large data frame with many columns and, by calculating the extent to which data moved together (i.e. moved in a positive or negative direction together), reduced the data to two columns. This method necessarily means some of the noise and complexity of the data is dropped, though much can be preserved. In this case, I preserved 82% of the variance in the data even though I moved from 12 columns to 2. Now with two columns of data, I could visualize the data on a two-dimensional grid. Finally, to help understand the data, I used the K-Means algorithm³⁴ to group similar data points together. This simple algorithm attempts to cluster data by finding the data points that are closest to each other.

The resulting model was striking (Figure 7). Just by tallying and calculating the proportion of different word types, the data clustered in a manner that was reminiscent of

³² Matthew Honnibal & Ines Montani, “spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing” (2017).

³³ I used the implementation developed for F Pedregosa et al, “Scikit-learn: Machine Learning in Python” (2011) 12 *Journal of Machine Learning Research* 2825.

³⁴ *Ibid.*

the maps produced by interpreting the connections between statutes. Regulatory-heavy sections, the fringe of the original maps, clustered (in blue) together when considering

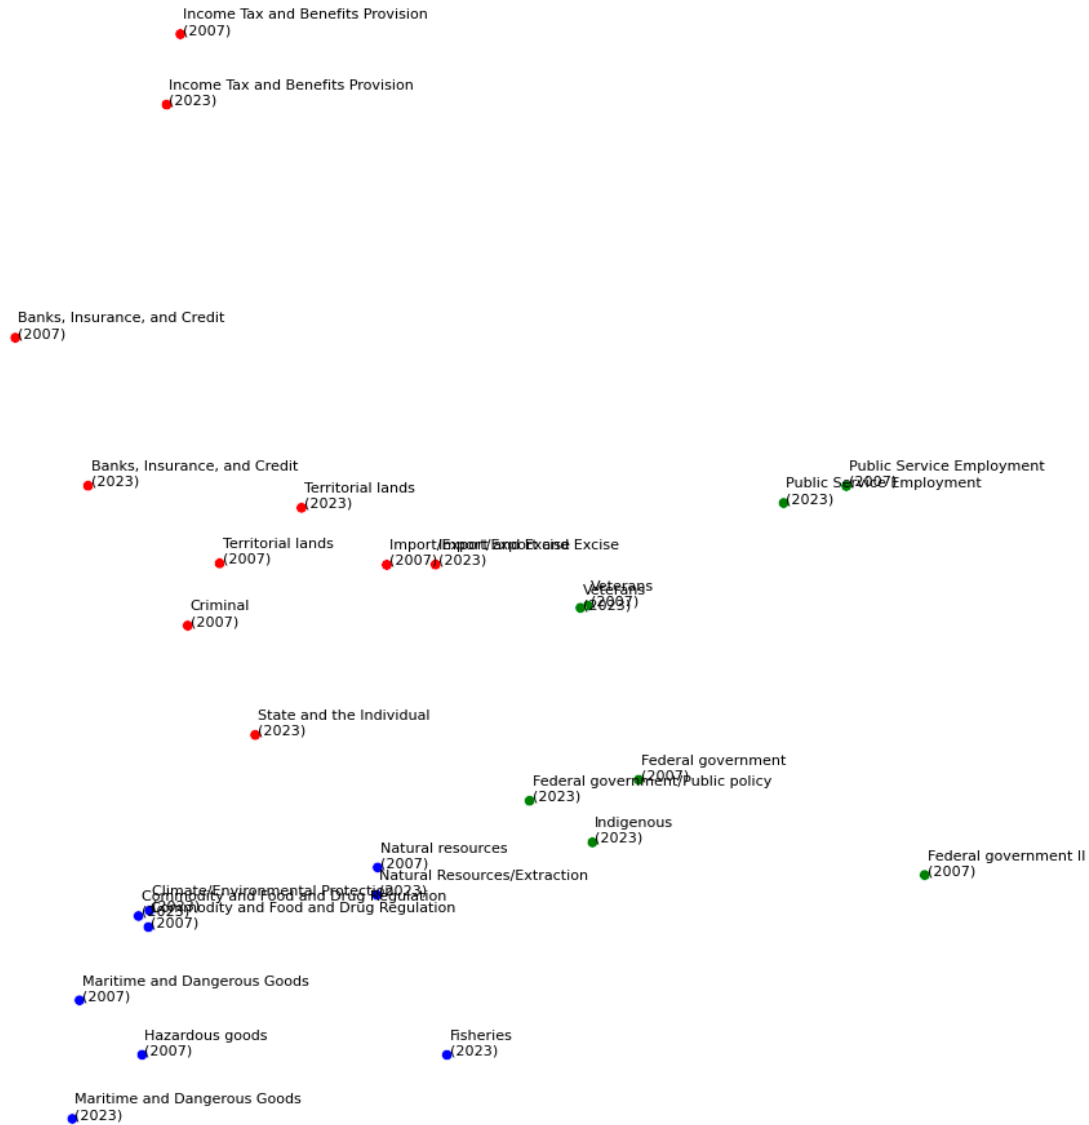


Figure 7: Classifying and comparing the grammatical structure of different genres of law.

grammar. Those sections found at the bottom right that primarily dealt with the federal government again clustered together (green). And then those areas that tended towards the upper right clustered together in red. The mapping is, to be sure, not perfect but it points to how the structure uncovered at the network level replicates, at least to a degree, the deep grammars of the sentences of each genre's law.

At the same time, we can spot a general drift in the language of law. Most genres moved, between 2007 and 2023, downwards while staying in relatively the same position with respect to the x-axis. But a few stayed roughly in the same spot or, in the case of the territorial lands genre, moved up. Determining what this means is challenging because, as discussed above, the x and y axis are amalgams of a dozen different columns. The more important insight, however, is that this method suggests that the drift of law is regulatory: most different genres, moved, from 2007 to 2023, closer to the regulatory corner of the map.

2.5 Discussion and conclusions

This study presents Canadian law in a new and unexpected way. Laws, like all texts, are designed to be closely read. Lawyers and legal researchers examine them for their content and their meaning and, from those close readings, build out arguments. This project took a different tack. Instead of reading law closely, it read it from a distance. Here, the object was not to develop a better analysis of a text—to better understand what factual intervention a discrete law sought to make in the world—but to attend to law's sizes, its interconnections, and its forms.

At bottom, this is a story about stability and instability. By every quantitative measure, Canadian law is a dynamic text. In just sixteen years, it grew almost 1.3 times, adding almost 4 million words. More than that, it renewed: comparing the law in 2023 to 2007, most sections of statutes and regulations were either new or amended in some way during that sixteen-year period. All told, Canadian law is an extraordinarily large text.

This scale precludes conventional systemic and holistic analyses. But comparing discrete statutes, networks, and years yields empirical proof of intuitive suppositions. There are, for example, grammatical changes afoot in legal sentences. It is meaningful that the law in 2023 contains proportionally more words that suggest conditional sentences or descriptive and declarative sentences. What might this mean? Here, further study is required but I can suggest a hypothesis. Law, it appears, is trending towards a more prescriptive, regulatory disposition. More and more terms and circumstances are, I

suspect, being defined, with more and more exceptions being identified. Indeed, it strikes me as significant that, when I compare the forms of different genres of law, those that changed the most were largely regulatory. This is hardly proof, but a suggestion of a drift that warrants further exploration.

But the most interesting conclusion here concerns the maps. Canadian law features enough explicit interconnections between laws and regulations to allow for a compelling spatialization of law. And once spatialized, genres and zones of law come into view. If we each intuitively appreciate that there are different genres of law, it is remarkable to see that an algorithmic approach uncovers a structure we assume but cannot necessarily describe. Of course, this algorithmic analysis is neither the first nor last word regarding the different genres of law (especially because it does not consider how, for example, other branches of government, other institutions, or other actors may organize their analyses of law), but it helps us understand how the decisions of drafters to cite different laws might indicate law's deep structure.

Watching these communities grow, change, and emerge helps to tell an empirical story about law. Law is objectively more complex and more interconnected, but we can also spot new zones of law that speaks to new social concerns. The emergence of autonomous emissions law and Indigenous self-government law speaks to how, in the past 16 years, the federal government has adopted new rules to address social concern. The general growth of regulatory zones speaks to an increasing tendency to prefer precise definition. The stability of revenue raising zones shows, as we would expect, that there was not a massive change in the way government raises or manages revenue streams.

As an administrative law lawyer, however, my eye keeps getting drawn to the centre of the map. I wonder if here we see the most interesting story. In law schools, in texts on law, public law is sub-divided into many fields. Criminal law is its own class, as is immigration law, as is natural resources law, as is administrative law. But, at least from the perspective of statutes and regulations, it looks as if there is a consolidation at work. Will it, I wonder, always make sense to speak of a distinct criminal law or, as the linkages between laws that implicate the use of state coercive power increase in number, will it become better to think of a penal core? To see an algorithmic tool find, over sixteen years,

a consolidation around two poles—the laws that describe the operation of government and the provision of service merge into one core, and penal laws resolve into another—makes me wonder whether there is a larger drift here. The law of veterans benefit provision is noticeably closer in 2023 to the federal government core than it was in 2007. Will it, too, merge? Perhaps these statutory and regulatory movements, detectable with math, tell us about a drift that one day soon will just be our new common sense.

2.6 Appendix A: Genres of law (2007)

Community/Legal genre	# of words	# of laws
Income tax and social benefits	2,252,245 (16%)	200 (5%)
Natural resources/Extraction/Indigenous law	1,268,767 (9%)	238 (6%)
Import and export/Excise tax	1,167,887 (8%)	387 (10%)
Banking, insurance, and credit	1,150,598 (8%)	260 (7%)
Commodity regulation/Food and drug regulation	1,089,476 (8%)	145 (4%)
Maritime/Dangerous goods	1,026,182 (7%)	120 (3%)
Criminal law	732,068 (5%)	88 (2%)
Employment entitlements (federal government)	348,814 (2%)	91 (2%)
Hazardous goods	244,443 (2%)	43 (1%)
Federal government	227,770 (2%)	96 (3%)
Territorial lands	176,598 (1%)	58 (2%)
Copyright and broadcasting	141,617 (1%)	41 (1%)
Veterans	139,259 (1%)	36 (1%)
Financial administration	99,331 (1%)	40 (1%)

2.7 Appendix B: Genres of law (2023)

Community/Legal genre	# of words	# of laws
Income tax and social benefits	2,769,776 (15% of all words)	231 (5% of all laws)
Banking, insurance, and credit	1,476,433 (8%)	339 (7%)
Import and export/Excise tax	1,457,928 (8%)	462 (10%)
Maritime/Dangerous goods	1,308,992 (7%)	107 (2%)
State and the individual (criminal)	1,246,916 (7%)	207 (4%)
Commodity regulation/Food and drug regulation	1,157,861 (6%)	157 (3%)
Federal government	922,879 (5%)	228 (5%)
Natural resources/Extraction	811,279 (4%)	161 (3%)
Climate/Environmental protection	704,417 (4%)	82 (2%)
Employment entitlements (federal government)	416,414 (2%)	115 (2%)
Fisheries	397,744 (2%)	59 (1%)
Indigenous law/self-government	262,630 (1%)	95 (2%)
Copyright and broadcasting	187,872 (1%)	63 (1%)
Veterans	166,758 (1%)	41 (1%)
Territorial lands	158,579 (1%)	67 (1%)

3. Speaking Like a Judge

Using Artificial Intelligence to Empirically Assess Judicial Speech in Supreme Court of Canada Hearings by Language Spoken and Gender of the Speaker

Americans know a lot about SCOTUS oral hearings. For ten years, Justice Clarence Thomas never spoke.¹ Justices who appear more responsive to public opinion use more politically charged language.² Women judges are more frequently interrupted than justices who are men.³ Judges use humour to “signal their power” and “direct their most humorous comments at the advocates with whom they disagree, the advocates who are losing, and novice advocates.”⁴ Justice Scalia’s courtroom comments were more tweetable than Justice Alito’s.⁵ The shift to telephonic hearings during the early days of the COVID-

¹ Adam Liptak, “Clarence Thomas Breaks 10 Years of Silence at Supreme Court”, *The New York Times* (29 February 2016), online:

<<https://www.nytimes.com/2016/03/01/us/politics/supreme-court-clarence-thomas.html>>.

² Noah Bergam, Emily Allaway & Kathleen McKeown, “Legal and Political Stance Detection of SCOTUS Language” (2022) arXiv, 2211.11724.

³ Adam Feldman & Rebecca D Gill, “Power Dynamics in Supreme Court Oral Arguments: The Relationship between Gender and Justice-to-Justice Interruptions” (2019) 40:3 *Justice System Journal* 173.

⁴ Tonja Jacobi & Matthew Sag, “Taking Laughter Seriously at the Supreme Court” (2019) 72 *V and L Rev* 1423 at 1423.

⁵ Jack Metzler, “Most Tweetable Justice: An Empirical Study” (2015) 18 *Green Bag’s Micro-Symposium on Supreme Court Rankings*.

19 pandemic changed the nature of judicial interventions and empowered the Chief Justice.⁶ Judges use information obtained during oral arguments to begin building coalitions for opinions.⁷ The words, and quantity of words, uttered by a judge is often a predictor of how they will decide the case.⁸

What is known about Supreme Court of Canada (SCC) oral hearings? Not much. A core reason for the knowledge differential, however, is clear: Americans have access to Supreme Court transcripts and Canadians do not.⁹ While the American court lets people listen to its hearings and publishes transcripts of each oral hearing, the Canadian court only provides a live video stream. And for research purposes, audio is notoriously difficult to work with. In general, people can read twice as fast as people speak and reading comprehension is marginally better than listening comprehension.¹⁰ And, perhaps more to the point, from a computational or quantitative research methods perspective, there is a robust toolset of text analysis tools which do not yet exist for audio.¹¹ American researchers can use more advanced text analysis techniques because they simply have the text to work with.

This means that easy access to transcripts facilitates, in its own way, access to courts and non-textual law in ways that disseminating judgements and opening physical courtrooms cannot. But transcripts cost. Producing a transcript is a labour-intensive task that requires significant skill on the part of the transcriber. A transcript of a single day of hearings might cost close to one thousand dollars or, in terms more familiar to scholarly

⁶ Eve M Ringsmuth et al, “SCOTUS in the time of COVID: The evolution of justice dynamics during Oral arguments” (2023) 45:1 Law & Policy 66.

⁷ Ryan C Black, Timothy R Johnson & Justin Wedeking, *Oral Arguments and Coalition Formation on the U.S. Supreme Court: A Deliberate Dialogue* (University of Michigan Press, 2012).

⁸ Lee Epstein, William M Landes & Richard A Posner, “Inferring the Winning Party in the Supreme Court from the Pattern of Questioning at Oral Argument” (2010) 39:2 The Journal of Legal Studies 433.

⁹ Of course, there are other factors at play: the different media landscapes between the two countries, public interest in apex Court hearings, and so on.

¹⁰ Victor Kuperman et al, “A lingering question addressed: Reading rate and most efficient listening rate are highly similar” (2021) 47:8 Journal of Experimental Psychology: Human Perception and Performance 1103.

¹¹ Justin Grimmer, Margaret E Roberts & Brandon M Stewart, *Text as data: a new framework for machine learning and the social sciences* (Princeton: Princeton University Press, 2022).

research, a fair chunk of a SSHRC budget. As Elizabeth Sheehy explained in *Defending Battered Women on Trial*: researchers must accept that transcripts “require lengthy waits and are extremely expensive to acquire.”¹²

This leads to my first intervention. In the past twenty months, artificial intelligence technologies have dramatically improved and democratized. Many of us interact with our phones and computers using our voices and, it turns out, most of the underlying technology is open-source that can be used—not just by Big Tech—by any of us with access to a reasonably powerful computer. To be sure, there is a learning curve but with just a little bit of coding and know how, but it is possible for legal researchers to use the same tools that large corporations sell to us but at virtually no cost. In this chapter, I show how we might use artificial intelligence to solve the transcript problem by developing my own artificial intelligence transcriber which I then use to make transcripts of oral arguments of the 2021-2022 SCC term.

That methodological point, however, is in service of a narrower and more traditional research aim. Just like the American researchers I canvassed at the beginning of this paper, I am interested in better understanding the functions and nature of oral argumentation before the SCC. In the absence of a major body of Canadian empirical work about what happens in SCC hearings, I aim to answer two relatively straightforward quantitative questions. First, who speaks more: judges who are men or judges who are women and, second, in what language do judges speak? My main answers: judges who are men speak much more than judges who are women and judges who are not from Quebec barely speak French in Court at all.

As I write, the government is accepting applications for an appointment to the Court to fill Justice Brown’s vacant seat. One of the job requirements set out by the government for applicants is that each be functionally bilingual while an identified institutional need of the court is ensuring that its composition reflects the diversity of Canadian society. In recent years, some commentators have suggested that there is some

¹² Elizabeth A Sheehy, *Defending Battered Women on Trial: lessons from the transcripts* (Vancouver, BC: UBC Press, 2014) at 14.

tension¹³ between these goals: insisting on bilingualism might, for instance, prevent otherwise excellent and qualified candidates—who would address diversity concerns on the bench—from applying¹⁴, especially when it is unclear what level of bilingualism counts as functional.¹⁵ On the other hand, judicial bilingualism is unquestionably important to the federation.¹⁶ I do not propose to resolve these tensions here but I hope to empirically enrich our discussions about bilingualism at the Court by suggesting that, at least in pure quantitative terms and at least with respect to participation in oral arguments, functional bilingualism remains elusive.

One last point: this chapter is for a non-technical audience, but I do hope to showcase some of the promise of computational methodologies. My view is that much computational legal scholarship is oriented towards the computer scientist, the linguist, or the economist. As such, those of more humanist impulses sometimes miss out on how new modes of quantification and computational assessment could advance our research. I therefore begin with a lay methodological discussion about how I drew on open-source tools to code an AI transcriber. Readers who are uninterested in this discussion can safely skip over the first part and meet me at the second part, where I describe the core study and answer my research questions. I end with a discussion about how computational methods can advance legal scholarship.

¹³ Marie-Ève Hudon, *Bilingualism in the Federal Courts* (Library of Parliament, 2011).

¹⁴ Alexandra Nasager, “The Supreme Court, Functional Bilingualism, and the Indigenous Candidate: Reconciling the Bench” (2019) 57 *Alta L Rev* 797.

¹⁵ Jean-Christophe Bédard-Rubin & Tiago Rubin, “The Elusive Quest for French on the Bench: Bilingualism Scores for Canadian Supreme Court Justices, 1985–2013” (2022) 37:2 *CJLS / LRCDS* 249 and Jean-Christophe Bédard-Rubin & Tiago Rubin, “Assessing the Impact of Unilingualism at the Supreme Court of Canada: Panel Composition, Assertiveness, Caseload, and Deference” (2018) 55 *Osgoode Hall LJ* 713.

¹⁶ Matthew Shoemaker, “Bilingualism and Bijuralism at the Supreme Court of Canada” (2012) 35:2 *Canadian Parliamentary Review* 30 and Sébastien Grammond and Mark Power, “Should Supreme Court Judges be Bilingual?” in Nadia Verrelli ed., *Democratic Dilemma: Reforming Canada’s Supreme Court* (McGill-Queen’s Press - MQUP, 2013).

3.1 Building an AI transcriber and transcribing the 2021-2022 SCC term

Behind this article, lies the larger research objective of making an accessible computational transcriber. The reason for ambition is simple: I believe that transcripts are an invaluable, but often out of reach, research tool. In Ontario, a professional transcriptionist charges at least \$6.30 a page.¹⁷ For anyone who wants to work with more than a few hours of audio, the cost can quickly become prohibitive. Moreover, because most Courts require parties to submit human certified transcripts for appeals, and because appellate courts are busy, there is a deficit of human capacity in the labour market available for legal researchers. While there are private corporations that provide automatic computational transcripts, the quality of these transcripts can be variable and they often do not get produced in the form lawyers and legal researchers are used to working with (they might not, for example, separate different speakers) and there are perennial concerns about privacy and data ownership.¹⁸

Researchers, I think, intuit that the open-source computational community might have something to offer them (especially given, for better or for worse, that major players in the tech industry – Alphabet, Meta, Microsoft, OpenAI, Amazon – often make some of their most exciting software freely available) but are turned off by the challenge of learning how to read and write code. Fair enough. The good news in this regard is that many law faculties are beginning to deliver computationally informed curricula and schools are graduating more and more computationally literate lawyers each year. Perhaps it is too much to ask law professors to code, but they can now find, provided the right tools are available, a research assistant who can. At least, this is the idea behind obiter.ai: make it easier to access the tools, and researchers might be able to use them. With this background

¹⁷ O Reg 145/22: Fees for Court Transcripts.

¹⁸ See for example “My journey down the rabbit hole of every journalist’s favorite app”, (16 February 2022), online: *POLITICO* <<https://www.politico.com/news/2022/02/16/my-journey-down-the-rabbit-hole-of-every-journalists-favorite-app-00009216>>.

in mind, I turn now to the technologies that I drew upon to make an artificial intelligence transcriber.

3.2 A gentle introduction to neural networks

What does a human transcriber do to produce a transcript? In a basic sense, their role is to translate data from one medium (sound) to another (text). Moving down one level of abstraction, we can appreciate that a transcriptionist is skilled at least three separate tasks: speaker diarization (who speaks when?), speaker recognition (who is speaking?), and speech recognition (what are they saying?). A high-quality transcript not only notes when a new speaker starts speaking, but can identify each speaker, and then represent each speaker's speech as text. Each of these tasks is challenging in its own right. Some speakers sound the same, sometimes speakers speak over each other, some people speak clearly, some do not.

Up until a few years ago, the task of accurately transcribing spoken language into written text was beyond the capabilities of computers. The diverse and dynamic nature of speech presented intricate challenges that could not be effectively tackled using traditional programming methods. However, recent advancements in the realm of computer science, particularly in the development of expansive machine learning systems known as neural networks, have significantly amplified the abilities of computers to decipher and interpret complex patterns, including those found in human speech.

A neural network operates on sophisticated mathematical principles and learns from data it is exposed to. Imagine that we needed to develop a model that can differentiate between images of apples and bananas. We would first need a dataset of digital images of these fruits where each image was properly labeled (as an apple or a banana). Because every digital image is a series of numerical values, these values can be fed into a math formula. A neural network is a type of formula composed of many interconnected variables, also known as neurons or weights. Each bit of the data from the image is pushed through the neural network where each node multiplies the data by some amount, before

passing the data to the next neuron. This sequence of operations continues across layers in the network.

At the end of these computations, the network arrives at an output. This output is a numerical value which the network uses to make a prediction (i.e “1” for banana and “0” for apple). But, and this is the important part, when the neural network is first used it will certainly be wrong. The initial “weights” were selected randomly because the network has yet to learn what mathematical representation an apple or a banana might have.

The labels or answer keys associated with each image therefore play a critical role in improving the network’s predictions. Whenever the network misclassifies an image (which we figure out by comparing the prediction to the answer key), the computer adjusts the weights slightly to push the output towards the correct answer. By adjusting the weights over a huge number of iterations, the neural network “learns” the distinct features of apples and bananas. After enough training, the neural network, evolves into something capable of completing its core task: distinguishing, in this case, between apples and bananas.

As algorithms have become more impressive, and compute costs have declined, researchers have developed impressive applications. One particularly impressive innovation is known as the transformer. In 2017, researchers from Google released a paper entitled “Attention is All You Need” that proposed a new mechanism for designing neural networks for machine translation tasks.¹⁹ This new mechanism has proved capable of recognizing patterns in all sorts of data, including speech data or unstructured text data. This said, the core functionality is not so different from the type of apple and banana detector described above. At bottom, each AI approach functions in the same way: using a mass of data to build and tune a massive and complex algorithm that we call a model.

Many (but certainly not all) of these models are, in turn, released to the public for use under permissive open-source licenses.²⁰ While it is beyond the scope of my project

¹⁹ Ashish Vaswani, et al, “Attention Is All You Need” (2017) arXiv, 1706.03762.

²⁰ The main public repository for artificial intelligence and machine learning models is, at the time of writing, hosted by the private company *HuggingFace*. At the time of writing, *HuggingFace*

here, it is important to critically think through why large corporations might release the products of their work to the public for use. In the main, large models require massive amounts of data and computer power to train—well beyond the resources of most scholarly research labs. Indeed, modern artificial intelligence models and other technologies can be thought of as industrial-level projects because the energy required to make them and run them can have measurably large carbon footprints.²¹ For many big technology firms (and here I am thinking of Amazon, Microsoft, and Alphabet) it will be in their financial interest to release and improve open-source software so that they can sell ancillary goods, such as hardware or hosting services.²² Put differently, large corporations might be incentivized to release digital technology for public use, not to contribute to the economy of public goods, but to build a market for their core offerings or to increase their prestige within the open source community.

Leaving these larger market dynamics to one side, by 2022, researchers in several domains independently created and open-sourced technologies that individually could approach, meet, or exceed the human capability of the three skills (diarization, speaker recognition, and speech recognition) required for transcription. Three technologies were of particular importance. First, in 2022, OpenAI (the company behind ChatGPT) released Whisper: a neural-network framework, that along the lines proposed in Google's 2017 paper, allowed for the transcription and translation of audio into text. Whisper, which can be run on a reasonably powerful personal computer, was trained on 680,000 hours of audio and associated transcripts from multiple languages. After a lengthy and compute intensive process, the neural network proved capable of recognizing words and sentences in sound and making equivalent textual representations.²³

hosted over 250,000 models available for download: HuggingFace, “Hugging Face – The AI community building the future”, online: <<https://huggingface.co/>>.

²¹ Emily M Bender et al, “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜” (2021) FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency 610.

²² Avi Goldfarb & Catherine Tucker, “Digital Economics” (2019) 57:1 Journal of Economic Literature 3 at 13.

²³ Alec Radford et al, “Robust Speech Recognition via Large-Scale Weak Supervision” (2022) Arxiv, 2212.04356.

Second, Pyannote, a 2019 software package released by a French academic lab, allows for speaker diarization by recognizing when in an audio file a new voice is speaking while, after separating each speaker segment, clustering each individual speaker together. Together, these functions allow for an understanding of who speaks when because Pyannote can recognize when people speak and recognize identical speakers.²⁴

Finally, even though Pyannote can determine when sound segments contain the same speaker, it cannot identify the speaker. Here, the SpeechBrain collection of tools, made by a Montreal-based lab, allows for voice recognition or, more simply, the comparison of an identified audio sample against an unknown sample. This means that each speaker can be checked against labeled samples to see if, for example, they are a judicial speaker allowing for automatic identification.²⁵

When practitioners speak of open-source technologies, as I have here, they really mean that some technologists have developed code-bases and data that can be modularly and easily be incorporated into someone else's program. In some sense, coding of the sort that lawyers and legal researchers might be more interested in, is less computer science than it is crafting a recipe or set of instructions to follow. This means that it is easy to draw upon multiple specialized tools developed by others and piece them together in a pipeline for one's own purposes.

In practice, to make the transcriber for this project I developed a simple Python script that fundamentally just moved data between different open-source components, iteratively identifying and abstracting information from the data. In particular, the program:

1. Scraped the SCC website for the links to all webcasts;
2. Downloaded the video;
3. Separated the audio from the video file;

²⁴ Hervé Bredin et al, “pyannote.audio: neural building blocks for speaker diarization” (2019) arXiv, 1911.01255.

²⁵ Mirco Ravanelli et al, “SpeechBrain: A General-Purpose Speech Toolkit” (2021) arXiv, 2106.0462.

4. Diarized the speakers and made a data frame (a spreadsheet) explaining when each speaker spoke;
5. Split the audio file into chunks such that each chunk corresponded with each instance where someone spoke;
6. Transcribed each audio chunk and inserted the resulting text and language into the data frame;
7. Compared each audio segment to clips of each judge speaking in both French and English and, if the program found a match with a judge, inserting that identity into the data frame.

This left me with a spreadsheet file for each hearing. After that, so that I could manually verify that the program identified the correct judicial speaker, I wrote a small additional program that took screenshots of each speaker and, using an open-source tool designed for scientific publishing,²⁶ published a human-readable version of each transcript on the internet. These transcripts can be viewed at obiter.ai/scc/.

3.3 Analyzing and validating machine-generated transcripts of 51 SCC oral hearings

For the 2021-2022 term, the program found 51 oral hearings. This is fewer than the total number of cases that the Court heard public argument on (56 cases) because some matters were heard together, sometimes over multiple days.²⁷ The total length of the recordings was 6,639 minutes (or just under 111 hours). This means that the mean recording time was approximately two hours. Of the 56 cases that the Court heard oral argument on, 13 were from Ontario, 12 were from British Columbia, 11 were from Alberta, and 7 were from Quebec. 42 cases were classified as criminal (22 as of right and 20 by leave) and 14 were

²⁶ “Quarto”, online: *Quarto* <<https://quarto.org/>>.

²⁷ All the data were scraped from Supreme Court of Canada, “Supreme Court of Canada - Archived Webcasts” online: <<https://www.scc-csc.ca/case-dossier/info/webcasts-webdiffusions-eng.aspx>> Last Modified: 2022-12-14.

classified as civil. Given how the Court grouped some matters together, there were 37 sittings focused on criminal matters and 18 sittings focused on civil matters.

Summing all the words spoken in Court, by judges, lawyers, and (occasionally) Court staff, the machine transcriber estimates that over 960,000 words were uttered in Court during the 2021-2022 term. For reference, Hamlet (a play with a performance that can exceed five hours) is 29,551 words long. No transcription is, or ever will be, perfect and verifying a transcript for accuracy is both an art and science. Both humans and computers struggle, for example, to transcribe audio when people are speaking over each other. I additionally observed that the AI transcriber struggled when asked to transcribe a short, single word, utterance by a speaker. I hypothesize that this is because of how the program I wrote was designed: because I fed in small chunks of audio that were split to only contain one speaker, this means that the computer had no access to any contextual information surrounding the audio clip. In other places, it appears that the transcriber struggled when speakers overlapped or spoke over each other.

I also observed that in some cases the AI ‘hallucinated’²⁸ words that were not there. This phenomenon, which is well recognized in the scientific literature, relates to how the underlying neural network was trained. Sometimes the AI will, just as people do, draw on its previous experiences to make an educated guess about what was said and, sometimes, these educated guesses amount to significant errors. By and large, however, I found that the AI transcriber’s outputs approached the expectations for accuracy I had for human transcripts.

There were spots, however, where it failed catastrophically. Sometimes the transcriber appeared to get stuck in a loop and, like a broken record, repeat the same phrase over and over. In other places, it clearly chose the wrong word, misspelled names, or neglected to include punctuation. I did not correct these errors when I found them because I was unable to develop an effective method to identify each major error. Instead, I used the transcripts as the computer generated them and the reader may review each

²⁸ Ziwei Ji et al, “Survey of Hallucination in Natural Language Generation” (2023) 55:12 ACM Comput Surv 248:1-248:38.

individually to understand the extent to which they can safely be relied upon. In fairness to the machine, it is important to remember that all transcripts, including those made by humans, are necessarily partial and error-prone and will only ever give a researcher access to a “glimpse into what happened.”²⁹ Accordingly, instead of treating the transcriber’s outputs as gospel, I prefer to see the transcripts as useful guides to the oral hearing and as an evidentiary base from which to draw conclusions about larger trends.

There was one aspect of the data that I determined that I could effectively and usefully verify. Given that my interest for this project related to quantification and assessing the number of words spoken by different judges, I wanted to ensure that the transcriber correctly identified and recognized judicial speakers. Recall that the transcriber initially identified each speaker clip by comparing it with a pre-identified section of audio for each judge. To check whether these identifications were generally correct, I implemented a facial recognition algorithm. Since all the audio was drawn from a digital video posted on the SCC’s website, I could extract the image that was displayed contemporaneously with each clip of each speaker. I reasoned that the SCC video feed will usually (except in cases of quick back and forth exchanges) show the speaker who is speaking.

Using the FastAI library,³⁰ I trained a computer vision AI model (the facial recognition model) to identify each speaker and, once this model was ready, I wrote code to automatically compare its assessments of the speaker with that of the transcriber for sections of audio where a person spoke longer than five seconds. I manually reviewed each discrepancy and concluded that the AI transcriber misidentified a judicial speaker (i.e. said Justice Wagner was speaking when, in fact, Justice Brown was speaking) 9 out of the 2783 identifications and missed a judicial speaker approximately 40 times out of 3206 non-judicial speaker identifications. Put differently, this validation method suggests that the transcriber was accurate 5940 times out of 5989 instances.

²⁹ Elizabeth A Sheehy, *Defending battered women on trial: lessons from the transcripts* (Vancouver, BC: UBC Press, 2014) at 14.

³⁰ Jeremy Howard and Sylvain Gugger, “Fastai: A Layered API for Deep Learning” (2020) 11:2 Information 108.

To summarize, no transcript is perfect. Every transcriber, human or otherwise, will struggle when people speak over each other, microphones become muffled, or people's voices change. This uncertainty makes it impossible, then, to quantify exactly the extent to which any transcript is accurate. In the transcripts produced by the machine, there are errors, but there are indicia that suggest, on the whole, that the transcripts provide an acceptable base from which to work with for social scientific inquiry.

3.4 Who speaks, when, and in which language?

The Rules of the Supreme Court of Canada give litigants a right to address the Court in either official language.³¹ Until 1987, the Court convened an oral hearing for almost all matters, including applications for leave. An amendment to the *Supreme Court Act* that year authorized three member panels of the Court to decide leave applications on written materials alone.³² Today the Court hears oral arguments in virtually every case (aside from leave applications) that it decides. Lawyers for each side generally have an hour each to present their arguments and interveners usually have five to ten minutes each. The judges of the Court often interrupt a lawyer's presentation to pose questions or invite comments.

Little is known empirically about the function and effects of oral arguments. Emmet Macfarlane, in a comprehensive institutional study of the Court, suggests that the “emphasis on written material in the contemporary period is regarded as having diminished the importance of the oral hearing stage.”³³ Indeed, in an age when most of the court's work is done behind closed doors—reading judgements below, reading records, reading facta, deciding leave, discussing the case with clerks, discussing the case with colleagues, drafting opinions, revising opinions—it is worth asking whether oral arguments matter much at all.³⁴

³¹ *Rules of the Supreme Court of Canada*, SOR/2002-156, s 11.

³² Emmett Macfarlane, *Governing from the Bench: The Supreme Court of Canada and the Judicial Role* (UBC Press, 2013) at 78.

³³ *Ibid* at 99.

³⁴ Within the American literature, this is a recurring question; see as examples: Robert J Martineau, “The Value of Appellate Oral Argument: A Challenge to the Conventional Wisdom”

Canadian Supreme Court judges suggest that oral hearings, while important, impact outcomes only a small percentage of the time but that they remain important because judicial questions “allow counsel to focus specifically on what is troubling the justice and give them the opportunity to resolve it.”³⁵ As one judge explained to Macfarlane:

[t]he oral hearing is a useful exercise in a significant number of cases – not the majority – but even if it doesn’t change the Court’s opinion about which way a case should go, it may enlighten the issues further and may bring a better or different focus on the issues and sometimes that focus may be different from the focus that appeared from the factums.³⁶

Along these lines, there is American scholarship that, along these lines, suggests that oral hearings provide a rare opportunity for judges to introduce ideas and concepts not canvassed in the written record. One study found that 80% of questions posed by American judges, for example, referred to matters that were not in litigants’ briefs and that 70% of all policy questions were not explicitly raised in any party’s written materials.³⁷ This allows judges to “gain information beyond the briefs that they believe will help them make decisions in line with their preferred goal.”³⁸

Within the hearing, different judges interact with counsel and each other differently. Justice Bastarache explains “[t]here are the ‘get on with it’ judges, the patient judges, the chatty judges and the quiet judges.”³⁹ While there is no literature that empirically considers how Canadian judges behave and talk in the courtroom, there is

(1986) 72 Iowa L Rev 1; Warren D Wolfson, “Oral Argument: Does It Matter” (2001) 35 Ind L Rev 451; John G Roberts Jr, “Oral Advocacy and the Re-emergence of a Supreme Court Bar” (2005) 30:1 Journal of Supreme Court History 68; Timothy R Johnson, Paul J Wahlbeck & James F Spriggs, “The Influence of Oral Arguments on the U.S. Supreme Court” (2006) 100:1 American Political Science Review 99; Michael Duvall, “When Is Oral Argument Important - A Judicial Clerk’s View of the Debate” (2007) 9 J App Prac & Process 121; Lawrence S Wrightsman, *Oral Arguments Before the Supreme Court: an empirical approach*, (New York: Oxford University Press, 2008).

³⁵ Macfarlane *supra* note 32 at 98.

³⁶ *Ibid* at 94-95.

³⁷ Timothy R Johnson, *Oral Arguments and Decision Making on the United States Supreme Court* (SUNY Press, 2004) at 126.

³⁸ *Ibid*.

³⁹ *Ibid* at 97.

some analysis of gendered speaking patterns among counsel who appear before the Court. A recent independent study found that lawyers who are perceived as men, despite accounting for roughly half of all lawyers in Canada, address the SCC in oral argument approximately 75% of the time.⁴⁰

The gender gap regarding counsel who appear before the Court, however, appears to be narrowing. The Court includes in its online dockets the identities of each party and the membership of their counsel team. I wrote a small program to scrape these names from the Court website and developed a comprehensive picture of all lawyers who appeared before the Court since 1987. I then compared each lawyer's first name to the entries in the World Gender Name Dictionary 2.0⁴¹ to make predictions about the lawyer's gender. This showed that since the late 1980s, the gap between men and women counsel who represent a party (which is distinct from appearing for a party or speaking on behalf of a party) in a case that is finally appealed to the Court has narrowed.

⁴⁰ Arthur Peltomaa, "Gender and Racial Diversity of Counsel at the Supreme Court of Canada: An Empirical Study" (2021), online:

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3930191.

⁴¹ Gema Lax-Martinez et al, "Expanding the World Gender-Name Dictionary: WGND 2.0" (2023) World Intellectual Property Organization (WIPO) Economic Research Working Paper Series No. 64.




Figure 1: Predictions regarding the gender of each lawyer on a Supreme Court of Canada case.

Less is known about the languages used in oral hearings. There is empirical evidence that shows that a judge’s language, or Québécois origin, impacts Court operations. In a 2018 study, Jean-Christophe Bedard-Rubin and Tiago Rubin found that unilingual judges were less likely to be assigned to panels concerning cases coming from Quebec, they are less likely to write judgements on cases arising from Quebec, and that judges from Quebec end up writing most judgements concerning cases from Quebec courts.⁴² From the perspective of the government, oral proficiency in both official languages is important for appointment purposes. The 2023 invitation for applications to

⁴² Jean-Christophe Bedard-Rubin & Tiago Rubin, “Assessing the Impact of Unilingualism at the Supreme Court of Canada: Panel Composition, Assertiveness, Caseload, and Deference” (2018) 55 Osgoode Hall LJ 713.

fill the vacant Western Canada seat explains that “[i]deally, the judge can converse with counsel during oral argument and with other judges of the Court in French or English.”⁴³

With these background questions that help us orient to the problems of bilingualism and equity in the appointment process for judges in mind, I turn to the central questions of this study: who speaks, on what cases, and in which languages?

3.5 Judges who are men speak more than judges who are women

In the 2021-2022 SCC term, the Court’s oral hearings were over 51 hours long. As discussed above, the artificial intelligence transcriber concluded that 960,881 words⁴⁴ were uttered by all speakers over the course of the term. Of those words, the program estimated that 23% of all words (222,169 words) were spoken by judges. Importantly, the Court does not sit nine judges on each case. My dataset does not include information about which judges did not participate in a particular hearing. In future, it would be useful to develop a dataset of judicial participation to determine the extent to which judges participate equally in hearings.

Justice	Gender	Number of words	% of judicial total
Justice Moldaver	M	44957	20.235496
Justice Brown	M	36105	16.251142
Justice Kasirer	M	31379	14.123933
Justice Rowe	M	25352	11.411133

⁴³ Office of the Commissioner for Federal Judicial Affairs Canada, “Qualifications and Assessment Criteria”, online: <<https://www.fja-cmf.gc.ca/scc-csc/2023/qualifications-eng.html>> Last Modified: 2023-06-20.

⁴⁴ It is worth recalling that there are methodological limitations involved in counting words. For example, some people use more filler words (‘umm’, ‘hmm’, ‘ah’, ‘well’) than others. Do these count as words or not and, for transcription purposes, is it possible to either include or exclude all of them? I do not think so. This reminds us that transcription remains an art and, despite the real numbers produced by quantification, there a margin of ambiguity at play here. This said, my examination of the machine transcriber’s outputs shows that it rarely included filler words.

Justice Wagner	M	21478	9.667415
Justice Côté	W	18255	8.216718
Justice Karakatsanis	W	17977	8.091588
Justice Martin	W	14732	6.630988
Justice Jamal	M	11934	5.371586

Table 1: Words uttered by each judge (2021-2022 term).

To assess the gender identity of each judge, I examined their official biography and assumed that the pronouns used are an adequate proxy for their gender.⁴⁵ As Table 1 shows, the five most verbose judges were all men. Despite accounting for 66% of the bench, justices who were men accounted for 77% of all words uttered. Here Justice Jamal stands out as an outlier. It may be that as the newest member the Court he spoke the least. The clearest pattern in the data is the gendered split between the speakers: the three women justices all spoke roughly the same amount and ranked at the bottom of the list of loquacious speakers. Nor was the division small. Justice Moldaver, for example, spoke over 3 times as many words in Court as did Justice Martin.

Beneath this major trend, evidence of different judicial personalities on the bench emerges. Some justices spoke in longer chunks than others. As Table 2 shows, Justices Moldaver, Martin, and Kasirer spoke in longer chunks than their peers.

Justice	Average number of words per intervention⁴⁶	Total number of interventions
Justice Moldaver	102	441
Justice Martin	84	175
Justice Kasirer	84	374
Justice Jamal	77	156

⁴⁵ Supreme Court of Canada, “Supreme Court of Canada - Judges of the Court”, online: <<https://www.scc-csc.ca/judges-juges/index-eng.aspx>> Last Modified: 2023-06-12.

⁴⁶ For the purposes of this study, an intervention is defined as a chunk of speech that the machine transcriber attributes to a single judicial speaker.

Justice Rowe	76	335
Justice Karakatsanis	70	260
Justice Côté	65	283
Justice Brown	54	671
Justice Wagner	33	652

Table 2: Number of judicial interventions and average word count of judicial interventions

Justice Brown and Justice Wagner, in contrast, spoke in much shorter chunks. This may be a feature of their different roles. As the Chief justice, Justice Wagner often intervenes in the hearing to manage it by, for example, telling a party that they are out of time or inviting a party to begin their submissions. But for these managerial interventions, which are often quite short, I suspect that his total contribution to the number of words spoken would be lower. Justice Brown, who on average is the first to intervene after the Chief Justice, appears to be one of Justice Bastarache’s “chatty” judges who engaged with counsel frequently: the number of interventions he made was the highest of any member of the Court, while the average number of words used per intervention was the second lowest.

The judges of the SCC have different subject matter interests, but the overall gendered patterns persist across case types. As Table 3 shows, when civil and criminal cases are separated out, the overall picture changes.

Justice	Words spoken in criminal cases		Justice	Words spoken in civil cases
Justice Moldaver	33234		Justice Brown	13154
Justice Brown	22951		Justice Rowe	11726
Justice Kasirer	21203		Justice Moldaver	11723
Justice Wagner	16738		Justice Kasirer	10176

Justice Rowe	13626		Justice Karakatsanis	8500
Justice Martin	10092		Justice Côté	8419
Justice Côté	9836		Justice Jamal	5668
Justice Karakatsanis	9477		Justice Wagner	4740
Justice Jamal	6266		Justice Martin	4640

Table 3: Comparing words uttered in civil verses criminal matters.

3.6 Non-Québécois judges rarely speak French

Taking all the words uttered in the Court, English was spoken (by lawyers and judges) relative to French at a rate of 6.5:1.⁴⁷ The use of French in Court primarily turned on whether the case originated from Quebec. In the 2021-2022 term, there were seven cases from Quebec courts. Each of these the cases were criminal matters, five were heard because the Court granted leave and the other two came to the Court by way of right. There were no cases from Quebec that were heard concurrently with another matter.

When the Quebec cases were removed from the dataset, leaving only matters from the rest of Canada, French is spoken less than one percent of the time by judges and lawyers alike. In contrast, cases from Quebec feature both French (83% of words) and English (16% of words). Notably for a bilingual Court, as Figure 2 shows, there is a “missing middle” of cases in which English and French feature in oral argument evenly. As we can see, in most cases, English is the dominant language spoken in oral hearings while, in a smaller percentage, French is. There are no cases where English and French are spoken at roughly equivalent rates. Of course, this finding speaks to only the one year

⁴⁷ This ratio may be somewhat misleading because different languages and grammatical structures use words at different rates. Put differently, some languages, French compared to English for instance, are “wordier” than others. To test this, I took all the official judgements of the Supreme Court of Canada issued between 2010 and 2020, which are published in both official languages, and counted the number of words. With reference to this one decade alone, all of the French judgements were 1.1 times longer than all of the equivalent English judgements.

under study and, in other years, we may notice other trends. For example, we could reasonably hypothesize that cases from New Brunswick might feature a greater use of both official languages.

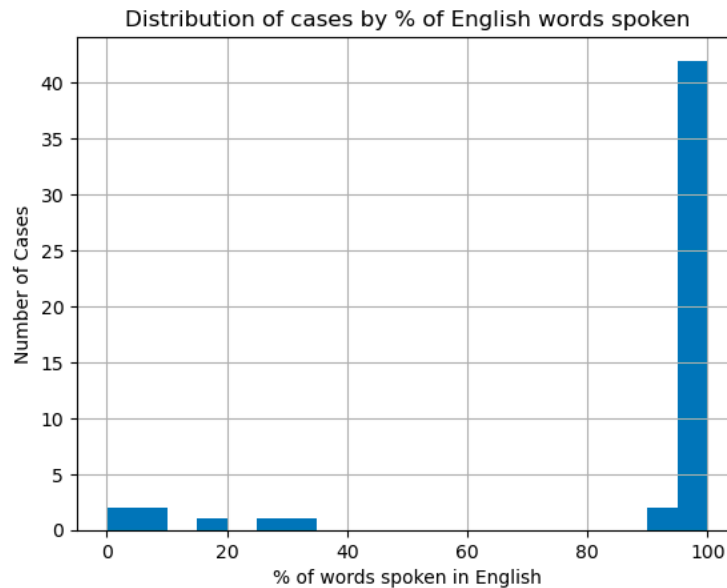


Figure 2: Distribution of cases by percentage of words spoken in English.

Officially, every judge on the SCC during the 2021-2022 term (except for Justice Moldaver) was functionally bilingual. Nonetheless, Table 4 shows that the Québécois judges spoke the largest share of French while the rest of the judges barely spoke or spoke not at all. This is likely attributable to comfort with another language as opposed to a judge’s specialty with Quebec law because, as noted above, each of the Quebec cases primarily heard in French concerned a *Criminal Code of Canada* matter as opposed to a provincial legal matter. This is an interesting point to pursue for future research because, while the Court’s francophones are currently all previously from Quebec, this was not historically the case. It would be interesting to see how, in different time periods and with different benches, these patterns present or do not present.

Judge	French words
Justice Kasirer	11278

Justice Wagner	8550
Justice Côté	6576
Justice Brown	427
Justice Jamal	262
Justice Martin	195
Justice Rowe	37
Justice Karakatsanis	0
Justice Moldaver	0

Table 4: Number of words spoken in French.

Examining the 2020-2021 cases from Quebec closely, a clear pattern is evident. The bulk of the hearing is conducted in French because the appellant and respondent make their arguments in French. During that phase of the hearing, the judges from Quebec ask most of the questions with other judges only occasionally interjecting. When, however, English language interveners begin to participate, the judges from the rest of Canada became more engaged.

3.7 The research opportunities and limitations of computational methodologies

The benefit of computational methodologies is that they allow new sorts of research questions and programs come into view. For sociologically minded scholars, this study shows two main benefits. First, digitization makes legal information more accessible. There is a tremendous amount of legal information out there, but much of it is practically out of reach for researchers. But we see here how modern technology can, for example, produce transcripts for functionally no cost—enabling all sorts of new approaches on research problems. Second, scale. Each scholar who studies the SCC (or any court for that matter) can only ever aspire to study a tiny fraction of the text and data produced for and by the institution. Accordingly, we must be strategic: we analyze a few canonical decisions, we read a few texts closely, we conduct some in-depth case studies, and so on. We do not

read legal texts or consider legal data at scale because how could we? But if we could, we would discover all sorts of new and interesting problematics for study.

New computational methods, sometimes referred to as methods of ‘distant reading,’ will help us sketch out a larger sociological picture of how legal institutions and cultures work, how legal ideas develop, and how influence is exercised. As the literary scholar Franco Moretti puts it “a field this large cannot be understood by stitching together separate bits of knowledge about individual cases, because it isn’t a sum of individual cases: it’s a collective system, that should be grasped as a whole.”⁴⁸ No, they will not replace research based on careful interpretations of core texts but supplement them by contextualizing information and arguments in a larger sociological milieu.

And, speaking only for myself, the most exciting aspect of this study is not what it suggests we might one day learn about the SCC, but what we might learn about all first-instance law. Bail, social benefits, small claims, evictions: now that we can develop affordable transcripts at scale, perhaps we can develop new ways of describing how the law functions in these important, large, and understudied fields. Perhaps, as we transform more legal information into bits and bytes, “big data could lead us back to big questions.”⁴⁹

But it is important to be restrained. This study asked narrow, preliminary quantitative questions, but its findings already point to important new areas for research. As I set out at the beginning, politicians, the profession, and the public remain seized by questions about the significance of bilingualism and the functional importance of a diverse bench. The data presented here do not answer these questions, but they at least show that there is a “there there.” In the term under study, it was judges from Quebec that did the heavy lifting on cases from Quebec. Perhaps this means that the federal government is not emphasizing the ‘functional’ part of ‘functional bilingualism’ enough when it assesses candidates for the top Court? Or perhaps the Court functions just fine and it is acceptable that those judges who are most capable using one language do end up using that language in Court.

⁴⁸ Franco Moretti, *Graphs, Maps, Trees: abstract models for literary history*, (London: Verso, 2007) at 4.

⁴⁹ Franco Moretti, *Literature, Measured* (Stanford Literary Lab, 2016) at 7.

Similarly, when all judges' interventions are totalled, there are clear gendered patterns in who speaks and for how long in Court. But what do these findings mean? It is hard to say. Perhaps, for example, women judges speak less because they are more concise and direct. Perhaps they are interrupted by their male colleagues. Perhaps they have more carefully read the written materials. Or perhaps there were idiosyncratic features of this term that led to this outcome. Quantitative and computational methods can pose questions, but it is qualitative methods—which I hope will be deployed regarding these matters one day soon—that can provide detailed and fulsome answers.

4. Persistent Discord: Inconsistent Adjudication of Terrorism Cases in Canada

This study finds that different adjudicators and different judges decide terrorism deportation cases differently. How big of a problem is this? “A lack of unanimity is the price to pay for the decision-making freedom and independence” of administrative adjudicators, the Supreme Court of Canada explained in *Domtar*.¹ But how much discord can the law bear? Parties reasonably expect that like cases will be treated alike and that outcomes will not turn on the identity of an individual decision-maker. The public expects the laws passed by legislatures to mean the same thing to litigants who appear in different places and at different times.

For these reasons, administrative bodies are empowered and directed to resolve internal conflicts.² Institutions may use tools and internal practises (distributing past reasons, holding plenary meetings and training, and using templates) to ensure coherence

¹ *Domtar Inc v Quebec (Commission d'appel en matière de lésions professionnelles)*, [1993] 2 SCR 756 at 800 (SCC).

² *Ibid.* See also, *Toronto (City) v CUPE, Local 79*, 2003 SCC 63 at para 51: “[I]f the result in the subsequent proceeding is different from the conclusion reached in the first on the very same issue, the inconsistency, in and of itself, will undermine the credibility of the entire judicial process, thereby diminishing its authority, its credibility and its aim of finality.” For a discussion on the obligation for administrative tribunals to address inconsistent decision-making see Joseph Robertson, “Administrative Deference: The Canadian Doctrine that Continues to Disappoint,” (18 April 2018) at 41-45, online: <canlii.ca/t/stvr> [perma.cc/VKC6-XEHJ].

and avoid conflicting results. When adjudicators depart from past practice, they must explain this departure in their reasoning. Ultimately, inconsistency may be acceptable, but it must be rare, principled, and considered.³ Where discord becomes persistent, reviewing courts “may find it appropriate to telegraph the existence of an issue in its reasons” and indicate that “it may become increasingly difficult for the administrative body to justify decisions that serve only to preserve the discord.”⁴

But what if neither the tribunal nor the courts notice the inconsistency? And what if the reviewing court’s own jurisprudence unwittingly fortifies the discord? This study of national security-based deportation orders finds that, between 2018 and 2020, both the Canadian deportation tribunal and the Federal Court of Canada treated members of the same alleged terror group differently, even though their cases were functionally the same. This study finds that the Immigration Division (‘the Division’ or ‘the ID’) of the Immigration and Refugee Board (“IRB”) finalized 73 deportation cases concerning members of the Bangladesh National Party (“BNP”), an opposition group and former governing party. In each case, the Canadian government alleged that the group used terror-based tactics to achieve political ends and all its members were therefore legally unwelcome in Canada. Even when the cases featured word-for-word identical evidence, some adjudicators always ordered deportation and some adjudicators did not. Would a particular member of the BNP be ordered deported? It depended on the identity of the adjudicator assigned to their case.

Through the mechanism of judicial review, the Federal Court of Canada is charged with supervising deportation decision-making to, amongst other functions, ensure the law’s coherence. During the period under study, the Court’s interventions entrenched inconsistency. Frequently, first-instance adjudicators recycled their reasons across different cases. This meant that some cases that appeared before the Federal Court of Canada had functionally identical sets of reasons and identical underlying records. Unfortunately, it appears that the Court did not realize that it was hearing many versions

³ *Canada (Minister of Citizenship and Immigration) v Vavilov*, 2019 SCC 65 at paras 129-132 [*Vavilov*].

⁴ *Ibid* at para 132.

of the same case. When different judges considered the same evidence and reviewed the same reasons, they reached different conclusions: sometimes the deportation order was properly issued, sometimes it was not.

This study raises rule of law concerns. The Supreme Court of Canada recently reaffirmed that “the rule of law breaks down where legal inconsistency becomes the norm and the law’s meaning comes to depend on the identity of the decision maker.”⁵ At least during the period under study, neither the Canadian public nor prospective deportees could be expected to understand the law of terrorism and deportation. Were members of the BNP deportable for security reasons? The answer was both ‘yes’ and ‘no.’

This conclusion is particularly unsettling when we consider the stakes. In almost every case examined here, the person was in Canada claiming refugee protection. In most cases, their claim was terminated after the Division concluded that they were deportable for belonging to a terrorist organization. Labelling a person a member of a terrorist group and terminating their claim for protection is a grave matter and this part of the law ought to be careful, not cavalier.

How did the discord develop and fester? My diagnosis directs our attention to problems occasioned by scale. When Parliament established the Immigration Appeal Board in 1966 (the predecessor body to the IRB), it capped the number of adjudicators at nine.⁶ After operating for a year, the tribunal issued a single slender volume of jurisprudence that was then distributed to each law school in the country.⁷ Today, the IRB is one of the world’s largest administrative tribunals, employing hundreds of decision-makers across a range of subject areas, and issuing tens of millions of words worth of decisions each year. No human and no court can be expected to synthesize and understand a jurisprudence of this scale: discordant signals get lost in a cacophony of noise.

When everything else is stripped away, this paper is first an empirical analysis of security-based deportation orders. This project began because I was curious, and this inspired my first research question, to learn how many people Canada sought to deport for

⁵ *Ibid* at para 71.

⁶ *Immigration Appeal Board Act*, SC 1966, c 90, s 3.

⁷ Immigration Appeal Board, *Immigration Appeal Cases: selected judgments* (Ottawa, 1969).

espionage, terrorism, or other national security reasons. I therefore begin with a high-level discussion of national security deportation law to orient the reader. In the following section, I outline my methodology and report on the study's findings. This part will be of particular interest to readers who are interested in the on-the-ground operation of Canadian national security deportation law. It was during this larger investigation that I developed serious concerns, the exploration of which inspired my second research question, about the consistency of decision-making regarding BNP cases. I turn to these administrative law concerns in my third section before concluding with recommendations.

Perhaps the most useful and practical remedial suggestions emerge, however, from this study's nature. This project was enabled by modern computational research methods, without which it would have been impossible to notice the connections between cases and across tribunals. We may not be able to read millions of words, but perhaps we can use computers to usefully organize legal data so that we can spot and surface inconsistencies and problems in jurisprudence that would otherwise go unnoticed. Ultimately, I encourage Courts and tribunals to supplement their traditional research toolkits with new computational methodologies.

4.1 Canadian national security deportation law

The *Immigration and Refugee Protection Act* (the 'IRPA') explains that all non-Canadians are either admissible or inadmissible.⁸ An admissible person may enter and stay in Canada and an inadmissible person may not. If a person already in Canada is inadmissible or becomes inadmissible, they must leave the country. If they do not leave on their own, the Canadian government (represented by the Minister of Public Safety and Emergency Preparedness) can initiate deportation proceedings.⁹

There are many bases for inadmissibility, ranging from non-compliance with immigration law to criminality to health reasons.¹⁰ Depending on the type of status a non-

⁸ *Immigration and Refugee Protection Act*, SC 2001, c 27, Part 1, Divisions 3 and 4 [IRPA].

⁹ *Ibid*, ss 4(2), 44.

¹⁰ *Ibid*, Part 1, Division 4.

citizen has (they may, for example, be a permanent resident, a temporary resident, or a protected person) they may be vulnerable to different types of inadmissibility. Regardless, all non-citizens are equally subject to national security inadmissibility law. Section 34 of the *IRPA* explains that all non-citizens are deportable for:

- (a) engaging in an act of espionage that is against Canada or that is contrary to Canada's interests;
- (b) engaging in or instigating the subversion by force of any government;
- (b.1) engaging in an act of subversion against a democratic government, institution or process as they are understood in Canada;
- (c) engaging in terrorism;
- (d) being a danger to the security of Canada;
- (e) engaging in acts of violence that would or might endanger the lives or safety of persons in Canada; or
- (f) being a member of an organization that there are reasonable grounds to believe engages, has engaged or will engage in acts referred to in paragraph (a), (b), (b.1) or (c).¹¹

The burden of proof for inadmissibility matters is attenuated, at least relative to criminal or civil matters. A person is deportable if evidence establishes that there are “reasonable grounds to believe” that the inadmissibility allegation is true.¹² The Supreme Court of Canada explains that this

standard requires something more than mere suspicion, but less than the standard applicable in civil matters of proof on the balance of probabilities. In essence, reasonable grounds will exist where there is an objective basis for the belief which is based on compelling and credible information.¹³ [citations omitted]

Unlike criminal or civil matters, the *IRPA* authorizes the government to initiate an inadmissibility proceeding for conduct and facts that “have occurred” or “are occurring” but also for conduct and facts that “may occur.”¹⁴ Put differently, a person can be removed

¹¹ *Ibid*, s 34.

¹² *Ibid*, s 33.

¹³ *Mugesera v Canada (Minister of Citizenship and Immigration)*, 2005 SCC 40 at para 114.

¹⁴ *IRPA*, *supra* note 8 at s 33.

from Canada if they are assessed as a future risk to Canadian security even if they have not committed a crime or engaged in identifiable problematic conduct.

Non-citizens who are members of organizations that have, are, or may engage in espionage, subversion, or terrorism are inadmissible. Noting that it is “trite to say that terrorist organizations do not issue membership cards,” the Federal Court concluded that the term membership ought to be given an “unrestricted and broad interpretation.”¹⁵ That said, the Federal Court reminds adjudicators that it is not appropriate “to classify anyone who has had any dealings with a terrorist organization as a member of the group” and that “[c]onsideration has to be given to the facts of each case including any evidence pointing away from a finding of membership.”¹⁶ For example, the Court says that it is usually not appropriate to find children inadmissible because of membership in a terrorist group or hold persons who were coerced to join terrorist groups responsible for their role in the organization.¹⁷

In some cases, the timing of a person’s decision to leave a problematic group may be relevant. Even though the *IRPA* makes people accountable for future acts, Justice Mandamin concluded that it is unfair to make people responsible for all future acts of organizations that they were ever members of:

If an individual joins an organization that is not engaged in terrorism or has not engaged in terrorism in the past, there cannot be any adverse implication that can be drawn from the individual’s membership in the organization. Where an individual becomes a member in an organization, then leaves and the organization subsequently becomes associated with terrorism, the nexus between the individual and terrorism is at best merely that of suspicion, less than the prescribed standard “reason to believe.”¹⁸

This broad approach to targeting members of problematic organizations is different than the mechanism used by Canadian criminal law, which does not make it an offence to be a

¹⁵ *Canada (Minister of Citizenship and Immigration) v Singh*, 1998 CanLII 8281 at para 52 (FC).

¹⁶ *Toronto Coalition to Stop the War v Canada (Public Safety and Emergency Preparedness)*, 2010 FC 957 at para 118.

¹⁷ For a larger discussion on this point, and an excellent summary of inadmissibility law, see Jamie Chai Yun Liew & Donald Galloway, *Immigration Law*, 2nd ed. (Toronto: Irwin Law, 2015) at 483-487.

¹⁸ *El Werfalli v Canada (Public Safety and Emergency Preparedness)*, 2013 FC 612 at para 76.

member of a group but criminalizes some behaviours that might assist terrorist organizations.¹⁹

4.1.1 Bases for security inadmissibility

The *IRPA* does not provide precise definitions for each category of inadmissibility. For example, the terms terrorism, security, and espionage are not defined in the statute. This means that the definitions have largely been elaborated by tribunals and courts.

Both immigration and criminal law sets out a high *mens rea*, or mental element, for terrorism, subversion, and espionage. For people to be inadmissible for security reasons, there must be reasonable grounds to believe that they must *intend* to commit the problematic activity. This high *mens rea* excludes from culpability people or groups that know that their acts might have problematic consequences or who act negligently.²⁰

4.1.1.1 Espionage against Canada or against Canadian interests

Espionage is the intentional “covert or surreptitious act of gathering information.”²¹ This immigration law definition is broader than the criminal law definition of espionage because it “does not require any element of hostile intent and can be occasioned even when carried out lawfully on behalf of a foreign government or agency.”²²

Not all surreptitious collections of information rise to the level of espionage. Because the statute only renders a person deportable for espionage conducted against Canada or Canadian interests, some collections of information are of no concern to immigration law. Justice Norris warns that adjudicators must not confuse espionage against Canada’s interests with espionage against “things Canada is interested in.”²³ For example, Canada may be interested in protecting the rights of all journalists around the

¹⁹ *Criminal Code*, RSC 1985, c C-46, ss 83.18 and ff.

²⁰ *Suresh v Canada (Minister of Citizenship and Immigration)*, 2002 SCC 1 at para 98 [*Suresh*].

²¹ *Afanasyev v Canada (Citizenship and Immigration)*, 2012 FC 1270 at para 19 [*Afanasyev*] (citing *Peer v Canada*, 2010 FC 752 at para 3, aff’d 2011 FCA 91).

²² *Afanasyev*, *supra* note 21.

²³ *Weldemariam v Canada (Public Safety and Emergency Preparedness)*, 2020 FC 631 at para 51.

world, but if a person or organization spies on a journalist in another country, this conduct will not engage the espionage provisions of the *IRPA*.²⁴

4.1.1.2 Subversion of any government by force or subversion of any democratic government by any means

Non-citizens are inadmissible for subverting any government or subverting democratic governments by force.²⁵ Subversion has been defined as “accomplishing change by illicit means or for an improper purpose related to an organization” and as “[a]ny act that is intended to contribute to the process of overthrowing a government.”²⁶ Force also has a broad definition. In *Oremade*, the Federal Court said that force was “coercion or compulsion by violent means, coercion or compulsion by threats to use violent means, [and] reasonably perceived potential for [the] use of coercion by violent means.”²⁷

Lawful acts taken against governments can, despite their legality, still render a person inadmissible for subversion. As Justice Gauthier of the Federal Court of Appeal explains:

As noted by the Division, the word “subversion” is not defined in the Act, and there is no universally adopted definition of the term. *The Black’s Law Dictionary’s* definition to which the Division refers at paragraph 27 (particularly, the words “[t]he act or process of overthrowing ... the government”) is very much in line with the ordinary meaning of the French text (“*actes visant au renversement d’un gouvernement*”). Although in certain contexts, the word “subversion” may well be understood to refer to illicit acts or acts done for an improper purpose, the words used in the French text do not convey any such connotation. I am satisfied that the shared meaning of the two texts does not ordinarily include any reference to the legality or legitimacy of such acts.²⁸ [emphasis added]

²⁴ *Ibid.*

²⁵ *IRPA*, *supra* note 8, ss 34(1)(b)-(b.1). For a discussion of this provision’s history see Jared Porter, “No Rebels Allowed: The Subversion Bar in Canada’s Immigration Legislation” (2018) 8:1 Sask L Rev 25.

²⁶ *Suleyman v Canada (Minister of Citizenship and Immigration)*, 2008 FC 780 at para 63; *Shandi (Re)*, [1991] FCJ no 1319, 17 Imm LR (2d) 54.

²⁷ *Oremade v Canada (Minister of Citizenship and Immigration)*, 2005 FC 1077 at para 27.

²⁸ *Najafi v Canada (Public Safety and Emergency Preparedness)*, 2014 FCA 262 at para 65.

Accordingly, people who join lawful campaigns against despotic governments can, for example, be found inadmissible.

4.1.1.3 Terrorism

In *Suresh*, the Supreme Court of Canada defined terrorism as:

[a]ny ... act intended to cause death or serious bodily injury to a civilian, or to any other person not taking an active part in the hostilities in a situation of armed conflict, when the purpose of such act, by its nature or context, is to intimidate a population, or to compel a government or an international organization to do or to abstain from doing any act.²⁹

This immigration law definition of terrorism is potentially different from the definition of terrorism set out in the *Canadian Criminal Code*. In the code, terrorism criminal offences are tethered to definitions of “terrorist groups” and “terrorist activity.” This means that terrorist offences in the *Code* and terrorism in immigration law may differ in at least two different ways. First, Canadian law does not criminalize *membership* in a terrorist organization. While non-citizens face immigration consequences for belonging to problematic groups, a person can only be convicted of a terrorist offence if they intentionally assist a terrorist group. Second, “terrorist activity” in the *Code* may have a broader reach than the immigration law definition. For example, “terrorist activity” includes “serious interference with or serious disruption of an essential service, facility or system” except for interference that is “a result of advocacy, protest, dissent or stoppage of work that is not intended to result in the conduct or harm referred [in other parts of the definition].”³⁰ This idea that “serious interference” could amount to terrorism is, at least in some readings, beyond that contemplated in the *Suresh* definitions.

Primarily for this reason, Justice Norris reminds adjudicators to attend to the differences between the two bodies of law: “[i]mporting criminal law concepts like ‘terrorist activity’ into the immigration context ... risks expanding the reach of ... the *IRPA* beyond what Parliament intended.”³¹ Importantly, however, the case law is

²⁹ *Suresh*, *supra* note 20 at para 98.

³⁰ *Criminal Code*, *supra* note 19, s 83.01(1)(b)(ii)(E).

³¹ *Rana v Canada (Public Safety and Emergency Preparedness)*, 2018 FC 1080 at para 49 [*Rana*].

ambiguous and it may be a mistake to overstate the differences between the definitions. Justice Gagné, writing after Justice Norris, explains that “I do not see a significant difference between these two definitions. In my view, the first definition is not broader or narrower than the other; the *Criminal Code* definition is simply more detailed while the *Suresh* definition is more general.”³²

4.1.1.4 Residual grounds

The *IRPA* makes people who are “a danger to the security of Canada” and people who engage “in acts of violence that would or might endanger the lives or safety of persons in Canada” inadmissible.³³ There is considerable uncertainty about the meaning of these sections. In 2020, the Federal Court of Canada observed that, at least until very recently, there was “very little jurisprudence considering the scope of” these residual grounds.³⁴ It appears that this is because the Canadian government made very few allegations under these sections.

Regarding the danger to the security of Canada provision, the Supreme Court of Canada has said in obiter that the scope of the provision reaches past security considerations. To decide whether it would be “detrimental to the national interest, the Minister must consider more than just national security and whether the applicant is a danger to the public or to the safety of any person.”³⁵

In 2022, the Supreme Court of Canada heard arguments regarding the specific residual section that allows for the deportation of persons who engage in “acts of violence that... might endanger the lives or safety of persons in Canada.”³⁶ At issue is whether these sections apply to all non-citizens who engage in dangerous conduct or whether the relevant “acts of violence” must have some nexus to Canadian security. In 2023, the

³² *Saleheen v Canada (Public Safety and Emergency Preparedness)*, 2019 FC 145 at para 38 [*Saleheen*].

³³ *IRPA*, *supra* note 8, ss 34(1)(d)-(c).

³⁴ *Dleiw v Canada (Citizenship and Immigration)*, 2020 FC 59 at para 5 [*Dleiw*].

³⁵ *Agraira v Canada (Public Safety and Emergency Preparedness)*, 2013 SCC 36 at para 82.

³⁶ *IRPA*, *supra* note 8, s 34(e).

Court found that the sections had to be interpreted narrowly, requiring a connection to a concern regarding Canadian security.³⁷

4.1.2 Adjudication

The *IRPA* sets out two procedural routes for Canadian immigration authorities to seek a finding of inadmissibility and to obtain a removal order. First, by bringing a case before the Immigration Division (the “ID” or the “Division”) of the IRB or, second, by asking the Federal Court of Canada to uphold a “security certificate” as reasonable.

This paper is only concerned with the first route and the reason for this interest is simple: the security certificate regime is a major, complex, but seldom used procedure. Very briefly, the security certificate process is designed to allow the government to obtain a removal order without disclosing all the evidence it relies upon to the potential deportee while installing several mechanisms to nonetheless protect the deportee’s right to a fair process.³⁸ Since 1991, Canadian authorities have issued certificates 27 times.³⁹ No certificates were issued during the period under study.

While it is beyond the scope of this article to explain why the government prefers one mechanism over another, I pause to note that the jurisprudence and scholarly literature regarding the security certificate regime is massive.⁴⁰ Many security certificate cases go on for years and require the expenditure of significant state resources. In contrast, most cases brought before the Division resolve quickly, often in a single sitting, and only sometimes are the subject of an appeal or judicial review. Security certificates are expensive, challenging, and the subject of significant public scrutiny, while proceedings before the Division are comparatively cheap, quick, and quiet.

There is a practical reason that explains why this is the case. As we shall see, the subject of almost every national security deportation case is a refugee claimant in Canada.

³⁷ *Mason v Canada (Citizenship and Immigration)*, 2023 SCC 21 [*Mason*].

³⁸ *IRPA*, *supra* note 8, Part 1, Division 9.

³⁹ Public Safety Canada, “Security Certificates” (last modified 1 December 2015), online: <publicsafety.gc.ca/cnt/ntnl-scrtr/cntr-trrrsm/srtr-crtfcts-en.aspx> [perma.cc/39G7-RCEK].

⁴⁰ See e.g. *Canada (Citizenship and Immigration) v Harkat*, 2014 SCC 37; *Charkaoui v Canada (Citizenship and Immigration)*, 2007 SCC 9.

When a person claims protection in Canada, they are required to self-disclose facts and evidence to substantiate their claims.⁴¹ Often in security cases, the evidence that the person discloses forms the basis of the government’s case against them. If a person, for example, claims that they are being persecuted by members of a group that they used to belong to, their own admission that they *were* a member of the group might establish an inadmissibility case. Security certificate proceedings make sense when Canada believes it is necessary to protect the means, methods, and evidence it obtains through covert channels, there is no such imperative when a person willingly turns over the evidence that will be used against them.

4.1.2.1 Adjudication before the Immigration Division

Every security inadmissibility case brought before the Division begins the same way. Section 44 of the *IRPA* allows immigration officers who are of the “opinion that a permanent resident or a foreign national who is in Canada is inadmissible” to write a “report setting out the relevant facts.”⁴² That report is referred to a delegate of the Minister of Public Safety and Emergency Preparedness who, if they determine that the report is well-founded, may refer the matter, along with all relevant evidence in the Minister’s possession, to the Division for adjudication.⁴³

The decision to refer may have collateral consequences for the prospective deportee. If the person has a pending claim for refugee protection, that application will be suspended pending the resolution of the inadmissibility case.⁴⁴ If immigration officials believe that the person poses a danger to the public or is unlikely to attend their own inadmissibility hearing, they may arrest the person and initiate immigration detention proceedings to hold the person in custody pending the resolution of their deportation matter.⁴⁵

⁴¹ *Refugee Protection Division Rules*, SOR/2012-256, ss 6–9, 11.

⁴² *IRPA*, *supra* note 8, s 44.

⁴³ *Ibid*; *Immigration Division Rules*, SOR/2002-229, s 3.

⁴⁴ *IRPA*, *supra* note 8, s 103.

⁴⁵ *Ibid*, ss 55–58.

Once the Division receives the report, it must hold a hearing where the immigration authorities begin by presenting their evidence and calling any witnesses. The prospective deportee may be represented by counsel, lead their own evidence, and make representations explaining why the Division ought not to accept the Minister's case.⁴⁶ People subject to inadmissibility proceedings do not have a right to silence and can be required to testify against their own interests.⁴⁷ If the Division determines that there are reasonable grounds to believe that the report's allegations are established with evidence, it must issue a deportation order.⁴⁸ If the Division finds that the report is not well-founded, it issues a "favourable decision."⁴⁹

4.1.2.2 Appeals before the Immigration Appeal Division

Deportees and immigration officials have asymmetric access to appeal rights. If the Division issues a favourable decision, the Minister of Public Safety may appeal to the Immigration Appeal Division (the "IAD"), another specialized section of the IRB.⁵⁰ The Minister's appeal rights are broad, and the IAD can allow the appeal, dismiss it, or substitute its own decision if it finds that the original decision was "wrong in law or fact or mixed law and fact" or made in breach of the duty to be fair.⁵¹ In contrast, deportees may not appeal to the IAD.⁵²

4.1.2.3 Judicial review before the Federal Court of Canada

If a deportee disagrees with the decision of the Division, or the deportee or the immigration officials disagree with how the IAD resolved its appeal, they may apply to judicially review the decision before the Federal Court of Canada.⁵³ A judicial review is not an appeal, but an administrative law process. Reviews are adjudicated with reference to the norms and principles of administrative law. Judges may not allow a judicial review

⁴⁶ *Ibid*, ss 162–167.

⁴⁷ *Bruzzese v Canada (Public Safety and Emergency Preparedness)*, 2016 FC 1119 at paras 17–18.

⁴⁸ *IRPA*, *supra* note 8, s 45.

⁴⁹ *Immigration Division Rules*, *supra* note 43, s 7.

⁵⁰ *IRPA*, *supra* note 8, s 63(5).

⁵¹ *Ibid*, s 67(1)(a)–(b).

⁵² *Ibid*, s 64(1).

⁵³ *Ibid*, s 72.

simply because they think the underlying decision was wrongly decided, instead it is their remit to assess whether the decision was reached through a fair process and whether the decision is reasonable. The Supreme Court of Canada explains that a reasonableness review requires judges to determine whether there “is a failure of rationality internal to the reasoning process” in the decision or if the decision is “untenable in light of the relevant factual and legal constraints that bear on it.”⁵⁴

Neither immigration officials nor deportees have an automatic right to judicial review. The *IRPA* requires parties to first apply to the Court for leave to bring an application. The purpose of the leave requirement is to screen out frivolous applications for judicial review. Applications are to be granted when there the case discloses “a fairly arguable case.”⁵⁵

Parties that disagree with a leave decision have, except in the rarest of circumstances, no right to further review or appeal.⁵⁶ Parties that disagree with a final judicial review decision of the Federal Court can only appeal that decision if the judge that decided the case certifies that the case raises “a question of general importance.”⁵⁷ If the judge does certify a question, parties are authorized to make an appeal to the Federal Court of Appeal.

4.1.3 Relief available for inadmissible non-citizens

Once the Division (or the IAD on a government-initiated appeal) finds a person inadmissible, a removal order is made and it becomes immediately enforceable.⁵⁸ If the deportee has permanent or temporary status, it is immediately revoked and, if a person is

⁵⁴ *Vavilov*, *supra* note 3 at para 101.

⁵⁵ *Bains v Canada (Minister of Employment and Immigration)*, [1990] 47 Admin LR 317, 109 NR 239 (FCA).

⁵⁶ *Canada (Citizenship and Immigration) v Tennant*, 2018 FCA 132, Stratas J (the Federal Court of Appeal can hear some cases in the absence of a certified question if the case discloses “a flaw that is ‘fundamental,’ strikes at ‘the very root’ of the judgment or ‘the very ability’ of the Court to hear the case, in some circumstances has ‘substantial particularity,’ and raises ‘serious concerns’ regarding the rule of law” at para 18).

⁵⁷ *IRPA*, *supra* note 8, s 74(d).

⁵⁸ *Ibid*, ss 45(d), 66, 69.

in Canada seeking refugee protection, their claim is automatically terminated.⁵⁹ The law says the person “must leave Canada immediately.”⁶⁰ If they do not, the government must enforce the order “as soon as possible.”⁶¹

People who wish to nonetheless stay in Canada have two procedural options available to them. First, the *IRPA* allows deportees to apply directly to the Minister of Public Safety and Emergency Preparedness for a declaration that the matters at issue do not, despite the issuance of a removal order, constitute inadmissibility. The Minister may only issue this declaration if they are satisfied that to do so would not be “contrary to the national interest.”⁶²

Second, Canadian and international law generally forbids the deportation of persons to places where they face risk.⁶³ If a refugee is found inadmissible for security reasons, their case is referred to the Minister who, after hearing from the deportee, determines whether “the person should not be allowed to remain in Canada on the basis of the nature and severity of acts committed or of danger to the security of Canada.”⁶⁴ If this risk-balancing process resolves in favour of the refugee, their removal order is stayed.

National security deportees who have not been recognized as refugees or protected persons are entitled to advance a claim that they will be killed, seriously harmed, or tortured if deported. This claim is advanced immediately prior to their deportation through a mechanism called the pre-removal risk assessment application (“PRRA”). If their claim is assessed positively, it is referred to the risk-balancing process described above.⁶⁵

⁵⁹ *Ibid*, ss 46(c), 103(1)(a).

⁶⁰ *Ibid*, s 48.

⁶¹ *Ibid*.

⁶² *Ibid*, s 42.1(1).

⁶³ *Ibid*, ss 96, 97, 115; *Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment*, UNGA, 39th Sess, UN Doc A/RES/39/51 (1984) GA Res 39/46.

⁶⁴ *IRPA*, *supra* note 8, s 115.

⁶⁵ *Ibid*, ss 112-115.

4.2 Empirical study of security deportation cases (2018–2020)

Two research questions inspired this study. First, I was interested in continuing an investigation begun by Dr. Angus Grant. Dr. Grant’s doctoral dissertation examined all national security inadmissibility cases decided by the division between 2002 and 2012. This major longitudinal study found: (1) the number of terrorism deportation cases steadily increased during the period under study; (2) that terrorism allegations were usually made against refugee claimants with citizenships in the Global South; and (3) that terrorism allegations almost never related to in-Canada conduct.⁶⁶ I was curious to determine whether, almost a decade later, these trends continued.

Second, immigration law practitioners were suggesting that the Federal Court of Canada was deciding BNP terrorism cases in an unpredictable and problematic manner.⁶⁷ Sometime in or around 2015, immigration officials started to regularly allege that members of the BNP, a massive Bangladeshi political party, were all inadmissible because the party had a history of using terror-based tactics.⁶⁸ Eventually, some of these cases reached the Federal Court of Canada on judicial review and, on its face, it did appear that something was amiss.

In 2017 and 2018, the Court decided eight BNP-related terrorism cases, most of which concerned deportation decisions made by the Division. The Court upheld decisions

⁶⁶ Angus Grant, *Confronting (In)Security: Forging Legitimate Approaches to Security and Exclusion in Migration Law* (PhD dissertation, Osgoode Hall Law School, 2016), online: <digitalcommons.osgoode.yorku.ca/phd/24/> [perma.cc/7RG7-YAKF].

⁶⁷ For examples of discussions on immigration law blogs regarding this issue see Steven Meurrens, “The Bangladesh Nationalist Party” (13 August 2019) online (blog): <meurrensonimmigration.com/the-bangladesh-nationalist-party/> [perma.cc/HKP5-FUJS], Raj Sharma, “The Bangladesh Nationalist Party (BNP) is (Not) a Terrorist Organization” (18 April 2017) online (blog): <sshlaw.ca/the-bangladesh-national-party-bnp-is-not-a-terrorist-organization/> [perma.cc/F72N-ML2F].

⁶⁸ The earliest reference to a BNP-related terrorism proceeding in public jurisprudential databases indicates that an s 44(1) report was signed on January 23, 2015, alleging the applicant was inadmissible for their membership in the BNP (see *Chowdhury v Canada (Citizenship and Immigration)*, 2017 FC 189 at para 4 [*Chowdhury*]).

concluding that the BNP engaged in terrorism in five cases.⁶⁹ There were three cases where the Court remitted the decision back as unreasonable.⁷⁰ In one of those three cases, Justice Mosley explained his rationale for finding the decision unreasonable this way:

I have considerable difficulty with the notion that a general strike called by a political party in an effort to force the party in power to take steps such as proroguing Parliament or convening by-elections, falls within the “essence of what the world understands by ‘terrorism’”. It is not an overstatement to suggest, as the Applicant has in these proceedings, that the Respondent’s interpretation of the statute could capture political activities which, if carried out in Canada, would be protected under s 2 of the *Canadian Charter of Rights and Freedoms*, absent an intention to use violence to achieve the political ends.⁷¹

For the most part, however, the judges of the Court largely insisted that there was nothing improper about the apparently conflicting outcomes. As discussed above, the institution of judicial review tolerates some discord because judges are not asked to determine whether each case was rightly decided, but whether each case’s decision comports with the requirements of reasonableness. If different cases featured different facts or different reasoning chains, it would make sense for different judges to assess different cases differently. And to be sure, in their decisions on BNP cases, this is how the judges explained any apparent discrepancies:

- **Justice Gagné:** “[E]ach case must be decided on its particular record and on the findings of fact made in the impugned decision.”⁷²
- **Justice MacDonald:** “Upon closer examination of these decisions however, it is clear that they are made in relation to particular findings and the particular evidentiary record before the Court. They are not broad proclamations on the status of BNP that bind future decisions.”⁷³

⁶⁹ *Gazi v Canada (Citizenship and Immigration)*, 2017 FC 94; *SA v Canada (Public Safety and Emergency Preparedness)*, 2017 FC 494; *Kamal v Canada (Immigration, Refugees and Citizenship)*, 2018 FC 480; *Alam v Canada (Citizenship and Immigration)*, 2018 FC 922 [*Alam*]; *Intisar v Canada (Citizenship and Immigration)*, 2018 FC 1128.

⁷⁰ *Chowdhury*, *supra* note 68; *AK v Canada (Citizenship and Immigration)*, 2018 FC 236 [*AK*]; *Rana*, *supra* note 31.

⁷¹ *AK*, *supra* note 70 at 41.

⁷² *Saleheen*, *supra* note 32 at para 26.

⁷³ *Rahman v Canada (Public Safety and Emergency Preparedness)*, 2019 FC 807 at para 33.

- **Justice Fothergill:** “[E]very case is determined based on the record before the tribunal. This accounts for the different outcomes.”⁷⁴
- **Justice Norris:** “While perhaps regrettable, it is inherent in the nature of judicial review under the reasonableness standard that perfect consistency across cases on questions of mixed fact and law will not always be achieved.”⁷⁵
- **Justice Grammond:** “These different outcomes may perhaps be explained by the different reasoning adopted by the ID in each case or by differences in the record before the ID.”⁷⁶

These points would be doctrinally sound, if it was indeed true that the underlying records and findings were dissimilar. I decided to investigate the extent to which they were.

My interest was also motivated by earlier scholarship about inconsistent, and potentially unfair adjudication before the Federal Court. In a pair of studies, Professor Sean Rehaag quantitatively analyzed tens of thousands of Federal Court of Canada leave decisions and found wide discrepancies between different judges’ leave grant rates: some judges frequently granted leave, some rarely did. This led Professor Rehaag to conclude that at least some outcomes before the Court turned on the “luck of the draw.”⁷⁷

Not everyone was persuaded by this type of analysis. At least one judge disputed the value of this sort of quantitative methodology. Justice Zinn, in a decision where he was asked to consider a similar study by Professor Rehaag regarding refugee decision-making, dismissed the value of quantitative studies. While he acknowledged that Professor Rehaag’s studies might “raise an eyebrow,” he said that the judiciary and the public should approach this sort of research cautiously, because:

the informed reasonable person, thinking the matter through, would demand a statistical analysis of this data by an expert based upon and having taken into consideration all of the various factors and circumstances that are unique to and impact on determinations of refugee claims before he or she

⁷⁴ *Alam*, *supra* note 69 at para 45.

⁷⁵ *Rana*, *supra* note 31 at para 7.

⁷⁶ *MN v Canada (Public Safety and Emergency Preparedness)*, 2019 FC 796 at para. 9 [*MN*].

⁷⁷ See Sean Rehaag, “Judicial Review of Refugee Determinations: The Luck of the Draw?” (2012) 38:1 *Queen’s LJ* 1; and Sean Rehaag, “Judicial Review of Refugee Determinations (II): Revisiting the Luck of the Draw” (2019) 45:1 *Queen’s LJ* 1.

would think it more likely than not that the decision-maker would not render a fair decision.⁷⁸ [emphasis added]

To anticipate and answer Justice Zinn's concern, I decided that my investigation ought to supplement the quantitative with the qualitative. For this reason, I elected to review every case, and every decision, in detail.

4.2.1 Methodology

This study has three stages, (1) data collection through access to information requests, (2) high-level analysis of the data using a customized coding program, and (3) a qualitative review of the cases. First, using access to information law, I obtained every national security inadmissibility decision issued by the Division between January 1, 2018, and July 15, 2020. I chose this time range because of its currency and because I wanted to work with a manageable dataset. I assumed that two years would give me good coverage while not inundating me with data. I extended the period slightly beyond the two-year mark to see whether there was any meaningful impact on tribunal operations because of COVID-19, given that I knew some tribunals paused their adjudication during period of the pandemic. Second, using a large database of Federal Court of Canada dockets and a customized computer program, I traced the outcomes of each case as it moved, via judicial review, from the Division to the Federal Court. Third, once I associated each case before the Division with its counterpart file (if any) at the Federal Court, I obtained and reviewed the full Federal Court file.

4.2.1.1 Access to information request

The *Access to Information Act* gives all Canadian citizens and permanent residents a general right, subject to specific statutory qualifications, to access and review records under the control of Canadian federal institutions.⁷⁹ One such qualification precludes bodies from releasing private information. Because the IRB is a public adjudicative body, its proceedings and decisions are presumptively public and accessible, with one major caveat:

⁷⁸ *Turoczi v. Canada (Citizenship and Immigration)*, 2012 FC 1423 at para 15.

⁷⁹ *Access to Information Act*, RSC 1985, c A-1, s 4.

when a case concerns a refugee or refugee claimant, the proceeding is held *in camera* to protect the person from malignant actors.⁸⁰ To ensure that protected information is not inappropriately released, analysts at the IRB review each request and, if necessary, apply redactions. For this reason, it can take some time for the IRB or any governmental body to process a request.

In August 2020, I asked the IRB to forward me all decisions associated with security and, for a separate study, war crimes and crimes against humanity inadmissibility cases, finalized by the Division between January 1, 2018, and July 15, 2020. In March 2021, I received 3,628 pages of disclosure from the IRB.⁸¹ Once I received the package, I read each decision and coded it for the name of the decision maker, the date of the decision, the prospective deportee's name, their counsel's name (if any), their country of citizenship, their reason for being in Canada, the allegation made against them, and the outcome.

As I was reviewing the dataset, I became concerned for two reasons that it was incomplete. First, the IRB publishes high-level statistical information about its cases on its website. This data suggested that 123 national security deportation cases were finalized in 2018 and that 97 cases were finalized in 2019. In contrast, the disclosure package provided included 58 decisions for 2018 and 44 for 2019. This suggested that there was a large body of jurisprudence the IRB had not disclosed. I made inquiries with the IRB and was advised that the website data included cases that were “withdrawn or [where] the individual failed to appear.”⁸² This addressed my concern as my study is focused on cases where a final decision was issued.

Second, as I reviewed each decision, I checked for cross-references in CanLii, Canada's main public repository for case law. During my examination, I found two Federal Court of Canada decisions that concerned national security cases decided by the Division

⁸⁰ *IRPA*, *supra* note 8, s 166.

⁸¹ Simon Wallace, “The New Canadian Law of Refugee Exclusion: An Empirical Analysis of International Criminal Law Deportation Orders, January 2018 to July 2020” (2022) 22:4 Intl Crim L Rev 721.

⁸² Jessica Arrechi, Immigration and Refugee Board analyst, “A-2020-00383” (28 May 2021) via email [communicated to author].

during the period under study that had not been disclosed.⁸³ I made further inquiries and learned that the IRB disclosed all cases marked in its internal database as “closed.” It appears that because the Federal Court returned the matters for re-adjudication during the period under study, the IRB’s database re-coded these decisions as “open.” I understood that these were the only two cases that were not disclosed because of this feature of the IRB’s database. With answers to these concerns, I resumed my analysis, confident that the dataset is, although not fully complete, all but complete.

4.2.1.2 Development of a Python program to analyze Federal Court of Canada dockets

Until recently, it was not possible to comprehensively associate first-instance IRB cases with the associated (if any) judicial review file at the Federal Court of Canada. The Court does not publish its decisions on leave applications. This means that researchers can only determine how a leave application was disposed of if they already know the court file number or party name, which can be used to query the Court’s docket. This information is often difficult to obtain because many of the IRB’s decisions concern refugee claimants. When these decisions are disclosed to researchers, names and identifying information are, as discussed above, redacted from the decision, making it impossible to look up how an individual case was treated.

As I analyzed the decisions, however, I observed that the Division’s internal case file number was properly never redacted from decisions. I also knew from experience that when litigants file applications for leave to judicially review a decision of the Division, they are required to provide that file number to the court.

During Professor Rehaag’s research into Federal Court of Canada decision-making, he developed his own massive database of information extracted from the Court’s official dockets. I approached Professor Rehaag about accessing this database of dockets that he reconstructed. Using a customized program coded in the Python programming language, the computer searched the database of every Federal Court of Canada docket to

⁸³ *Dleiw*, *supra* note 34; *SR v Canada (Public Safety and Emergency Preparedness)*, 2019 FC 1118.

find all cases with a similar internal Division file from one of the cases under study here.⁸⁴ This customized program was necessary because the metadata required was neither on CanLii or another public facing database, and because the program needed to search the textual data of the dockets. Because the Federal Court of Canada supervises many different tribunals with thousands of decisions, there was some overlap between file naming conventions. I manually reviewed each docket to isolate cases that concerned judicial reviews against Division matters.

In this study, using a minimal amount of case metadata and a small computer program, a cross-adjudicative database of cases was created to enable a previously impossible type of qualitative analysis. It is worth pausing to consider the significance of this methodological innovation. A feature of Canadian federalism and the separation between adjudicative institutions means that each department, each ministry, each court, and each tribunal maintains its own separate databases, often concerning the same people or the same matter. Because information is kept in silos, it is difficult for researchers to reconstruct a full picture.

4.2.1.3 Reviewing each Federal Court of Canada case file

Finally, I sought access to each Federal Court of Canada case file. In principle, all files in the Federal Court of Canada are accessible to the public.⁸⁵ Unfortunately, it took considerable time to access the files discovered by the Python program because COVID-19 restrictions prevented me from attending Court facilities until the Fall of 2021. In the meantime, Court staff were kind enough to scan and digitally forward several small files for my review.

⁸⁴ Python is an open-source programming language (see “Python,” (date last accessed 25 November 2023), online: <www.python.org> [perma.cc/C55D-X7XK]).

⁸⁵ Federal Court, “Policy on Public and Media Access” (last modified 29 March 2023), online: www.fct-cf.gc.ca/en/pages/media/policy-on-public-and-media-access [perma.cc/NM6U-X2TS].

4.3 High-level analysis of the dataset

Between January 2018 and July 2020, the Immigration Division adjudicated 128 national security inadmissibility allegations.⁸⁶ Some cases featured multiple allegations, explaining why the number of allegations (128) exceeds the number of cases (125). The Division found the person inadmissible and issued a removal order in seventy-two per cent of the matters. As Table 1 shows, most allegations turned on a person’s alleged membership in a group that engaged in espionage, terrorism, or subversion.

	Deportation order	Favourable decision	Total
Espionage	0	1	1
Subversion by force	3	0	3
Democratic gov’t subversion	0	0	0
Terrorism	0	0	0
Being a danger to the security of Canada	1	0	1
Engaging in acts of violence that would or might endanger people in Canada	0	1	1
Membership in a problematic organization	89	33	122

Table 1: National security inadmissibility cases, by allegation.

⁸⁶ This is the number of national security deportation cases in the dataset disclosed to me. It does not include the two additional cases I discovered in public jurisprudential databases.

4.3.1 The Immigration Division is Canada’s busiest national security tribunal

The Immigration Division is Canada’s busiest national security law tribunal and it has been getting busier. Dr. Grant’s earlier study of security-based deportation orders found that the number of cases increased from 2 in 2002 to 27 in 2012.⁸⁷ This study finds that the pace at which Canada has brought migration-related terrorism cases has accelerated further, finding that 59, 44, and 22 cases were finalized respectively in 2018, 2019, and the first half of 2020. This volume likely makes the Division the busiest national security tribunal in Canada. Indeed, the Division finalized more terrorism cases (59) in 2018 than Canadian criminal courts did between 2001 and 2018 (54).⁸⁸

4.3.2 Most national security cases were brought against asylum seekers from the Global South

Most national security cases involved refugees or refugee claimants. In 110 cases, the person’s identity was redacted from the decision. The fact of an *in camera* hearing is usually a good indicator that the person has a claim for protection or an appeal outstanding.⁸⁹ In a further eight cases, there were indications that the person was granted protection before Canada made deportation allegations. This means that ninety-four per cent of the cases concerned people who claimed that they would face persecution, torture, death, or cruel and unusual treatment if deported from Canada.

As Table 2 shows, only a small percentage of persons subject to national security inadmissibility allegations had citizenships in countries from the Global North.

Country of citizenship	Deportation ordered	Favourable decision	Total
Bangladesh	55	19	74

⁸⁷ Grant, *supra* note 66 at 136.

⁸⁸ Michael Nesbitt, “An Empirical Study of Terrorism Charges and Terrorism Trials in Canada between September 2001 and September 2018” (2019) 67:1/2 Crim LQ 95.

⁸⁹ Hearings are presumptively conducted in public unless a person has an outstanding claim for protection pending or there are reasons to believe that a person could be endangered, there is a risk that a public security matter will be discussed or holding the hearing in public would prejudice the fairness of the proceeding (see *IRPA*, *supra* note 8, s 166).

Ethiopia	8		8
Iraq	4	1	5
Côte d'Ivoire	5		5
Burundi	4		4
Nigeria		4	4
Zimbabwe	1	2	3
El Salvador	2	1	3
India	1	1	2
Chad	1		1
Congo/Germany	1		1
Egypt	1		1
Ghana	1		1
Jordan	1		1
Sudan	1		1
Turkey	1		1
UK	1		1
USA	1		1
Venezuela	1		1
Colombia		1	1
Eritrea		1	1
Mali		1	1
Russia/USA		1	1
Saint Lucia		1	1
Somalia		1	1
Yemen		1	1

Table 2: Citizenships of persons subject to national security inadmissibility allegations

The cases concerning nationals of the United Kingdom and the United States both involved members of far-right racist groups. The case of a Russian-American dual-national concerned a woman accused of engaging in espionage.

There were only two cases where the government alleged problematic in-Canada conduct. One concerned a person from St. Lucia whose criminal charges associated with a shooting in Canada were stayed (a deportation order was not issued).⁹⁰ The other concerned a citizen of Jordan who used social media to encourage people to engage in

⁹⁰ This case has been the subject of significant subsequent litigation, see *Mason*, *supra* note 37.

terrorism against Canada (a deportation order was issued). Both individuals were long-term residents of Canada.

Not including the BNP (discussed below), the groups Canada most often made membership allegations against were the *Fédération étudiante et scolaire de Côte d'Ivoire*, with five terrorism cases against nationals from the Ivory Coast, the Ethiopian Information Network Security Agency, with four espionage cases for Ethiopian nationals who allegedly worked for the Agency, and Ginbot 7, with three subversion by force cases against Ethiopian nationals.

A surprising class of allegations concerned the Iraqi KDP or Peshmerga. In four cases, people were ordered deported for belonging to groups that helped the allied American forces overturn Saddam Hussein's government. As one Member reasoned, helping American forces topple the Iraqi government made members of these groups inadmissible because they engaged in subversion by force against a government. He explained: "[t]he Kurdish Peshmerga assisted by the US military were finally able to defeat the Iraqi military and topple its oppressive leadership."⁹¹ It appears that no American soldiers were ordered deported from Canada for their participation in the same conflict.

4.4 Case study of deportation cases concerning members of the Bangladesh National Party

The most notable finding of this project is that most national security cases during the period under study turned on the same allegation: the deportee was a member of the BNP and the BNP engaged in terrorism. Of the 74 cases involving Bangladeshi nationals, Canada alleged that the person was a member of the BNP or a BNP-affiliated group 73 times.

⁹¹ B8-00663 (6 February 2019), Rempel Member.

In one other case concerning a Bangladeshi national, Canada alleged that the person was a member of the Awami League, the current governing party in Bangladesh, and that that party also engages or engaged in terrorism. Ultimately, the Division determined that the person in that case was not, in fact, a member of the party.

4.4.1 Evidentiary basis for the BNP terrorism allegation

Bangladesh is a populous country located in the North-East of India. It is a predominantly Muslim nation. It was part of colonial India until 1947 when the country was partitioned. Modern Bangladesh was, in British India, part of the province of Bengal. The western part of the province had a Hindu majority and became part of India. The eastern part of the province, with its Muslim majority, became part of Pakistan. The new Pakistani state was not territorially contiguous. Even though East Pakistan was more populous, political, and economic power was centralized in West Pakistan. During the 1960s, the Bengali Awami League led a political movement that advocated for increased autonomy and independence for East Pakistan. In 1971, the Pakistan Army launched a large operation to suppress Bengali independence. In March 1971, the leader of the Awami League declared Bangladesh's independence. A war for liberation lasted until December 1971, when the Pakistani forces in Bangladesh surrendered.

The BNP was formed in 1978 and quickly became a bitter rival of the Awami League. Between 1975 and 1990, Bangladesh was under military government rule. In 1991, the country held its first democratic elections and the BNP formed government. The Awami League boycotted the next election, complaining that it was not fairly administered. To protest the election, the League called for a hartal, or a general strike, that paralyzed the country. Under pressure, the BNP agreed to a constitutional amendment that required the transfer of political power to a caretaker government before and during an election. In 1996, an election supervised by a caretaker government was held

and the Awami League won. In 2001, after another election supervised by a caretaker government, the BNP was re-elected. In 2008, the Awami League returned to power.⁹²

In 2011, the Awami League abolished the caretaker government mechanism. Led by the BNP, opposition parties protested the move. They called for general strikes that peaked in 2013 and 2014.⁹³ Some of these protests turned violent. In one case, Justice Norris summarized the evidence about the general strikes this way:

The strikes and blockades had a significant impact on the economy. Transport links to Dhaka were blocked and almost all travel outside the major cities was prevented. *Hartals* and traffic blockades frequently turned violent, with clashes between supporters of the [Awami League] on the one hand and supporters of the BNP and other opposition parties on the other. Numerous instances of opposition party members and activists throwing petrol bombs at trucks, buses, and other vehicles that defied traffic blockades were documented. As well, attackers in several locations reportedly vandalized homes and shops owned by members of Bangladesh's Hindu community before and after the election. Opposition leaders denied their parties were involved in the violence, blaming government agents instead.⁹⁴

After the opposition was defeated, it appears that many BNP members fled Bangladesh. The United States Department of State, Human Rights Watch and Amnesty International report that the Awami League targets and persecutes BNP members.⁹⁵ Many individuals claimed protection in Canada.

⁹² See generally Central Intelligence Agency, “Bangladesh” (last updated 16 November 2023), online: <cia.gov/the-world-factbook/countries/125uopa125esh/> [perma.cc/5W98-JZSW]; Australian Government, *DFAT Country Information Report: Bangladesh* (Department of Foreign Affairs and Trade, 2019), online: <dfat.gov.au/sites/default/files/country-information-report-bangladesh.pdf>.

⁹³ European Asylum Support Office, *Country of Origin Information Report: Bangladesh Country Overview* (December 2017) at 29-30, online: <coi.euaa.europa.eu/administration/easo/Plib/Bangladesh_Country_Overview_December_2017.pdf> [perma.cc/SK8C-J7TS].

⁹⁴ *Rana*, *supra* note 31 at para 14.

⁹⁵ See United States of America, Bureau of Democracy, Human Rights, and Labor, *Country Reports on Human Rights Practices for 2020: Bangladesh* (Department of State, 2021), online: <state.gov/reports/2020-country-reports-on-human-rights-practices/Bangladesh> [perma.cc/A7Q8-LU4Z]; Human Rights Watch, “Bangladesh: Events of 2019” (2020), online: <hrw.org/world-report/2020/country-chapters/bangladesh> [perma.cc/LS8F-CWTJ]; Amnesty International, “Bangladesh. Human Rights in Asia-Pacific: Review of 2019” (30

Sometime in or around early 2015, the Canadian government started to make national security deportation allegations against BNP members.⁹⁶ Before the Immigration Division, the government said that the BNP’s involvement in the hartals amounted to terrorism. In the words of one adjudicator:

The BNP intentionally called for hartals to effect political change, however, with the passage of time, hartals have become increasingly violent leading to deaths, property damage and impacting the economy of Bangladesh; hartals resulted in 50 percent of all violence; at a minimum, the BNP was reckless and wilfully blind with regard to the consequences of continuously calling for hartals to achieve their political ends, aware of the ensuing violence that would inevitably put in danger the safety of the civilian population.⁹⁷

For this reason, Canadian officials say that the BNP engages in terrorism and that, therefore, all the party’s members are inadmissible to Canada.

4.4.2 Adjudication before the Immigration Division

Most of the national security allegations made during the period under study were made against members of the BNP. It appears that in every case, Canadian officials brought the case after the person self-disclosed information about their memberships and associations in a claim for refugee protection. The Division ordered the deportation of 55 of the 73 (seventy-five per cent) persons alleged to be members of the BNP.

	Montreal	Toronto	Vancouver	Total
Deportation order	46	7	2	55
Favourable decision	11	3	4	18

Table 3: BNP terrorism cases by site of adjudication

Of the 18 cases where a favourable decision was issued, the Division rejected the government’s request for a deportation order because it found that:

1. The Minister was unable to prove that the BNP engaged in terrorism (six cases);
2. The person concerned was not a member of the BNP (six cases); or,

January 2020), online: <amnesty.org/en/latest/news/2020/01/2019-in-review-bangladesh/> [perma.cc/P7R2-JXXW].

⁹⁶ *Chowdhury*, *supra* note 68.

⁹⁷ B6-00407 (23 July 2018), Tordorf Member.

3. The person left the BNP before it was possible to believe that the party would engage in acts of terrorism (six cases).

There were regional dimensions to the type of favourable decisions adjudicators made. Some adjudicators in Toronto and Vancouver reasoned that Canada did not lead sufficiently compelling evidence to prove that the BNP engaged in terrorism and, therefore, that Canada could not discharge its burden of proof. That argument was, during the period under study, never successful in Montreal.

Adjudicators also disagreed somewhat on the “no temporality” defence. A person can defend against a membership inadmissibility allegation by proving that they left the organization before it embraced terrorism. In each case where a person successfully advanced this defence, they left the party before 2012. However, there were two cases where adjudicators concluded that deportees were inadmissible even though they left the party in 2002 and 2011, respectively.

4.4.2.1 Inconsistent outcomes and reasons between adjudicators

Seventeen adjudicators made BNP-terrorism decisions. Four members, all of whom were in Montreal, decided three or more cases. No adjudicator changed their mind on a substantive issue during the period under study. The major debate amongst adjudicators was whether the BNP engaged in terrorism. In the end, seventy-one per cent of the adjudicators who considered at least one BNP case found that the party engaged in terrorism while the remaining twenty-nine per cent of their colleagues disagreed. Put differently, even though some members heard many more cases than their colleagues because no member changed their mind on the core issue of whether the BNP engages in terrorism, we can see that there was a substantive split amongst adjudicators.

	Location	Deportation order	Evidence does not prove that BNP engaged in terrorism	No temporality	No membership	Total
Morin	Mtl	19			2	21
Milo	Mtl	10		1	2	13
Thibault	Mtl	10		2		12

Tordorf	Mtl	7		2	2	11
Seyan	Tor	2				2
Ko	Van	1		1		2
Kohler	Tor		2			2
Adamidis	Tor	1				1
Beecham	Tor	1				1
Cook	Van	1				1
Del Duca	Tor	1				1
Heyes	Tor	1				1
Seifart	Tor	1				1
Currie	Tor		1			1
McPhalen	Van		1			1
Rempel	Van		1			1
Tessler	Van		1			1

Table 4: BNP cases, by adjudicator and outcome.

Adjudicators knew that they sometimes reached different conclusions than their colleagues regarding BNP cases. In at least 21 cases, lawyers presented written reasons from a case where the BNP inadmissibility allegation was rejected by another Division member. In each of those 21 cases, the presiding adjudicator declined to engage substantively with their colleagues' reasons. Instead, Members explained that it would be improper to guess what evidence their colleagues had before them when they reached the conclusion that the BNP did not engage in terrorism. For example, in every one of her ten decisions under study here, Member Thibault wrote that she was "unaware if similar evidence was before the [other decision maker]. Furthermore, that decision is not binding on the ID." It does not appear that any adjudicator tried to see whether the evidence before their colleagues was different or not.

4.4.2.2 Canada led functionally identical evidence across cases

I analyzed the full evidentiary records of each case (discussed further below) that was brought before the Federal Court of Canada. In every case reviewed decided in Montreal, 458 pages of the government's evidence were identical. The remaining twenty or thirty pages of evidence were specific to the person concerned and were used to substantiate the claim that the person was a member of the BNP. In both Toronto cases reviewed; Canada

submitted the same 598-page package to substantiate its allegation that the BNP engaged in terrorism.

The evidentiary packages did not include traditional scholarly or expert evidence. This is how Justice Grammond described the evidentiary package led in Montreal:

Apart from a number of newspaper articles, the evidence consisted mainly of two reports. One was written in 2005 by Bangladeshi academics under the auspices of the United Nations Development Program [UNDP] and analyses the causes and consequences of the frequent use of *hartals* in the political life of Bangladesh. The other was written in 2015 by Human Rights Watch and is mainly concerned with human rights violations committed by the country's security forces, although it also contains a smaller section devoted to "opposition violence."⁹⁸

The same analysis holds for the Toronto package. It included, in addition to several human rights reports, news articles from the *Wall Street Journal*, the *Guardian*, *Huffington Post*, the *Toronto Sun*, *The Telegraph*, the *BBC*, *Reuters*, the *New York Times*, *Forbes*, the *Financial Times*, the *Dhaka Tribune*, and *Vice News*.

4.4.4.3 Montreal-based adjudicators recycled their reasons across cases

Montreal members did not re-write their decisions after each case. Rather, they tended to add and subtract paragraphs and modified sentences over time. A close examination of two of Member Milo's decisions demonstrates this point. During the period under study, Member Milo issued ten BNP-membership deportation orders. The first and last decisions were decided, respectively, on December 24, 2018, and July 14, 2021. In both cases, the only substantive issue Member Milo addressed was the question of whether the BNP engaged in terrorism.

Despite the passage of time, and even though different lawyers represented the people subject to the deportation allegation, Member Milo's reasons in these two decisions are fundamentally similar. While Member Milo did write bespoke reasons describing the person's background, the critical sections on whether the BNP engaged in terrorism are functionally the same. The 2018 analysis of whether the BNP engaged in

⁹⁸ *MN*, *supra* note 76 at para 14.

terrorism is 5,045 words long while the 2021 analysis is 5,190 words long. Each paragraph is virtually identical. The differences between the two sets of reasons are found in two additional paragraphs added to the latter decision. One paragraph references a recently issued Federal Court of Canada decision and the second dismisses a discrete argument made by the deportee's lawyer. The Federal Court of Canada only reviewed Member Milo's last set of reasons (it appears that no other person ordered deported as a member of the BNP by Member Milo sought judicial review of the reasons). This decision was remitted back as unreasonable.⁹⁹

Sometimes members made minor adjustments to their stock reasons to respond to developments in the case law. For example, in 2018, Justice Norris released *Rana*, which reminded adjudicators that the *Suresh* and *Criminal Code of Canada* definitions of terrorism are different and should not be simply conflated.¹⁰⁰

Immediately before *Rana*'s release, Member Morin ordered a person deported for their membership in the BNP.¹⁰¹ In this case, he explicitly said that he used both the criminal law and immigration law definitions to determine whether the BNP engaged in terrorism. After *Rana*'s release, he modified his legal definitions paragraphs and explained that he would assess the case using only the *Suresh* legal definition. Despite saying that he was now applying a different and somewhat narrower legal test, the rest of his analysis remained word-for-word identical.¹⁰² Just over a year later, Member Morin's stated assessment of the law changed again slightly. Pointing to the mixture of directions from the Federal Court, he explained:

On many occasions, the Federal Court concluded that the Criminal Code could be used to consider terrorism as stated recently in *Saleheen v. Canada (Minister of Citizenship and Immigration)* 2019 FC 145. Relying on *Saleheen*, the Federal Court recently affirmed the principle that "... the specific intention to cause death or serious injury must exist for a finding of terrorism, whether the decision-maker applies the Criminal Code or the *Suresh* definition (*Miah v. Canada (Public Safety and Emergency*

⁹⁹ *Foisal v. Canada (Citizenship and Immigration)*, 2021 FC 404.

¹⁰⁰ *Rana*, *supra* note 31.

¹⁰¹ B7-00792 (11 October 2018), Morin Member.

¹⁰² B8-00122 (6 December 2018), Morin Member.

Preparedness), 2020 FC 38). Given so, the foregoing analysis will take both definitions into account.¹⁰³

Lawyers spend a considerable time parsing words and definitions to make their arguments. This vignette shows how court-level disagreements occasion ground level confusion. Despite trying to attend to changing case law, Member Morin kept a similar analysis but tried to explain how the same analysis could correctly orient to directions coming from the Court.

Member Morin also modified his decisions in response to another Federal Court of Canada decision concerning the quality of the evidence led in BNP cases. In *MN*, discussed above, Justice Grammond explained that he was concerned that the evidence before the decision maker (i.e. the standard Montreal evidentiary package) was not sufficient to “address the full range of circumstances” necessary to warrant a terrorism finding.¹⁰⁴

In three of the seven BNP deportation decisions Member Morin made after the Court published *MN*, Member Morin included a new paragraph that appears to respond to Justice Grammond’s discussion about the evidence. Even though it does not appear that the evidentiary package had substantively changed, Member Morin explained:

The documents filed come from various sources such as newspapers, but also from human rights monitoring agencies. Those documents describe the situation in Bangladesh and cover roughly 2 decades, from 1996 to 2015. The whole of this evidence provides a diversified yet coherent sum of information to allow the tribunal to draw conclusions on the issues that need to be decided in the present case.¹⁰⁵ [emphasis added]

Member Morin does not explain how the ‘whole of this evidence,’ which troubled Justice Grammond, does not trouble him. In fairness, this may be a consequence of how judges issue reasons. While Justice Grammond criticized the package, he did not include an index of the documents considered. It may be that because the index was not highlighted in the

¹⁰³ B9-00685 (11 February 2020), Morin Member.

¹⁰⁴ *MN*, *supra* note 76 at para 14.

¹⁰⁵ B9-00611 (10 March 2020); B9-00625 (14 July 2020); B9-00707 (14 July 2020), Morin Member.

reasons, it did not occur to Member Morin that he and Justice Grammond were looking at functionally the same material.

In *MV*, Justice Grammond reminded adjudicators to make sure that they were assessing whether the BNP *intended* to engage in terrorism, and not decide the cases on a lesser form of *mens rea* (such as knowledge or wilful blindness). In the twelve cases that preceded Justice Grammond's reminder, Member Morin explicitly said that he was deciding the case on a lesser form of *mens rea* (knowledge and wilful blindness):

In the present case, the tribunal concludes that violence was used to achieve political objectives, the link between the calls for hartals and the perpetration of terrorist acts is established. Given the predictable consequences of calling a hartal, it is difficult to find that political leaders did not know that deaths amongst the civilian population or serious bodily harm would result. Calling for a hartal is almost synonymous to endangering people's lives. Political leaders bare a certain responsibility. Hartals, in the context of Bangladesh, go beyond the mere expression of political activity as generally understood. Because of the history of decades of violence associated with such political demonstrations, only willful blindness would explain that political leaders were not aware of the human rights violations associated with such actions. In fact, there is a clear and documented pattern that the hartals lead to violence and economic chaos. It is equally clear that the acts of violence perpetrated during those hartals amount to terrorism. Deaths, random bombings, economic shutdowns, serious injuries, all a direct result of a political decision to call a hartal.¹⁰⁶ [emphasis added]

After the Court published Justice Grammond's decision, Member Morin substituted this long paragraph about wilful blindness with two shorter ones about intention:

In the present case, the tribunal concludes that violence was used to achieve political objectives, the link between the calls for hartals and the perpetration of terrorist acts is established. Hartals, in the context of Bangladesh, go beyond the mere expression of political activity or advocacy as generally understood. Because of the history of decades of violence associated with such political demonstrations, there is a direct link between hartals and human rights violations. Hartals, more often than not degenerate and as a result, people are killed or seriously injured during those protests. A call for

¹⁰⁶ B6-00497 (14 May 2018); B7-00173 (28 June 2018); B6-00449 (19 July 2018); B6-00565 (15 August 2018); B6-00700 (28 August 2018); B7-00792, *supra* note 102; B8-00122, *supra* note 103; B8-00123 (7 December 2018); B8-00371 (9 January 2019); B6-00682 (29 January 2019); B8-00359 (29 January 2019); B8-00374 (27 February 2019), Morin Member.

a hartal, in the context of Bangladesh, is intended to cause death or serious bodily harm. There is a clear and documented pattern that the hartals lead to violence and economic chaos. It is equally clear that the acts of violence perpetrated during those hartals amount to acts of terrorism. Deaths, random bombings, economic shutdowns, serious injuries, all a direct result of a political decision to call a hartal.

The tribunal infers that by calling a hartal, the political leaders intended to cause chaos, social disturbances and violence, they also expected that their sympathizers or members would enforce hartals with lethal force if necessary. Given the predictable consequences of calling a hartal, it is difficult to find that political leaders did not know that deaths amongst the civilian population or serious bodily harm would result. Calling for a hartal is synonymous to endangering people's lives. Political leaders bear a certain responsibility.¹⁰⁷ [citations omitted, emphasis added]

Despite reaching a different legal conclusion, these are the only paragraphs that changed in Member Morin's legal analysis. Put plainly, after Justice Grammond warned the Division that it was mistakenly applying the law, Member Morin kept the identical legal analysis but changed the operable words at the conclusion.

On its face, this appears to raise an integrity concern. That said, it is important to appreciate the difficult circumstances that Member Morin and other adjudicators of the Division were placed in. The Federal Court reviewed or decided leave applications regarding five of Member Morin's cases, upholding his decisions four times and remitting only one back. Faced with conflicting outcomes from the Court on similar reasons, a first-instance adjudicator would have tremendous trouble.

4.4.3 Adjudication of judicial reviews before the Federal Court of Canada

During the period under study, only 12 persons ordered deported for their membership in the BNP, out of a total of 55 persons, perfected applications for leave to judicially review the Division's inadmissibility finding. In every case, the deportee alleged that the Division

¹⁰⁷ B8-01284 (1 October 2019); B7-00169 (3 February 2020); B9-00685, *supra* note 104; B6-00624 (11 February 2020); B9-00611, *supra* note 106; B9-00707, *supra* note 106; B9-00625, *supra* note 106, Morin Member.

unreasonably concluded that the BNP engaged in terrorism. Leave was granted in nine cases (seventy-five per cent). The final application for judicial review was granted in four cases and dismissed in five cases.

Nine of the perfected applications (see Table 5) concerned the reasons of two adjudicators: Member Morin and Member Thibault. Because Canadian officials led the same evidence across cases, and because adjudicators recycled their reasons across cases, the judges of the Federal Court of Canada reviewed functionally identical records.

ID disposition date	Changes in reasons	Leave Judge	Leave	JR judge	JR granted	Court disposition date
<i>Decisions of Member Thibault</i>						
September 25, 2018	Baseline: 4963 words	McDonald	y	Grammond	y	June 10, 2019 ¹⁰⁸
October 16, 2018	90 new words referencing case law	Brown	y	Shore	n	July 8, 2018 ¹⁰⁹
October 19, 2018	70 new words referencing case law	Diner	y	Roy	y	July 19, 2019 ¹¹⁰
December 20, 2018	Minor redrafting led to a reduction of 177 words	Annis	y	Shore	n	August 29, 2019 ¹¹¹
<i>Decisions of Member Morin</i>						
October 11, 2018	Baseline: 4436 words	Shore	n	-	-	March 19, 2019 ¹¹²
January 9, 2019	445 new words referencing new case law;	Gleeson	y	Mosley	y	February 2, 2021 ¹¹³

¹⁰⁸ *MN*, supra note 76 (judicial review IMM-4992-18).

¹⁰⁹ *Khan v Canada (Citizenship and Immigration)*, 2019 FC 899 (judicial review IMM-5450-18) [*Khan*].

¹¹⁰ *Islam v Canada (Public Safety and Emergency Preparedness)*, 2019 FC 912 (judicial review IMM-5497-18) [*Islam*].

¹¹¹ *Ferdous v Canada (Public Safety and Emergency Preparedness)*, 2019 FC 1115 (judicial review IMM-259-19) [*Ferdous*].

¹¹² *Rahman v Canada (Public Safety and Emergency Preparedness)*, (judicial review IMM-5383-18).

¹¹³ *Islam v Canada (Public Safety and Emergency Preparedness)*, 2021 FC 108 (judicial review IMM-701-19) [*Islam*].

	155 new words summarizing the evidence					
January 29, 2019	310 additional words dismissing case-specific arguments	Roy	y	Walker	n	January 13, 2020 ¹¹⁴
March 10, 2020	370 new words summarizing the history of BNP; 74 words explaining why the evidence led is credible and reliable; 699 word long new summation of the evidence (replaces ~175 words)	Pamel	n	-	-	January 27, 2021 ¹¹⁵
July 14, 2020	Minor typographical changes	Pamel	n	-	-	January 27, 2021 ¹¹⁶

Table 5: Federal Court treatment of decisions issued by Members Thibault and Morin

In *MN* (IMM-4992-18; Federal Court decision: June 10, 2019), Justice Grammond identified a discrete legal error in Member Thibault’s stock reasons. In eight of Member Thibault’s decisions in the dataset, she explained that “by calling for hartals, the BNP leadership knew or, at best, was wilfully blind to the fact that it would result in deaths and serious injuries.” This, Justice Grammond explained, was an error:

the fact that lethal violence takes place during protests called by a political party may or may not lead to a finding that the political party has engaged in terrorism. Such a finding would need to be based on an analysis of a number of factors, including the circumstances in which violent acts resulting in

¹¹⁴ *Miah v Canada (Public Safety and Emergency Preparedness)*, 2020 FC 38 (judicial review IMM-1201-19) [*Miah*].

¹¹⁵ *Uddin Mishu v Canada (Public Safety and Emergency Preparedness)*, (judicial review IMM-3345-20).

¹¹⁶ *Hossain v Canada (Public Safety and Emergency Preparedness)*, (judicial review IMM-3348-20).

death or serious bodily harm were committed, the internal structure of the organization, the degree of control exercised by the organization's leadership over its members, and the organization's leadership's knowledge of the violent acts and public denunciation or approval of those acts. In this case, it appears that the ID focused exclusively on the last factor.¹¹⁷

In *Islam* (IMM-5497-18; Federal Court decision: July 19, 2019), concerning another case decided by Member Thibault, Justice Roy considered the same textual and zeroed in on the same problematic language discussed above. He allowed the judicial review and, explicitly agreeing with Justice Grammond, explained that “a lower [legal] standard was applied, one that is arguably close to recklessness or negligence as to what might ensue, and quite removed from the actual intent to cause death and serious injury.”¹¹⁸

In *Khan* (IMM-5450-18; Federal Court decision: July 8, 2018) and *Ferdous* (IMM-259-19; Federal Court decision: August 29, 2019),¹¹⁹ Justice Shore analyzed two sets of reasons by Member Thibault that were virtually identical—and contained the same language identified as problematic by Justices Grammond and Roy—to those considered in *MN* and *Islam*. In both cases, Justice Shore dismissed the applications for judicial review and upheld the reasons. This outcome is puzzling. In *Ferdous*, Justice Shore cited the decision made by Justice Grammond, which considered the same line of reasoning made by the same adjudicator, with approval. Indeed, Justice Shore cited the exact paragraph where Justice Grammond directly quoted the specific problematic language that was in all of Member Thibault's decisions. To state it plainly, it appears that Justice Shore—even though he said that he agreed with Justice Grammond—did not realize that he was considering a virtually identical text to the one Justice Grammond analyzed.

After Justice Grammond's decision in *MN*, Member Thibault decided one more BNP case. In that decision, she partially rewrote her reasons and edited out the reference to wilful blindness. In her new reasons, she concluded that the evidence showed that “the BNP leadership had the intention to cause deaths and serious injuries to reach a political

¹¹⁷ *MN*, *supra* note 76 at para 12.

¹¹⁸ *Islam*, *supra* note 110 at para 31.

¹¹⁹ *Khan*, *supra* note 109; *Ferdous*, *supra* note 111 at para 6.

objective.”¹²⁰ The deportee did not apply to the Federal Court of Canada for judicial review of this decision so we do now know how members of the Court might have considered these revised reasons.

Judges of the Federal Court also disagreed about the reasonableness of Member Morin’s stock reasons. Of the 11 people Member Morin ordered deported for their membership in the BNP, five applied to the Court for judicial review. Justices Shore and Pamel dismissed three applications for judicial review, meaning they did not think the case disclosed a “serious issue to be tried.”

Of the remaining two Member Morin cases that were reviewed, different judges of the Court did grant leave, meaning that they thought that there was a “serious issued to be tried.” The first full judicial review of Member Morin’s reasons was heard by Justice Walker, who upheld the decision under review as reasonable. In her decision, Justice Walker quoted Member Morin:

As illustrated by the documentary evidence, the 2012-2014 period in Bangladesh was one of the most violent in the history of Bangladesh. During that period, the BNP was the organization calling the hartals and promoting social disturbance to achieve political objectives. It is not plausible that there was not an underlying intention to achieve these goals through violence. The consequences of calling a hartal as well as the use of such a method to achieve political goals leaves little doubt of the intentions of political leaders calling for such actions. [emphasis added by Justice Walker].¹²¹

Justice Walker concluded that there was “no basis for the Court to intervene” because even if the reasons were “expressed in the negative, the ID imputed to the BNP and its political leaders the requisite specific intention to cause death and bodily harm.”¹²²

When Justice Mosley reviewed the same reasons, he focused on the same aspect of the stock reasons. In contrast to Justice Walker, Justice Mosley found that the reasoning was problematic. Referring to the exact same part of the text, he explained that the Member ignored “that the law requires that the perpetrator intentionally caused death

¹²⁰ B8-00368 (8 July 2019), Thibault Member.

¹²¹ *Miah*, *supra* note 114 at para 42, citing B8-00359 at para 76.

¹²² *Ibid* at paras 43-44.

and serious bodily harm, and substituting [sic] a different element (the requirement that there was knowledge, or even wilful blindness, that the calling for hartals would result in death and injuries).”¹²³ He remitted the decision back for redetermination. Together, this means one judge of the Court found that the same reasons did not disclose a serious issue to be tried, another found that the reasons were proper, while another concluded the decision was unreasonable and needed to be reconsidered.

Timing matters. Both deportation orders were issued in January 2019. Justice Walker’s decision was issued on January 13, 2020, and Justice Mosley’s decision was issued over a year later, on February 2, 2021. In the meantime, *MN* was released by Justice Grammond on June 10, 2019. By mid-2020, Member Morin’s new reasons explicitly referenced *MN*. Two of these slightly revised decisions were considered by the Court. Justice Pamel dismissed leave applications, therefore upholding the decision, in January 2021. If part of the function of judicial review is to help adjudicators, here the adjudicator received multiple Federal Court directions, but in a delayed and unhelpful order. Are courts and tribunals supposed to engage in a dialogue? If they are, this is not a conversation that made much sense.

One final point. For all that was the same across cases, there was one important difference: many applications for leave and for judicial review were brought by different counsel. While every counsel did say that the Division’s analysis of whether the BNP engaged in terrorism was unreasonable, not all counsel made their cases equally well or with equal force. Previous studies have argued that there may be a competence problem in some corners of the immigration and refugee bar.¹²⁴ Although an inquiry into counsel quality was not the subject of this project and not one that I am prepared to definitively comment on, it would be a mistake not to acknowledge that some advocates provided more assistance to the judges of the Court than others. For now, we can only hypothesize what

¹²³ *Islam*, *supra* note 113 at para 22.

¹²⁴ Jamie Chai Yun Liew et al , “Not Just the Luck of the Draw? Exploring Competency of Counsel and Other Qualitative Factors in Federal Court Refugee Leave Determinations (2005-2010)” (2021) 37:1 *Refugee* 61; Sean Rehaag, “The Role of Counsel in Canada’s Refugee Determination System: An Empirical Assessment” (2011) 49:1 *Osgoode Hall LJ* 71.

might have been if the quality of lawyer's work was consistent and consistently high quality.

4.5 Discussion and Recommendations

Most people subject to BNP terrorism allegations were ordered deported. This general truth obscures a more important reality: outcomes in individual cases turned on the identity of the adjudicators and judges who considered the case. Some members thought the BNP engaged in terrorism, some did not. Some judges saw errors in identical reasons whereas others did not. It is worth pausing to note that despite all this evidence of discord, and all these competing directions, it does not appear that a single first instance adjudicator changed their mind once about the BNP during the time under study. Once a member found that the BNP engaged in terrorism, they were never shaken from that conclusion.

This project therefore qualitatively confirms the conclusions in Professor Rehaag's quantitative studies that show that outcomes are often tied to the identity of the decision maker. By qualitatively examining the underlying texts—just as Justice Zinn suggested researchers ought to—I found that the texts were often the same, sometimes identical, but that different adjudicators and judges looked at those identical records differently.

On one level, this study suggests that individual adjudicators and judges ought to take more care with their decisions. There is nothing intrinsically wrong with first-instance adjudicators reusing reasons across cases when those cases are, in actual fact, largely the same. Indeed, we should hope that the same adjudicators would treat, and reason about, the same cases the same way. Yet there are moments, for example when a decision is overturned by a higher court, that require adjudicators to go back to the drawing board. The Canadian public should be concerned that some adjudicators, even when told by Federal Court of Canada judges that their decisions were problematic, only tinkered around the edges, changed a few keywords, and left the rest of their stock reasons untouched.

We must also appreciate the difficult position that the adjudicators of the Division found themselves in. It would be exceptionally challenging to re-engage in good faith when different judges of the Federal Court sent different messages about functionally the same texts. For deportees, their cases turn on the luck of the draw. In some senses, the same can be said for first-instance adjudicators who have their decisions reviewed by difference judges of the Federal Court of Canada, who might look at similar reasons differently. Whether the function of the judicial review is to strengthen reasoning, or to force re-evaluations of problematic reasons, here the project of effective judicial review was undermined by inconsistency.

For this reason, the primary conclusion of this paper concerns the Court. A core function of the institution of judicial review is to protect the consistency and coherence of the law. Here, the Court exacerbated inconsistency instead of improving it. Going forward, judges must take a greater interest in the decisions of their colleagues and closely examine the records before concluding that the case before them is idiosyncratic, unique, or unlike those considered by their colleagues. As this study shows, there is a chance that they all might be seeing the same reasons and the same evidence. To help judges and adjudicators better understand their own jurisprudence, I make several recommendations.

4.5.1 Recommendation 1: Tribunal and Court administrators should help decision-makers look beyond the four corners of each case

In *Vavilov*, the Supreme Court of Canada reminded courts that a “review of an administrative decision can be divorced neither from the institutional context in which the decision was made nor from the history of the proceedings.”¹²⁵ Instead of focusing exclusively on the textual material before them on judicial review, judges should also consider “publicly available policies or guidelines that informed the decision maker’s work, and past decisions of the relevant administrative body.”¹²⁶ So armed, courts will be

¹²⁵ *Vavilov*, *supra* note 3 at para 91.

¹²⁶ *Ibid* at para 94.

equipped to understand whether “a particular decision is consistent with the administrative body’s past decisions.”¹²⁷

In fairness to lower courts and judges, this is no easy task. As this study shows, many decisions are only publicly available in redacted form after an access to information request is actioned. It is not as if a lawyer can simply present the full jurisprudence to a judge to point out problems. While it is tempting to say that the answer is to just release more documents in publicly available datasets, there are practical reasons that make this difficult, including legitimate concerns about the free circulation of personal information, time-consuming and expensive process of redacting decisions, and the impossibility of translating every decision made by a federal Canadian tribunal or court for publication as required by the *Official Languages Act*.¹²⁸ In these circumstances, it is not reasonable to ask, as the Supreme Court of Canada does, for judges to familiarize themselves with the practises and decisions of each tribunal. There is simultaneously too much to read and not enough access to information.

For this reason, tribunals, and perhaps courts, should establish research departments to analyze jurisprudence at scale. If adjudicators and judges cannot in their day-to-day work read everything, tribunal administrators and executives should hire analysts to produce high-quality research products that present decision-making consistency and inconsistency alike. To a degree, this recommendation calls for a departure from a classic model of advocacy that makes lawyers responsible for framing the issues and presenting the relevant evidence. But, given how sprawling administrative law has become and how the public’s access to some decisions is limited, lawyers will have less of a line of sight on consistency issues that tribunals or courts.

In practice, this means that tribunals should be ready to present basic statistical information about each case to advocates, parties, and courts. How many other similar cases has the tribunal decided? How were those cases decided? Are outcomes consistent or inconsistent? Equally, tribunals should ensure that parties and courts can at least

¹²⁷ *Ibid* at para 131.

¹²⁸ *Official Languages Act*, RSC 1985, c 31 (4th Supp), ss 14-20.

partially access the records and decisions in other cases to see where they are the same and where they are different. Recently, the Immigration and Refugee Board established “quality centres” to “mine... decisions, as well as Federal Court cases in order to identify specific legal questions that repeatedly occur.”¹²⁹ This sort of initiative is positive and ought to be encouraged.

At the same time, judges should think about how their writing can be more useful to adjudicators. If specific language is problematic, it should be excerpted. If a collection of evidence is wanting, its index should be included in the judgement. This is to say that it would be useful for adjudicators and lawyers to have more direct access to the problematic documents that concern judges on judicial review and judges could immediately assist by explicitly highlighting and quoting problems in their reasons. If everyone can see and not just hypothesize about the problems in a first-instance text, we can better understand them.

4.5.2 Recommendation 2: Supplement traditional research methodologies with computational methodologies

The nature of this study points to one methodological avenue available to researchers, tribunals, and courts: computational methodologies for the study of law. This study was enabled by a simple but powerful network analysis that drew connections between case files across tribunals. Given Canada’s federal structure and the separation between many tribunals, networking analysis techniques may be of particular use to shadow databases that link information across institutions. At the same time, recent innovations in artificial intelligence have allowed researchers to use computers to examine the contents of legal decisions and extract, at an extremely high degree of accuracy, actionable legal insights.¹³⁰

¹²⁹ Immigration and Refugee Board, *Quality Assurance Framework for Decision-Making* (Ottawa, 2021) at 19, online (pdf): <irb.gc.ca/en/transparency/qa-aq/Documents/Quality_Assurance_Framework_for_Decision_Making_2021.pdf> [perma.cc/J7JP-B8PP].

¹³⁰ Indeed, at the time of writing, some of the most exciting technology is so new that the legal publication cycle has not caught up to the modern capabilities of artificial intelligence. See e.g., Sean Rehaag, “Luck of the Draw III: Using AI to Examine Decision-Making in Federal Court

Computational methods have arrived, and not a moment too soon. Consider this fact: in 2022, 90,000 people claimed refugee protection in Canada. Assuming that some people claimed together as family units, this means that the Refugee Protection Division will be asked to decide approximately 65,000 separate cases based on this intake alone. If each decision issued by the tribunal is 2,000 words long, the resulting refugee law jurisprudence will be 130 million words long. No research team, no matter how well resourced, can read and synthesize this much text. Computers, however, can assist by categorizing, clustering, and extracting information from decisions.

The scholarship on how computational methods can assist researchers remains in its infancy, but these approaches are already opening “important new research opportunities for law scholars by expanding the analytic methods that can be applied to legal texts.”¹³¹ Courts, law schools, firms, and adjudicative bodies should all dedicate resources to train a generation of lawyers and tribunal administrators with the hard skills required to code, productively interact with modern artificial intelligence, and develop programs that can ingest and synthesize legal meaning at scale.

4.5.3 Recommendation 3: Consolidate appeal routes

As currently written, the *IRPA* does not facilitate effective judicial review. As discussed earlier, when the government disagrees with a decision, it may appeal to the Immigration Appeal Division. When a deportee disagrees with a decision, they can only apply for leave for judicial review in the Federal Court of Canada. Practically, this means that the Federal Court only sees cases where a deportation order was issued while the IAD only sees cases where there was no deportation order issued. This no doubt partly explains why the judges of the Federal Court struggled with the state of the jurisprudence: they only saw one side of it.

Stays of Removal” (2023) Osgoode Legal Studies, Refugee Law Lab Working Paper No 4322881, online: <papers.ssrn.com/sol3/papers.cfm?abstract_id=4322881> [perma.cc/8VKL-C9FY]. For a recent, but already out of date, summary of legal computational methods research see Jens Frankenreiter & Michael A Livermore, “Computational Methods in Legal Analysis” (2020) 16 Annual Rev L & Soc Science 39.

¹³¹ *Frankenreiter & Livermore, ibid* at 40.

From an integrity of the justice system perspective, this is an undesirable arrangement and Parliament should consolidate appeal routes. A good template for legislators to use can be found in the refugee adjudication context where almost all decisions are appealed to the Refugee Appeal Division. A similar review process should be set up in the deportation context to help foster the development of a coherent jurisprudence. Alternatively, Parliament should re-route governmental appeals into the judicial review process so, at minimum, there is only one review body considering these cases.

5. Modeling Canada’s Refugee Law Jurisprudence

An underexplored aspect of the experience of forced migration is textual. My point is not that scholars have not explored some of the discursive elements of forced migration (they have), rather, my point is that we have not thought much about the fact that forced migration generates an enormous quantity of text. When people cross borders in a search of safe haven, and when they engage the mechanisms of refugee protection, the fact of their movement can generate millions of words of text. Now multiply the text produced by an individual claimant by the total number of refugee claimants. An unfathomable quantity of text.

Let us think this through in a bit more detail. Imagine a single refugee who makes a claim in Canada. First, they tell an officer—a border official—that they want to make a claim. This will be documented.¹ They will then be examined.² Part of this process will include an interview.³ Notes will be taken, perhaps a transcript will be made. Next, the claimant will fill out forms setting out the reasons for their claim.⁴ Perhaps they will hire a

¹ *Immigration and Refugee Protection Act*, SC 2001, c 27, ss 99(3) and 99(3.1). [IRPA]

² *Ibid*, s 15.

³ Minister of Citizenship and Immigration, “Intake of claims for refugee protection at ports of entry” (24 December 2023) online: < <https://www.canada.ca/en/immigration-refugees-citizenship/corporate/publications-manuals/operational-bulletins-manuals/refugee-protection/canada/intake-claims-refugee-protection-ports-entry.html#>>.

⁴ Refugee Protection Division Rules, SOR/2012-256, s 6. [RPD Rules]

lawyer to help them draft it. Perhaps they will develop multiple drafts. Eventually, their claim document (their “Basis of Claim” form) will be forwarded to the Immigration and Refugee Board, the independent tribunal charged with adjudicating most claims for protection made in Canada. A new file will be opened, new entries in a database will be made, and new administrative data will attach to the case. Before their hearing, the claimant will forward any documentary evidence they intend to rely on to support their claim to the Board.⁵ This will be added to their file. If they intend to rely on the evidence of a witness, they must write to the Board with information about the witness.⁶ But there is more evidence to be included. The Board, on its own initiative, will file a country conditions package prepared by its own research department (the “National Documentation Package”) for each country the person claims protection against. This package can be thousands of pages long.⁷ Most claimants will go to a hearing⁸ and at that hearing they will be questioned, given the opportunity to lead evidence, and to make legal arguments explaining why their claim should be accepted.⁹ At the end of the hearing or sometime after, the person deciding the case will issue a decision and reasons for that decision.¹⁰ Then, if either the government or the claimant thinks the wrong decision was made, they can appeal—which is another-text heavy process.¹¹

The point: a single case, even a simple one, is bound to produce hundreds of thousands of words of evidence and material. More complex cases, and cases that are subject to further review, will generate even more text.

Now let us consider the question of scale. In 2023, over 143,770 people claimed refugee protection in Canada.¹² How many words will their claims generate? The total

⁵ *Ibid*, s 34.

⁶ *Ibid*, s. 44.

⁷ Immigration and Refugee Board, “National Documentation Packages” (23 Jan 2024) online: < <https://www.irb-cisr.gc.ca/en/country-information/ndp/Pages/index.aspx>>.

⁸ *IRPA*, *supra* note 1, s 170.

⁹ RPD Rules, *supra* note 4, s 10.

¹⁰ *Ibid*.

¹¹ *IRPA*, *supra* note 1, s 110.

¹² Immigration, Refugees, and Citizenship Canada, “Asylum claims by year – 2023” (22 April 2024) online: < <https://www.canada.ca/en/immigration-refugees-citizenship/services/refugees/asylum-claims/asylum-claims-2023.html>>.

number of words is too large to even be guessed at. But to make an obvious point, this mass of text ought to interest scholars interested in Canadian refugee law. In aggregate, this corpus of text will be the best available record of what Canadian refugee law was, and how it was actually experienced, by this great quantity of people who engaged it. What types of claims do people make? How do they frame them? How are they considered? How do different people consider them? What are common leitmotifs? How do some leitmotifs track with some concepts? How do some track with others? Why? The best answers are all there, waiting to be read. Even the most normative and conceptual among us must feel the gravitational pull of this corpus. To change the law, to even say it must, could, or should change, is to make a claim about its actual on-the-ground shape. What is the law? The best answers to this question are always grounded in at least some empiricism. And in the refugee law field there is enough data an empiricist could work for lifetimes.

And it would take lifetimes, wouldn't it? No one works with this corpus of text because of the challenges associated with reading. You and I are physiologically limited. We read one word after the other, over and over. This takes time and this takes energy. Maybe on a good day we can read and analyze several tens of thousands of words, but that is a very good day—and one without any scheduled teaching or meetings. Hundreds of billions of words? Yes, valuable, but beyond our grasp and beyond our comprehension.

Now let me turn to my first argument. Maybe we cannot read that much, but could we teach a computer to? And if we did, could we develop a line of sight on some of the large empirical insights present in masses of text that are presently inaccessible? The example set by a small but growing cohort of digital humanists suggests that we can. By using modern open-source technologies, these scholars have developed a new sort of method that legal scholars can learn from. Each of us is used to close reading, but what if we used a computer to read from a distance? A geographer can both see a hill and read a map; why couldn't we both read a case and see the massive patterns of text produced by movements of people?

Exploring and expanding on this idea is the first objective of this paper. In the first part of this text, I introduce the reader to a set of approaches, perspectives, and methods that can be used to computationally elicit qualitative insights from large corpuses of text.

Artificial intelligence, machine learning: for many of us these words suggest research agendas well beyond our grasp. I show that they are not.

But it is better to show and not tell. This will bring me to the second part of the paper where I report on a study of my own. The Federal Court of Canada is this country's effective court of last resort for refugee claimants whose applications for protection have been refused. This means that its decisions form the most authoritative refugee law jurisprudence in the country. What, I wonder, was this jurisprudence about between 2013 and 2023? To answer that question, I show how researchers can take thousands of decisions representing millions of words of text and extract variables of interest and use advanced machine learning methods to classify texts.

The approaches—which are, in fact, quite simple, at least compared to what some of the major artificial intelligence research centres are now doing—reveal interesting and surprising insights about Canada's refugee jurisprudence. First, it is increasingly negative: the proportion of cases where the refugee claimant succeeds is steadily decreasing. Second, it is topically simple: when all is said and done, there are three major categories of refugee law case. Third, it is changing: ten years ago, cases that turned on the adjudicator's assessment of state's ability to provide protection (a key part of the refugee test) accounted for a good bulk of the Court's jurisprudence; in 2023, state protection barely registers as a matter of concern. In its place, a jurisprudence that asks questions about the viability of internal flight alternatives has developed.

Lawyers and professors tend to assess cases in doctrinal terms. This study challenges, or at least questions, the utility of these sorts of analyses. Large computational analyses encourage particularity to fade while inviting regular and repeating patterns to emerge to the fore. In the final analysis, it is difficult to spot locations where judicial or lawyerly creativity could play a role in decisively shifting or changing the law. Instead, the view we are left with suggests that large material forces are what compels jurisprudential change. To know what refugee law jurisprudence will be about, it is best to know what countries people are fleeing: the legal concepts follow. The role played by legal actors looks much more modest. Here judges and lawyers look, at least to me, like interpreters whose function it is to translate the diversity of human experience into a few legally intelligible

categories. Bourdieu, who believed that law was the result of a competition between different actors for the right to define law, would I think look at the maps and diagrams produced here and nod. In the end do I present a picture of the field of law embedded within the larger field of power? Yes, I think so. This, and a few other small practical matters, I pick up in the final discussion section.

5.1 Using computational methods to read from a distance

This much is obvious: digitization has changed our relationship with text. No doubt the transition from print to digital will occupy media scholars for lifetimes, but one specific change interests me here. Digitization has allowed us to *read* differently. Think about how you are reading this text. You are (I hope) reading it closely. You read each sentence, you digest the meaning that I have communicated, you read the next sentence, and so on. Reading is progressive: we move from word to word, sentence to sentence, allowing each element of the text to build on the previous one. Words, characters, paragraphs, grammar: together they allow for the communication that is more than the sum of its parts.

But let us stop here for a moment. What do digital techniques now make possible that was previously impossible? Well, now we can do the summing. If yesterday we read closely to understand larger meanings that transcended the discrete parts of a text, today we can read mathematically to count, divide, and multiply text so that we can spot patterns in it. Yes, our access to the meanings in text that are more than sums may (although, I emphasize, may) be lost, new perspectives on text come into view. Which authors use more violent and aggressive words and why? Well, let's begin to answer the question by counting.

The reason why summing is possible today, but that it was not practically possible before, is obvious. How good are you at counting to one million? How good are you at counting to one billion? On your own, this task is impossible. As the authors of a recent text on computational methods point out, researchers rarely engaged with textual data at

scale because it was “impossibly time consuming” to read comprehensive collections of text, it was intellectually difficult to “organize the texts into relevant categories,” and, because there was too much information to keep track of, no one could “measure the presence of concepts of interest” at scale. Moreover, until very recently it was “difficult to acquire documents because there was no clear way to collect and transcribe all the things people had written and said.”¹³ Now that most documents are born-digital, and an increasingly large proportion of the world’s information is stored in computer accessible formats, it is possible to use computational tools to interact with text in new ways. And who got us to first think deeply about those new ways? The socialists.

5.1.1 The study of the “great unread”

The first major scholarly movement to think through how computational tools could be used to aid investigations into large masses of text developed in the early 2000s in, somewhat surprisingly, the pages of the *New Left Review*. Franco Moretti, a Marxist student of literature, was dissatisfied with the state of the sociological study of literary texts. In the 1990s, literary scholars were divided by the canon wars, with scholars debating the value of the canon, the significance of its exclusions and its elisions, and the promise of new sorts of canons.

For Moretti, these debates about the canon seemed to miss the point. As a sociologist, he wanted to explore the social conditions that produced canonical texts in the first place, that led to a change in the canon, or invited the rise of new sets of canons. Beneath those few canonical texts, Moretti saw that there were thousands, hundreds of thousands, of other creative endeavours—some failed, some ignored—that were the precondition to the canon. Genius does not emerge fully formed from Zeus’ head: it comes to us as an iterated creature of history and social relations. That we do not examine these social conditions and these histories is partly a feature of the method of close reading: “the trouble with close reading... is that it necessarily depends on an extremely small canon.

¹³ Justin Grimmer, Margaret Robbins & Brandon Stewart, *Text as Data: A New Framework for Machine Learning and the Social Sciences* (New Jersey: Princeton University Press, 2022) at 3. [Text as Data]

This may have become an unconscious and invisible premiss by now, but it is an iron one nonetheless: you invest so much in individual texts *only* if you think that very few of them really matter.”¹⁴

Instead, Moretti wondered, could texts be read at a distance? He explained:

Distant reading: where distance, let me repeat it, is a condition of knowledge: it allows you to focus on units that are much smaller or much larger than the text: devices, themes, tropes—or genres and systems. And if, between the very small and the very large, the text itself disappears, well, it is one of those cases when one can justifiably say, Less is more.¹⁵

Drawing on a training informed by the natural science scholarship, Moretti set out to model literature. Just as scientists use diagrams, charts, and concepts to explain the natural world, Moretti wondered if he could productively graph and map literary change. Using network theory, for example, he presented new readings of Hamlet by showing how the different characters belonged to different sub-groups.¹⁶ Using trees inspired by evolutionary biologists, he showed how a series of writers slowly developed the concept of the clue in mystery fiction, until the idea was realized in its modern canonical form by Sir Arthur Conan Doyle.¹⁷

Following Moretti’s path, a new intellectual movement developed: digital humanism. As computers became increasingly powerful and accessible to university-based researchers, as more texts and artistic productions were digitized, technically skilled students of the humanities explored how they could model large artistic change using methods developed by computer scientists.

In some cases, these studies focused closely on changing word patterns. A study of World Bank annual reports co-authored by Moretti, for example, traced how the bank’s neoliberal turn was reflected in the sentence structure of its reports. Where the bank’s language was once focused on the development of heavy industry in the Global South (and

¹⁴ Franco Moretti, “Conjectures on World Literature” in Franco Moretti, *Distant Reading* (London: Verso, 2013) at 48.

¹⁵ *Ibid.*

¹⁶ Franco Moretti, “Network Theory, Plot Analysis” in *Ibid* at 211.

¹⁷ Franco Moretti, “The Slaughterhouse of Literature” in *Ibid* at 63.

therefore used many verbs and nouns), by the 1980s the texts focused more on partnerships and flexibility. Sentences became longer and the use of the word “and” increased markedly.¹⁸

Researchers of forced migration took note. In 2005, two linguists examined the discursive patterns associated with the terms “refugee” and “asylum seeker” in over 75,000 words of news articles and 265,000 words of text from the UNHCR that were published in 2003. Among their many interesting findings, they discovered that the expression “asylum seeker” had more of a negative inflection than the word “refugee.”¹⁹ Indeed, computationally enabled media discourse analyses proliferated in the mid-2000s, especially when scholars started to obtain access to raw social media data.²⁰ Historians have also used computational methods to conduct large longitudinal studies about how narratives about refugees have developed and changed.²¹

Recently, in large part because of the work of the Nordic Asylum Law and Data Lab,²² the Refugee Law Lab,²³ and Kaldor Centre Data Lab,²⁴ legal scholars have begun to explore how computational methods can be used to monitor and better understand decision making in refugee cases. Here, the work of Professor Sean Rehaag is emblematic.

¹⁸ Franco Moretti & Dominique Pestre, “Bankspeak: the Language of World Bank Reports” (2015) 92 *New Left Review* 75.

¹⁹ Paul Baker & Tony McEnery, “A corpus-based approach to discourses of refugees and asylum seekers in UN and newspaper texts” (2005) 4:2 *Journal of Language and Politics* 197.

²⁰ See as examples: Nazan Öztürk & Serkan Ayvaz, “Sentiment analysis on Twitter: A text mining approach to the Syrian refugee crisis,” (2018) 35:1 *Telematics and Informatics* 136; Adina Nerghes & Ju-Sung Lee, “Narratives of the Refugee Crisis: A Comparative Study of Mainstream-Media and Twitter” (2019) 7:2 *Media and Communication* 275; Claire Kelling & Burt Monroe, “Analysing Community Reaction to Refugees through Text Analysis of Social Media Data” (2023) 49:2 *Journal of Ethnic and Migration Studies* 492; and

²¹ Brandon Green & Antoine Pécoud, “Talking about Migration in Times of Crisis: A Textual Analysis of Narratives by IOM and UNHCR on Migrants and Refugees” (2023) *American Behavioral Scientist*.

²² University of Copenhagen, “Nordic Asylum Law & Data Lab” (last accessed 25 June 2024), online: <<https://asylumdata.ku.dk>>

²³ York University, “Refugee Law Lab” (last accessed 25 June 2024), online: <<https://refugeelab.ca>>.

²⁴ UNSW Sydney, “Kaldor Centre Data Lab” (last accessed 25 June 2024), online: <<https://www.unsw.edu.au/kaldor-centre/our-resources/kaldor-centre-data-lab>>.

In three papers, colloquially known as *Luck of the Draw*,²⁵ *Luck of the Draw II*,²⁶ and *Luck of the Draw III*,²⁷ Professor Rehaag uses increasingly complex (and increasingly computer assisted) methods to extract data about decisions from court dockets containing semi-structured textual data. He used this data to calculate the rates at which refugee applications for judicial review and emergency motions to stay removal were successful and, in the process, determined that there were wide discrepancies between different judges' grant rates.

Research out of the Nordic Asylum Law and Data Lab points to at least some of the directions researchers are now oriented towards. In a preview of a larger research project, doctoral researcher Asta Sofie Stage Jarlner uses an advanced computational technique to identify different topics within a collection of Danish asylum law studies. This method analyzes all the decisions to identify those themes and words latent within texts, allowing a researcher to make predictions about what various documents might be about.²⁸ This cutting-edge approach is helpful because it can help make sense of large amounts of text by showing researchers how like documents might be discovered and classified together.

To return to the original point, these examples show how computational methods, by treating text as data as opposed to something that can only be traditionally read, can be used to model large corpuses of text to develop insights that would be unavailable to us if we only looked at a few texts closely.

²⁵ Sean Rehaag, "Judicial Review of Refugee Determinations: The Luck of the Draw?" (2012) 38:1 Queen's LJ 1.

²⁶ Sean Rehaag, "Judicial Review of Refugee Determinations (II): Revisiting the Luck of the Draw" (2019) 45:1 Queen's LJ 1.

²⁷ Sean Rehaag, "Luck of the Draw III: Using AI to Extract Data About Decision-Making in Federal Court Stays of Removal" (2024) 49:2 Queen's LJ 73.

²⁸ Asta Sofie Stage Jarlner, "Using topic modelling to explore case summaries" (28 March 2023), Nordic Asylum Law & Data Lab (blog), online: <<https://asylumdata.ku.dk>>.

5.1.2 Using computational approaches to uncover patterns in text data

But how do computational methods actually work? To varying degrees, each researcher canvassed here followed the same multi-step process. First, they developed a dataset, or a corpus, of texts that they were interested in. Then, second, they used structured code-based methods to (i) extract information from the text, (ii) applied *supervised* machine learning methods, and/or (iii) applied *unsupervised* machine learning methods to classify texts within their dataset. In this section, I briefly discuss what each of these steps involves.

Let us begin at the beginning: dataset acquisition. Collecting texts to form a dataset for study can be a challenging—sometimes the most challenging—part of a research project.²⁹ This is primarily for two, largely opposed, reasons. On the one hand, even though data is everywhere, sometimes states, governments, corporations, and entities guard it closely. To take a Canadian example, many courts are reluctant to grant researchers bulk access to their decisions, citing privacy concerns.³⁰ This means that even if there are lots of digital material out there, legal or institutional barriers may prevent researchers from accessing and working with it. On the other hand, when large datasets of decisions are available, it can be difficult to narrow in on only those documents that are of interest or that are responsive to a research question.

Sometimes researchers can develop their own datasets. For example, in some cases large collections of text documents might be hosted on someone else’s website or server. In these cases, it is possible to “scrape” the decisions, or to write a computer program that systematically visits and downloads the contents of the webpage onto a researcher’s

²⁹ Niels Brügger, “Digital humanities and web archives: Possible new paths for combining datasets” (2021) 2 *International Journal of Digital Humanities* 145.

³⁰ Jo Sherman, *Guidelines For Canadian Courts: Management of Requests for Bulk Access to Court Information by Commercial Entities*” (Ottawa: Canada Judicial Council, 2021) and Jon Khan & Sean Rehaag, “Promoting Privacy, Fairness and the Open Court Principle in Immigration and Refugee Proceedings” (2023) 54:2 *Ottawa L Rev* 357.

computer, allowing them then to work with the data themselves.³¹ However, in some cases, web scraping can present ethical and legal problems.³²

Most textual data is referred to as unstructured data. Unstructured data is information that does not follow a predictable or consistent format.³³ At the beginning of any computational project, the researcher is required to think about how the data can be reformatted or reconceived to make it amenable for computational analysis, which, as we shall soon see, is always mathematical analysis. In each case, this means working with the data to uncover and work with patterns within it. Sometimes we are fortunate because some textual data comes to us with enough pre-existing structure that meaningful patterns can be easily uncovered. Consider a legal decision: often decisions from the same jurisdiction will feature the same sort of formatting. A file number might always be found on the third line, the name of the judge might always follow the same construction (e.g. every decision might feature a line “Adjudicator: x”). In these cases, we can write code to isolate for these constructions and extract the data that is of interest.

For more complex structures, computer scientists have developed tools to allow researchers to uncover “regular expressions” in data. For example, we might know that a date is always constructed the same way in a particular jurisdiction (perhaps, for example, a Court always writes the date as “dd/mm/yy”) and using a specifically crafted “expression” finds all instances where there are two digits followed by a backslash, followed by two more digits, followed by a backslash, followed by two final digits.³⁴ Sometimes, this sort of data extraction effort can be the major component of a research

³¹ M. Olmedilla, M.R. Martínez-Torres & S.L. Toral, “Harvesting Big Data in social science: A methodological approach for collecting online user-generated content” (2016) 46 *Computer Standards & Interfaces* 79.

³² Alex Luscombe, Kevin Dick & Kevin Walby. “Algorithmic thinking in the public interest: navigating technical, legal, and ethical hurdles to web scraping in the social sciences” (2022) 56:3 *Quality & Quantity* 1023.

³³ Adanma Cecilia Eberendu. “Unstructured Data: an overview of the data of Big Data” (2016) 38:1 *International Journal of Computer Trends and Technology* 46.

³⁴ Carl Chapman & Kathryn T. Stolee. “Exploring regular expression usage and context in Python” (2016) *Proceedings of the 25th International Symposium on Software Testing and Analysis*.

project. In *Judicial Review of Refugee Determinations (II): Revisiting the Luck of the Draw*,³⁵ for example, Professor Rehaag crafted a computer program that searched out many data structures he identified within the dataset. His final report was then based on an analysis of the computationally extracted data.

In many other instances, however, textual data defies easy categorization and extraction. In these cases, researchers will use either supervised and/or unsupervised machine learning methods to classify documents. Let me begin by discussing what machine learning is. Machine learning “is a branch of computer science that broadly aims to enable computers to ‘learn’ without being directly programmed.”³⁶ Inspired by statistical approaches, the goal of any machine learning algorithm is to find a way to fit a mathematical formula to a dataset.

The first step in any machine learning project is to convert the data into a form that the computer can work with. Or, to put it simply, to convert the data into numbers that the computer can build a mathematical formula around. Converting text to numbers can be counterintuitive and there are many different approaches researchers can use to convert the data. At its most basic, one method to convert a text into numbers can be to simply count the number of times a particular word appears in a document. This approach, a count vectorizer, will ultimately create a large table of data. Each row will represent a separate document, and each column will show the number of times a specific word appears in that document.

As can be readily appreciated, this approach is imperfect. The meaning associated with the order in which the words are used in is entirely dropped, long documents will dominate the dataset,³⁷ and this method has no way of appreciating what different words might mean because it is only counting instances. But, recall, this is just the simplest of

³⁵ Sean Rehaag, “Judicial Review of Refugee Determinations (II): Revisiting the Luck of the Draw” (2019) 45:1 Queens LJ 1.

³⁶ Qifang Bi et al, “What is Machine Learning? A Primer for the Epidemiologist” (2019) 188:12 American Journal of Epidemiology 2222 at 2222.

³⁷ Consider: a document with ten thousand words might use the word “refugee” 30 times, while a document with one thousand words might use the word “refugee” 29 times. We recognize that the shorter document is more likely to be concerned with the concept of the “refugee,” but this will not be apparent from a straight count vectorization.

many different methods available to a researcher to transform a document into numbers. In the next part, we will consider a slightly more complex process.

Once a dataset has been transformed into numbers, it is ready for either a supervised or an unsupervised machine learning project. In a supervised machine learning project, a researcher will first review a selection of their data (in our case, a selection of their texts) and identify how they think those decisions ought to be organized. For example, if a researcher is interested in dividing legal decisions into one group where an appeal was granted and another where the appeal was denied, they will label each of the selected decisions as either “granted” or “denied.” The computer will then be asked to find a way of mathematically developing a formula from the original data (in this case the text) that gets to the desired outcome (“granted” or “denied”).

This approach can be helpful when the data, which many text documents are, is complex and inconsistently organized. In *Luck of the Draw III: Using AI to Extract Data About Decision-Making in Federal Court Stays of Removal*,³⁸ Professor Rehaag used a supervised learning approach to classify the outcome of emergency motions to stop deportations from Canada. After reviewing and labelling a selection of Court dockets, he then directed the computer to develop a model that could distinguish between dockets where a motion was granted or a motion was refused. Once all the dockets were classified, he was then able to assess overall grant rates for motions to stay a removal.

Unsupervised approaches are useful when the researcher does not have a clear idea of how the dataset, or in this case documents, might be organized or classified. In these cases, a researcher will specify a few parameters (a few settings) and direct an algorithm to search out patterns within the data. There are many different use cases for unsupervised algorithms. Let us consider two here. Dimensionality reduction, first, is the process of taking a complex dataset with many variables (one that, for example, contains many thousands of different words) and combining them into a new dataset of fewer variables. This can be useful for making it easier to work with a dataset or to visualize patterns within

³⁸ Sean Rehaag, “Luck of the Draw III: Using AI to Extract Data About Decision-Making in Federal Court Stays of Removal” (2024) 49:2 Queen’s LJ 73.

the data (in text data, for example, there might be tens of thousands of variables, but dimensionality reduction enables us to represent at least some of the variability in two dimensions, in turn allowing us to plot the relative locations of data points).

Clustering, second, is when researchers direct algorithms to search for groupings of communities of data within a dataset. For example, we might hypothesize that decisions about criminal law will be about similar themes and use similar words, at least compared to decisions about civil law. A clustering algorithm might be able to detect that some data points are similar to each other, while others belong to other communities. Soon we will see how these methods can be combined to uncover interesting patterns within textual data. But before that, we must briefly consider some of the dangers of uncritically relying on algorithmic returns.

5.1.3 The value of validation

Computers are excellent pattern detection machines, but there is no guarantee that the patterns mathematics can uncover will be of interest or use to a social scientist. When using supervised methods, for example, there is a risk that the algorithmic model might “overfit” to the data. This is when the statistical model learns so much about the sample of data it is shown that it develops a formula that accounts for all the noise within the sample, and disregards the signal that interests the researcher. Put differently, it learns the sample so well, it becomes good at recognizing the sample, and bad at recognizing why the researcher selected and labeled the sample. A model that is “overfit” will perform badly when it is asked to label or make predictions about unseen data.³⁹

Other times, the categorizations that a social scientist is interested in might be so closely related that the computer has difficulty understanding the differences between data points within different categories, and wrongly assigns them to the same categories. The problem of “overlapping classes,”⁴⁰ my experience has shown, is a real challenge

³⁹ X. Ying, “An Overview of Overfitting and Its Solutions” (2019) 1168 J Phys: Conf Ser 022022.

⁴⁰ For a discussion on overlapping classes see Jerzy Stefanowski, “Overlapping, Rare Examples and Class Decomposition in Learning Classifiers from Imbalanced Data” in Zachary A. Pardos &

when working with textual data, particularly legal textual data. I hypothesize that this is because legal texts by nature share a similar vocabulary and part of the function of legal reasoning is to make dissimilar subjects or actions comprehensible in similar terms.

For these reasons and others, it is imperative that the researcher validate the returns of any algorithmic conclusion. Computer scientists have developed a range of metrics that researchers can use to assess the quality of any algorithmic conclusion, but there is no substitute for actually engaging with the data and the computational conclusions. At the end of the day, it is the researcher's responsibility to ensure that their work product is useful, reproducible, and understandable in terms that make sense, are related to the research question, and are interpretable.

5.2 The study: measuring change in the Federal Court of Canada's refugee law jurisprudence

To demonstrate how the computational methods canvassed here can be used together to develop new classes of insights on large bodies of refugee-related text, I report on a study of a large collection of Canadian refugee law documents. The Federal Court of Canada is functionally the court of last resort for refugee claimants in Canada whose claims have been declined. In Canada, most claims are first decided by an independent tribunal called the Immigration and Refugee Board (the "IRB"). Since 2012, the Board has had a two-level adjudicative process in which most cases are heard in the first by the Refugee Protection Division (the "RPD") and, if either the claimant or the government disagree with the RPD's decision, the Refugee Appeal Division (the "RAD") hears appeals. This description is something of an over-simplification because not every claimant and not every case is eligible to be appealed to the RAD *but*, and this is the key point, every claimant whose case is ultimately refused by the IRB, may apply for the Federal Court of Canada to review their case.

Tiffany Barnes, eds, *Emerging Paradigms in Machine Learning* (Berlin: Springer Berlin Heidelberg, 2013) 277.

Strictly speaking, the Court does not hear refugee appeals. Judicial reviews are decided on Canadian administrative law principles, usually turning on an assessment of the “reasonableness” of the underlying decision. In a reasonableness review, the question is not whether the decision is correct, but whether “the decision as a whole is transparent, intelligible and justified.”⁴¹

The Court does not review every case that it is asked to review. Before a full judicial review can be heard, a single judge of the Court must grant a party’s application for “leave” or permission to advance the review.⁴² This criterion for the leave requirement is not explicitly spelled out in statute and, given that leave decisions are issued without reasons and cannot be appealed, there is not an elaborated body of case law that delineates the shape and nature of the test to be applied. To the extent that there is a leading case, it is *Bains* in which Justice Mahoney said that to be granted leave a party must only establish “a fairly arguable case.”⁴³

Most applications for leave fail. In 2022, only 24.8% of applications for leave were granted by the Court.⁴⁴ If leave is granted, an application can resolve in several ways. The party bringing it may abandon it or withdraw it, the matter may be settled and resolved by way of a consent order (where the parties jointly ask the Court to issue an order directing the Board to reconsider the application), or it may be sent to a full hearing at the end of which a judge will issue a decision remitting the matter back for reconsideration or dismissing the appeal. If a decision reaches a hearing and a full merits review, the Court will issue justificatory reasons explaining its decision.⁴⁵ For most claimants, the Court is their last chance to have decisions about their refugee cases reversed. A handful of cases

⁴¹ *Canada (Minister of Citizenship and Immigration) v. Vavilov*, 2019 SCC 65 (CanLII), [2019] 4 SCR 653 at para 15.

⁴² *IRPA*, *supra* note 1, s 72.

⁴³ *Bains v Canada (Minister of Employment and Immigration)*, [1990] 47 Admin LR 317, 109 NR 239 (FCA) at para 1.

⁴⁴ Refugee Law Lab, “Federal Court” (visited 28 May 2024), online: <rlp.ca/federal-court>.

⁴⁵ For a discussion on the legal status of reasons, and the Federal Court of Canada’s changing approach to the issuance of reasons, see Sean Rehaag and Pierre-André Thériault, “Judgments v Reasons in Federal Court Refugee Claim Judicial Reviews: A Bad Precedent” (2022) 45:1 Dal LJ 185.

are, each year, appealed to the Federal Court of Appeal, but the number of these cases is diminishingly small.

This means the decisions and reasons released by the Court form the most substantive and significant body of decisions regarding refugee matters. While the Board does release a selection of its decisions, its decisions (are at least in theory) less significant than the Court's because they do not bind and influence decision makers in the same way. Understanding the Court's body of refugee related work, the object of this study, is to understand Canada's most significant refugee law jurisprudence. But understanding this jurisprudence is no easy task because the Court decides cases, and not just refugee cases, at scale. It is therefore not surprising that, to date, explicit studies on the Court's jurisprudence—as opposed to focused studies on discrete issues the Court considers—are absent.

Here, my object is to develop an approach on the Court's jurisprudence so it can be conceived of and written about despite its enormous size. While there is novelty to the approach, I argue that the new vantage points and perspective will strike subject matter experts as intuitive. To work closely with a developing body of text, as many refugee advocates and government lawyers do, is to appreciate some of its larger patterns and textures, and to understand that it is a changing set of documents, responding to social developments. Put differently, the computational presentation I make here will resonate with the impressionistic accounts of the people who work in the field, while showing them patterns that they may have forgotten or not been able to notice.

5.2.1 Collecting the dataset

For the purposes of this study, my goal is to critically describe ten years of the Court's refugee decisions. The first task therefore was to collect the decisions. Despite recent scholarly efforts to make computational analysis of Canadian jurisprudence easier, this remains a difficult task because the Court does not provide direct bulk access to its internal databases or its classifications of cases. Put differently, researchers cannot simply ask the Court for its decisions, and they certainly are not able to just ask the court to provide it

only its refugee law decisions. This meant that I needed to collect a full set (or as close as possible to a full set) of the Court’s decisions and isolate for the refugee law decisions.

To build my collection of decisions, I relied heavily on two datasets released by Professor Sean Rehaag. The first, the “Federal Court Bulk Decisions Dataset”⁴⁶ contains over 30,000 plain text decisions, with associated metadata, of Court decisions released online between 2001 and 2024. The Court translates each decision from English to French, and *vice versa*, to comply with the *Official Languages Act*⁴⁷ and the dataset contains (almost) all of the translations so, in truth, the dataset covers approximately 15,000 original decisions. This dataset is not, in a strict sense, comprehensive. The Court does not publish every decision it renders, often only retaining file copies for consent orders or various interim decisions. Moreover, between 2015 and 2018, the Court did not publish every final decision it rendered, choosing instead to only publish a selection of decisions thought to have precedential value.⁴⁸ These gaps are of no moment to this study. My interest is less in describing the totality of the Court’s work, but in better understanding the shape of the decisions that were one day released and made readily available to the public and do, today, form its jurisprudence.

To download and work with this dataset, I developed code and scripts using a programming language called Python. Python is an open-source project maintained by a non-profit foundation.⁴⁹ While it is difficult to strictly quantify how popular a programming language is at a given moment, I decided to work in Python because—at least at the time of writing—it was thought to be one of the best supported programming languages for data science work. This meant that for most of the more challenging and advanced methods I wished to employ, I could easily plug in implementations already developed by expert computer scientists instead of trying to reinvent the wheel myself.

⁴⁶ Refugee Law Lab, “Federal Court Bulk Decisions Dataset” online: <<https://refugeelab.ca/bulk-data/fc/>>.

⁴⁷ *Official Languages Act*, RSC 1985, c 31 (4th Supp).

⁴⁸ Sean Rehaag & Pierre-André Thériault, “Judgments v Reasons in Federal Court Refugee Claim Judicial Reviews: A Bad Precedent” (2022) 45:1 Dal LJ 185.

⁴⁹ “Python Software Foundation US”, online: Python <www.python.org/>.

After the script I developed downloaded the decisions, I broke out the English and French language versions of each decision. I found that the dataset contained 15,657 English language decisions and 13,878 French language decisions. I attribute the difference to the fact that it takes time to translate documents. It appeared that a backlog in translation work meant that many English language decisions had not yet been translated into French: I found that 1,948 English language decisions had yet to be translated, while only 184 French language decisions were waiting for translation. From my perspective, this meant that working with the English decisions meant that I would have a near complete collection of the Court's jurisprudence. I finally filtered for those decisions released between January 1, 2013, and December 31, 2013.

5.2.2 Isolating for refugee related decisions

Now that I had this dataset of decisions, the next task was to isolate for cases that concerned refugee claims that originated from the IRB. After studying a selection of the decisions, I noticed that there were patterns in how the documents were constructed. I wrote a program to take advantage of these patterns and to pull out variables of interest from each decision. This program reviewed the text of each case to identify (i) the judicial official who presided, (ii) the file number of the case, and (iii) the header, body, and footer text of each decision. By chance, this method helped me narrow in on the decisions that interested me. On occasion, the Court published an interlocutory decision (a non-final decision) in a case. My program attempted to split each part of the text by finding the first numbered paragraph of each decision. As it turned out, many interlocutory decisions did not have numbered paragraphs while all final decisions did, meaning that the program inadvertently identified most of the few non-final decisions in the dataset, allowing me to easily weed them out.

The fact that I was able to extract the file numbers from each decision was useful. All file numbers associated with immigration-related judicial convention. Each immigration judicial review file number begins with the same prefix ("IMM"), meaning that I could filter for file numbers with the prefix to find only immigration related judicial reviews. After doing this, I was left with 8,884 decisions. While this dataset was just over

50% of the size of the original English language dataset (n=15,657), it remained massive. I wrote a program that counted the number of words in the dataset and found that this immigration judicial review dataset exceeded 25 million words in length.

But this dataset remained overinclusive. Refugee cases are a subset of all immigration law judicial reviews the Court considers and it became necessary to develop a new method to identify only those decisions that concerned refugee matters from the IRB. This proved to be complex. Within the decisions, I could not find consistent patterns of words that clearly identified and separated refugee related matters from, for example, visa related judicial reviews.

This is when I relied on the second dataset published by Professor Rehaag. As part of his study on emergency motions to stop deportations, Professor Rehaag published a dataset of the Court's dockets, or its internal administrative database of cases. This administrative dataset included classifications for each case pulled from the Court's own registry system. It appears that a court administrator identified each type of different judicial review and entered this classification into the docket system.⁵⁰ I computationally compared the file numbers I extracted from each case with the file numbers in this second dataset and, where there was a match, extracted the classification from the Court's system. This method—marrying up two records—classified 96% of the 8,884 decisions contained in the decisions dataset. I wrote some short code that took these classifications and identified whether the case concerned an IRB refugee matter (True) or not (False).

To classify the remaining 4%, I used a supervised machine learning method. Recall how supervised machine learning methods work: the researcher shows the computer data and a pre-assessed outcome. The computer then develops a formula that associates the data with the outcome and, when this formula is shown new data, it makes a prediction about what the outcome for that new case or observation might be. In this case, 96% of the cases were already labeled and my goal was to identify the remaining 4%. To do this, I took a random selection of 3000 cases where I knew the classification. I then identified the two

⁵⁰ Refugee Law Lab, "Luck of the Draw III Dataset" (2023), online: Hugging Face <<https://huggingface.co/datasets/refugee-law-lab/luck-of-the-draw-iii>>.

columns of data I wanted to work with: the decision text for each decision and the final True/False refugee classification. Then I split the data into two sets: one training dataset with 2,400 cases and one test dataset with 600 cases that I would hold back, not train the algorithm with, but see whether it could correctly identify the class of the decision text.

Before training the algorithm, I needed to transform all the text into a format the computer could work with (that is, I had to change each text into numbers). I converted the data into a numerical representation with an algorithm called term frequency-inverse document frequency (“tf-idf”).⁵¹ In essence, this algorithm counts the instances of each word or expression⁵² in a document and divides it by the overall importance of the word in all the documents. Words and expressions that are relatively unique or important to a specific document get a higher score, while words that are common across all documents get lower scores.

Now that the data was prepared, I used a machine learning model called Linear Support Vector Classification⁵³ to search the data for patterns that could be used to predict the classifications of each case. This model appealed because it is widely used in the literature for text-based classification tasks.⁵⁴ After it was trained, I tested it on the 600 cases in the test dataset. The algorithm turned out to be highly effective, correctly classifying 99% of all cases. I reviewed a proportion of the instances where the algorithm was found to have erred. Interestingly, I found that in these few instances I often agreed with the algorithm over the human classification. It appeared that in some cases the court staff person entering the classification erred and that the algorithm was able to notice at least some of these instances. Put differently, the algorithm may outperform humans with respect to this classification task.

⁵¹ Juan Ramos, “Using tf-idf to determine word relevance in document queries” (2003) 242:1 Proceedings of the First Instructional Conference on Machine Learning 29.

⁵² In this case, I set the algorithm to search for expressions up to five words in length.

⁵³ For all machine learning algorithms, I used the implementations set out in F Pedregosa et al, “Scikit-learn: Machine Learning in Python” (2011) 12 Journal of Machine Learning Research 2825.

⁵⁴ See, as an example, this early paper: Jason Rennie & Ryan Rifkin, “Improving multiclass text classification with the support vector machine” Massachusetts Institute of Technology, Department of Brain & Cognitive Sciences and Artificial Intelligence Laboratory, 2001.

Given that I knew that this text-based classification method was effective, I decided to repeat the process to obtain another variable of interest from the decisions: the final judgement. This time I did not have pre-existing labels for each case that identified how it resolved. To develop the labels, I took a random selection of two hundred cases, read each case, and assigned my own label indicating whether the application for judicial review was dismissed or granted. I then used the exact same supervised method learning above, but instead of using the entire text, I only used a final part where, my own experience taught me, the Court always indicated what its decision was. After training the algorithm on 80% of the data (160 cases), I tested it on the remaining 40. Again, the algorithm was quite effective, correctly classifying 95% of all cases.

With these two models, I predicted whether each case in the dataset concerned a refugee matter and identified those applications that were dismissed or allowed. I removed all cases from the dataset for which the algorithm indicated the decision did not concern a refugee law matter. My final dataset, which began with over 15,000 decisions, now contained 3,991 decisions and totalled 11.3 million words.

5.2.3 Partitioning the decisions into clusters

One methodological hypothesis that guided this study is that similar documents, similar cases, will use similar words and expressions. A set of cases that are about similar themes and topics would, I reasoned, look similar from a semantic perspective. This intuition was born out of my experience as a lawyer. In my professional capacity, I can quickly read a decision and recognize whether it concerns a particular country or claim profile, engages a key case or a key concept, and see how it fits into the overall structure of Canadian refugee law. If I could spot these sorts of patterns, could a computer as well?

The primary difficulty with this task is developing (perhaps divining is a better word) a classification schema. While it is possible to separate cases out where the application was dismissed or granted or whether it concerned a refugee matter or not, jurisprudence is not so binary. Even as a subject matter expert, I cannot simply say that there are a certain number of categories of refugee law decisions, nor can I reasonably estimate how decisions might fall into different categories. Just as a student of literature

would be at a loss if they were asked to identify all the different categories of literature, I also am at a loss to define all the categories of refugee law.

This means that supervised learning will not do. An example reiterates the point and explains why. In *Computational Formalism: Art History and Machine Learning*, Amanda Wasielewski reports on a study undertaken by machine learning computer scientists. A research group sought to develop an artificial intelligence model that could distinguish between different genres of paintings. To train the algorithm the scientists showed it a large volume of pictures with genre labels assigned by various experts. At the end of the project, they reported that their model could accurately predict the genre of different paintings, although there were a few areas where it struggled. One was distinguishing between the paintings that were *abstract expressionism* and others that were from the *action painting* genre. The scientists ended their report hypothesizing that a new method might later be able to accurately distinguish between the two genres. But to Wasielewski, who is by training an art historian, this confusion made perfect sense: the distinction between abstract expressionist images and action painting images was the feature of a dispute among critics. Some believed that a subset of images that might be impressionist ought to be classified as action art, or *vice versa*. The mistake was to believe that the labels themselves were accurate, pure, and not a site of contestation—which, of course—they often are.⁵⁵

Instead of building my own schema, I decided to let the computer build its own using an unsupervised machine learning approach. Then, after the machine built its schema, I could evaluate it to see whether its decisions helped me understand the jurisprudence better. For this task, there are many different potential algorithms and

⁵⁵ Amanda Wasielewski, *Computational Formalism: Art History and Machine Learning* (Cambridge: MIT Press, 2023) at 72-74.

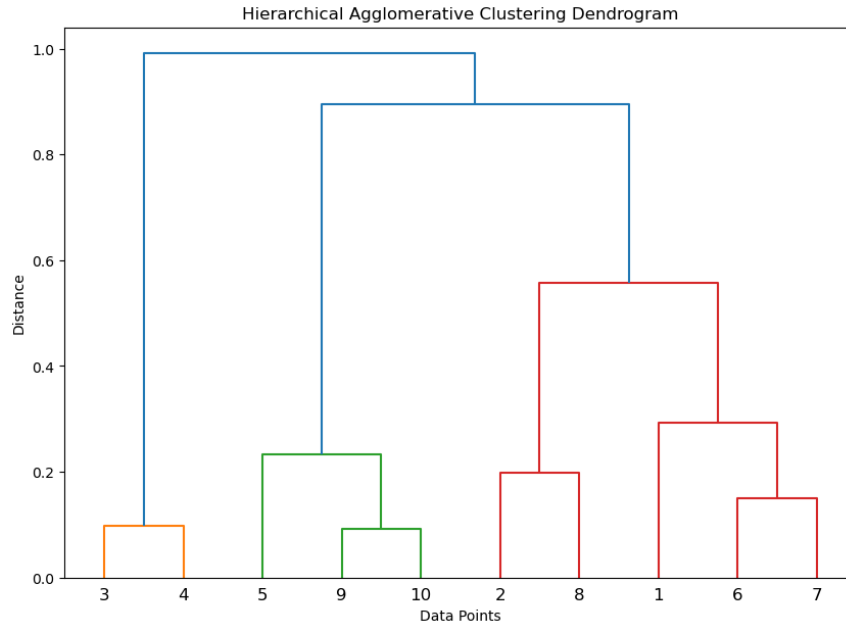


Figure 1 Example of how the agglomerative clustering algorithm groups data points together.

techniques to choose from, but I decided to use one of the simplest approaches. Hierarchical agglomerative clustering works by following a simple method. It begins by treating all data points (decisions texts in this case) as distinct. It then identifies the two data points that are most like each other and merges them into a cluster. It repeats this process over and over, until all the data points are in a single cluster. This is why this method is called hierarchical: it orders the data into categories that can contain sub-clusters.⁵⁶ Figure 1 shows an example clustering and hierarchy based on synthetic (or made-up) data for demonstration purposes. As we can see, at the bottom each individual data point exists on its own, but at each step up, the algorithm merges data points into a cluster until, at the very top, all the data points are merged. There are a number of ways of interpreting a hierarchical clustering. If we began at the very top level, we might say that there are two clusters. But if we went down one level further, we might say that there are three clusters or, one level further, four.

⁵⁶ “2.3.6. Hierarchical clustering.” scikit-learn User Guide. 2024. online: https://scikit-learn.org/stable/unsupervised_learning.html#hierarchical-clustering.

Figuring out how to interpret a clustering and how many clusters to ask the model to return is largely a qualitative exercise. The researcher must inspect the algorithmic returns and determine whether they are useful and responsive to the research question. It may well be, and as we shall soon see, that the computer picks up on patterns in the text that are not of intrinsic interest to the researcher. It is only by working with the data and taking responsibility for its use that a researcher can say whether a particular clustering is useful or not.

That said, there are quantitative metrics that can assist. In this case, I used a calculation called a silhouette score to assist me as I determined the most useful number of clusters to study. A silhouette score can be used to assess the coherence of a cluster and to determine whether a data point is closely related to the other members of a cluster, or whether it might be closer to data points in other clusters. If this sounds complex, it should not because the math is in fact quite simple. The score is calculated by (i) calculating the average distance between a data point and all the other data points in the same cluster and subtracting that figure from (ii) the average distance between the single data point and all the data points in the next closest cluster, before (iii) dividing that figure by the larger of (i) or (ii). The closer the score is to 1, the more defined and distinct the cluster. To get a sense of the overall coherence of a cluster, all the silhouette scores can then be averaged.⁵⁷

So now let us put all these pieces together. For this study, I used the tf-idf algorithm to redescribe each legal decision in numerical terms. Then I used the agglomerative clustering algorithm to create the hierarchical clustering. Before examining the data, I calculated the average silhouette scores for different numbers of clusters, ranging from 3 clusters to twenty-four clusters. In general, the silhouette scores were low, indicating that many data points were close together and that there may be some overlap between the clusters. But this did not surprise me because this is also my subjective assessment of refugee law jurisprudence: the cases written by subject-matter expert judges, trying to rationalize discrete phenomena in terms recognizable to law, are bound to use similar language and similar recurring patterns. Nonetheless, I observed that the

⁵⁷ “2.3.11.5. Silhouette Coefficient,” *Ibid.*

silhouette scores began to rise when the algorithm was asked to return eighteen clusters and, as we shall see, these clusters lent themselves to a qualitative analysis that was responsive to my research question. I therefore decided to explore how the algorithm divided the jurisprudence into eighteen different zones.

5.2.4 Visualizing multi-dimensional data

If data-informed scholarship on text has a look, it is visual. As Franco Moretti and Oleg Sobchuk observe: “Histograms, scatterplots, time series, diagrams, networks . . . ten, fifteen years ago, studies of film, music, literature or art didn’t use any of these. Now they do...”⁵⁸ But visualizing text is a challenging and counterintuitive endeavour. Imagine a simple scatter plot or a line graph. At the bottom, the x-axis, one value is represented; along the side, the y-axis, there is a second value. This visualization is a two-dimensional plot because, well, it visually shows how two dimensions, two values, relate to each other.

The difficulty with text-based documents is that they are definitionally multi-dimensional. Consider the problem mathematically: each word is a variable, the way the words are ordered are in turn new variables, the use of punctuation, allegory, and metaphor each increase the variance of a text. On the one hand, text is extraordinarily complex, and it is, quite frankly, astonishing how we have developed systems to communicate with it: out of a massive number of permutations of information, we have divined ways to order massive and complex data in specific ways so that others may understand it. Given the potential variance in a text, it seems to be beyond visualization: we can understand a plot with two, perhaps three dimensions, but not more and certainly not the billions or trillions of dimensions that might come together to form a text.

Yet, and now we look to the other hand, writing and reading works because we can use, understand, replicate, modify, and develop sensible patterns. And patterns can be mathematically understood and mathematically synthesized, sometimes all the way down to two or three dimensions. Millions of words can, and here is the bottom line, find

⁵⁸ Franco Moretti & Oleg Sobchuk, “Hidden in Plain Sight” (2019) 119 *New Left Review* 86 at 86.

representation in a simple plot. Yes, much meaning will be lost in the transition, but some new perspectives will be gained.

So, if a visualization can be found, it can be useful. For this project, visualizations serve two primary purposes. The first is the most intuitive. By plotting some trends I've spotted, we can see how they change, develop, or grow over time (i.e. the proportion of cases that were granted each year). These simple line graphs are familiar to each of us. The second and more complex visualization used in this study aims to produce a map of Canadian refugee law jurisprudence. Now that we are familiar with the idea that we can transform a text into numbers, and that we can group those texts based on distance, we can sense that it might be possible to simplify those representations of distance.

This may sound complex, but it is not: once each text is transformed into numbers we can appreciate that it becomes possible to calculate how close texts are to each other. Just as we can use clustering algorithms to locate texts that are similar to each other, we can use algorithms to spatialize different documents to produce a visual representation of where different texts might be in relation to each other. To do this, I followed a two-step process. First, I took the same tf-idf values that I used for the clustering. Recall, each value indicated the relative importance of a discrete word or term found in a decision, with higher values signalling that a term or word was more important for an individual text and lower values indicating it was less important to that text. For this study, the matrix of values was large. Each document was its own row (so there were 3,991 rows) and each the most important terms and words each had their own columns (of which there were 1,990). If this matrix were a spreadsheet, it would have 7,942,090 cells. This is an unwieldy amount of information, even for a computer to process and to make it more manageable I used an algorithm called principal component analysis to consolidate columns that rose or fell together into new columns.⁵⁹ The point of this was to reduce the complexity of the calculations that would follow in the second step and weed out some columns that largely contributed noise to data.

⁵⁹ scikit-learn, “sklearn.decomposition.PCA”, online: <<https://scikit-learn.org/stable/modules/generated/sklearn.decomposition.PCA.html>>.

Second, I used a method developed by Geoffrey Hinton, the so-called Godfather of Artificial Intelligence, to reduce this complex matrix to two, visualizable dimensions. This technique, known as t-SNE,⁶⁰ is impressionistic. Not only does the researcher feed in the data, but they must also set a few parameters, each of which will change the nature of the final visualization. By turning some knobs up or down, different pictures or arrangements of the data can present themselves. Yet, as we shall soon see, in this case the method presented a visualization that largely coordinated with the clustering. This resulting map—which, just like every other map, is a subjective representation of the world—shows us a new way of looking at law and helps us better see where clusters are diffuse and where they are tightly formed. It will also gesture, though not empirically prove, to a few core dynamics that haunt Canadian refugee law.

5.3 Findings

5.3.1 Canadian refugee law jurisprudence is increasingly claimant-negative

In general, and at least from a macro perspective, the Canadian refugee law system has been good to claimants over the past decade. Overall, the rate at which applications for protection were granted at the first-instance refugee protection hearing rose from 47% in 2013 to 79.5% in 2023.⁶¹ At the same time, the system scaled. In 2013, the first-instance tribunal (the RPD) declared that 6,159 people required protection; in 2023, it found that 37,222 people needed protection.⁶² During the same period, the appeal tribunal's

⁶⁰ scikit-learn, “sklearn.manifold.TSNE”, online: <<https://scikit-learn.org/stable/modules/generated/sklearn.manifold.TSNE.html>>.

⁶¹ These numbers are drawn from two sources: Sean Rehaag, “2013 Refugee Claim Data and IRB Member Recognition Rates” (14 April 2014), online: ccrweb.ca/en/2013-refugee-claim-data; and Sean Rehaag, “2023 Refugee Claim Data and IRB Member Recognition Rates” (29 April 2024), online: <https://refugeelab.ca/refugee-claim-data-2023>. Grant rates are calculated by dividing all positive cases by all cases that reached adjudication where a final decision on the merits was reached.

⁶² *Ibid.*

allowance rate increased from 28% to 38%, while the total number of cases the RAD heard increased from 414 to 9,334.⁶³

Statistics compiled by Professor Sean Rehaag and the Refugee Law Lab show that claimants also enjoyed relatively consistent success before the Federal Court of Canada. While in 2013, only 14% of applications for leave were granted, this rose to 34% in 2015, before settling in at around 25% in 2017. For applications that were granted leave, the overall grant rate between 2013 and 2023 hovered between 40% and 57%.⁶⁴ Putting it all together, refugee claimants in Canada are getting protection at higher rates than ever before, at higher numbers than ever before, and doing better with respect to applications for leave than ever before. And even for those who lose their cases, the Federal Court of Canada has granted reviews at roughly the same rate, even though the overall rate of refusals is down.

Yet, and here is the first finding of this study, refugee law *jurisprudence* is increasingly claimant-skeptical. Going into this study, I presumed that the decision texts published by the Court would match, in terms of volume and in terms of outcomes, the Court's overall work. If, for example, each year the Court released 300 refugee law decisions and granted judicial reviews in half of them, I assumed that there would be 300 cases, of which 150 would be positive decisions. This assumption was wrong. Using the assessments of outcomes of the jurisprudence assessed by my machine learning model, I

⁶³ Immigration and Refugee Board of Canada, "Refugee appeals statistics" (15 May 2024), online: Immigration and Refugee Board of Canada <<https://www.irb-cisr.gc.ca/en/statistics/appeals/Pages/index.aspx>>.

⁶⁴ Refugee Law Lab Portal, "Federal Court" (visited 13 June 2024), online: Refugee Law Lab <<https://rllp.ca/federal-court>>.

plotted the annual grant rate of the Court’s jurisprudence. Figure 2 shows a general, though somewhat bumpy, downwards trend. Between 2013 and 2015, the grant rate of all the cases published by the Court approached 50%, but by 2023 it was 28.5%, well below the



Figure 2: For refugee law cases that the Court released a public decision for, the grant rate has trended down.

overall grant rate of 40% for that year assessed by Professor Rehaag. To state it clearly, the jurisprudence of the Court—the body of law it produces—is less claimant positive than the overall decision-making posture of the Court.

What accounts for this discord between the jurisprudence and the Court’s work? It appears to be the product of two interlocking phenomena. First, my assessment of the data says that it matters whether a case has been scrutinized once or twice before reaching the Court. Recall that some, but not all claimants, are entitled to appeal negative first instance decisions to another tribunal with the Immigration and Refugee Board. Those claimants who may not bring their cases to the RAD—mainly those whose claims were found to have no credible basis, or were manifestly unfounded, or who were allowed to make a claim at a land border crossing between Canada and the United States—may only apply for leave and for judicial review. It appears that in cases not reviewed first by the RAD (i.e. those cases that come to the Court directly from the RPD) are, if they are granted leave, more likely to receive positive final adjudication from the Court.

In 2018 and 2019, the Federal Court’s registry began to categorize claims that originated from the first-instance tribunal separately from those that came to the Court after second-level appellate review. As Figure 3 shows, the overall grant rate for cases that




Figure 3: In 2018 the Court began to classify cases that originated from the RAD differently than those cases that originated from the RPD. The Court’s jurisprudence suggests that the Court grants RAD reviews at a lower rate than RPD reviews.

went through two levels of review was markedly lower than cases that came directly from the first-instance tribunal. On its own this is, at least as I read the data and understand the point of appellate review, an indicator of overall system health. Our expectation should be that cases that have gone through two levels of review feature, at least when compared to cases that have only gone through one level of review, fewer problems warranting judicial intervention. When we think back to the overall grant rate within the jurisprudence, and we appreciate that each year more cases are coming from the RAD, we should therefore be unsurprised that the overall grant rate is decreasing.

This dynamic on its own, however, fails to explain why the jurisprudential grant rate is lower than the Court’s overall grant rate identified by Professor Rehaag’s Refugee Law Lab. To explore this question, Professor Rehaag forwarded me a sample of entries

(n=688) from an unpublished dataset that tracks outcomes in judicial review cases.⁶⁵ After comparing the jurisprudential dataset with Professor Rehaag's dataset, I determined that there were 179 instances in the overall sample where Professor Rehaag had information about a case that was not in my jurisprudential dataset. In 24 of those cases, the reason for this discrepancy seemed clear: in some cases, Professor Rehaag had information regarding French language cases that had not yet been translated, and therefore were not included in my dataset. In other cases, the discrepancy arose from the fact that the Court sometimes heard and decided multiple files at the same time, and later issued a single decision that spoke to several different matters. This means that in Professor Rehaag's dataset, each *file* is counted while in the dataset under study here, each *decision* is counted, no matter how many individual files the one decision speaks to.

The remaining 155 unmatched cases were more interesting. In each of these cases, Professor Rehaag's analysis showed that the matter resolved on consent. A consent resolution happens when both the claimant and the government agree that there is an error or problem in the decision that warrants reconsideration and so the parties jointly ask the Court to issue an order granting the judicial review. When this happens, the Court rarely issues a set of reasons explaining its decision because it is essentially endorsing a settlement reached by the parties. This explains the largest part of the discrepancy between Professor Rehaag's overall statistics about the Court's resolution patterns and the resolution patterns in the jurisprudential dataset: the jurisprudential dataset simply does not and cannot account for settlements.

This explains why the Court's jurisprudence is disproportionately composed, relative to its overall work product, of refused applications for judicial review: a good chunk of the Court's positive decisions are in fact settlements that will never form part of a jurisprudence. While it appears that the discrepancy between Professor Rehaag's dataset and the one developed here is growing each year, I leave for another day and for another study the question of whether the government is agreeing to settlements at increasing rates

⁶⁵ This is the dataset that was originally developed (and has since been updated) for Sean Rehaag, "Judicial Review of Refugee Determinations (II): Revisiting the Luck of the Draw" (2019) 45:1 Queens LJ 1.

and whether it is only lately that its settlement practices are tangibly changing the shape of the Court’s jurisprudence. I note in passing that the Court introduced a pilot policy in 2019 designed to encourage settlement in immigration matters.⁶⁶ Perhaps we are looking at evidence that this pilot has been successful.

One final observation. Focusing only on cases where a judge issued a set of reasons enriches our analyses of judicial behaviour. Figure 4 shows the grant rates per year for each judge of the Court. Overall, there appears to be a gradual downward slope in the grant rate but, of equal interest, we see that there is considerable variability. Save for a few outliers, the grant rates of different judges change markedly between years.

To summarize, this exploration of the grant rates of the refugee cases issued by the Court, its jurisprudence, was enabled by a machine learning technique to extract the judge’s final decision. This method showed that refugee law jurisprudence is getting more claimant negative over time because it is tending to be composed of proportionally fewer




Figure 4: Each line shows the refugee law grant rate of a different judge, when only looking at cases that resulted in the issuance of a final decision.

allowed applications each year. This was surprising because the jurisprudential grant rate

⁶⁶ Federal Court of Canada, “Notice: IMM Settlement Pilot (Toronto)” (4 July 2019), online: [https://www.fct-cf.gc.ca/content/assets/pdf/base/Notice%20IMM%20settlement%20pilot%20\(Toronto\)%20-%20AMENDED%20July%204-2019%20FINAL%20ENG.pdf](https://www.fct-cf.gc.ca/content/assets/pdf/base/Notice%20IMM%20settlement%20pilot%20(Toronto)%20-%20AMENDED%20July%204-2019%20FINAL%20ENG.pdf).

appears to be dropping, while the Court's overall grant rate is thought to be stable. But upon closer examination we saw that this discrepancy between the Court's overall work and its jurisprudence is a feature of settlement practices. While we cannot know whether the Court would have granted the application for judicial review in every case where the parties reached a settlement, it is probable that it would have many of those times. Settlement, especially settlement at scale, therefore, appears to produce a jurisprudence that presents as less favourable to claimants, despite what the on-the-ground reality might be.

5.3.2 Visualizing and clustering refugee law jurisprudence shows that it is not topically diverse

5.3.2.1 Mapping Canadian refugee law

Let us now turn to some maps. Figure 5 shows all immigration law (not just refugee law) decisions in the dataset issued by the Court between 2013 and 2023. Every dot is a decision. On the left, darker coloured dots represent decisions issued closer to 2023 while lighter dots represent decisions issued closer to 2013. On the right, the colour of the dot corresponds to the classification of the decision: does it concern visa denial (dark green)? An emergency application to stop a deportation because of a new or unassessed risk (turquoise)? An application to immigrate to Canada on humanitarian grounds (orange)? A first instance refugee appeal (pink and red) or a second level refugee appeal (purple)?




Figure 5: Every dot is an immigration law decision. On the left, darker dots represent cases closer to 2023. On the right, the different colours represent the classifications made by the Court. We can see clearly that the computationally produced maps represent chronological and subject matter patterns.

Both maps (Figure 5) were created using the visualization method described above: each decision was transformed into numbers (tf-idf), columns that rose or fell together were merged (principal component analysis), and then the relationships between all the dots were visualized in two dimensions (t-sne). Without much analysis we can see that this algorithmic approach displays intuitive and sensible patterns in the data. Decisions that were issued around the same time often group together. Decisions that concern the same legal subject matters also group together. Refugee law decisions separate themselves into a separate island, with first instance and second instance decisions in turn separated by a small gulf. Let us pause: millions and millions of words mapped into two dimensions—and somehow some sense remains. Despite all the variation in each text, the algorithm arranged the data in a way that recognized and distinguished between legal classifications of matters.

Now a second set of maps (Figure 5). These maps show only refugee law decisions and, with fewer data points and fewer relationships, the presentation has changed shape. Again, on the left, the data points' colour refers the year the decision was issued. And, again, chronological patterns are evident: more recent decisions cluster together, within clusters we see a time-based movement, and newer decisions are focused in the top-left quadrant of the map.

On the right, we see a relationship between the decisions and geography. I wrote a small function that counted the number of times each decision referred to a country and

identified the most referred to country. A second function then grouped decisions by continent. Blue decisions refer to cases that where the country most referenced was in Asia; orange, countries from Europe; purple, Africa; green, North America; red, South America. On the one hand, everything looks a bit jumbled up; on the other, there are clear spots where each decision refers most to an Asian country, and one clear spot where decisions engaged most with European and South American countries. Comparing the two maps, European cases are clustered in regions that are closer to 2013 while African and Asian cases are newer.




Figure 6: These maps show all refugee law decisions (2013-2023). On the left, darker dots show cases that were decided closer to 2023. On the right, the different colours represent the countries (grouped together by continent) that are most mentioned in each decision.

5.3.2.2 Clusters of refugee law

Finally, a third map. Figure 7 shows the same map, but now the colours refer to one of the eighteen clusters identified by the agglomerative clustering algorithm. These clusters fall into three categories: (i) not descriptive, (ii) idiosyncratic, and (iii) broadly thematic.

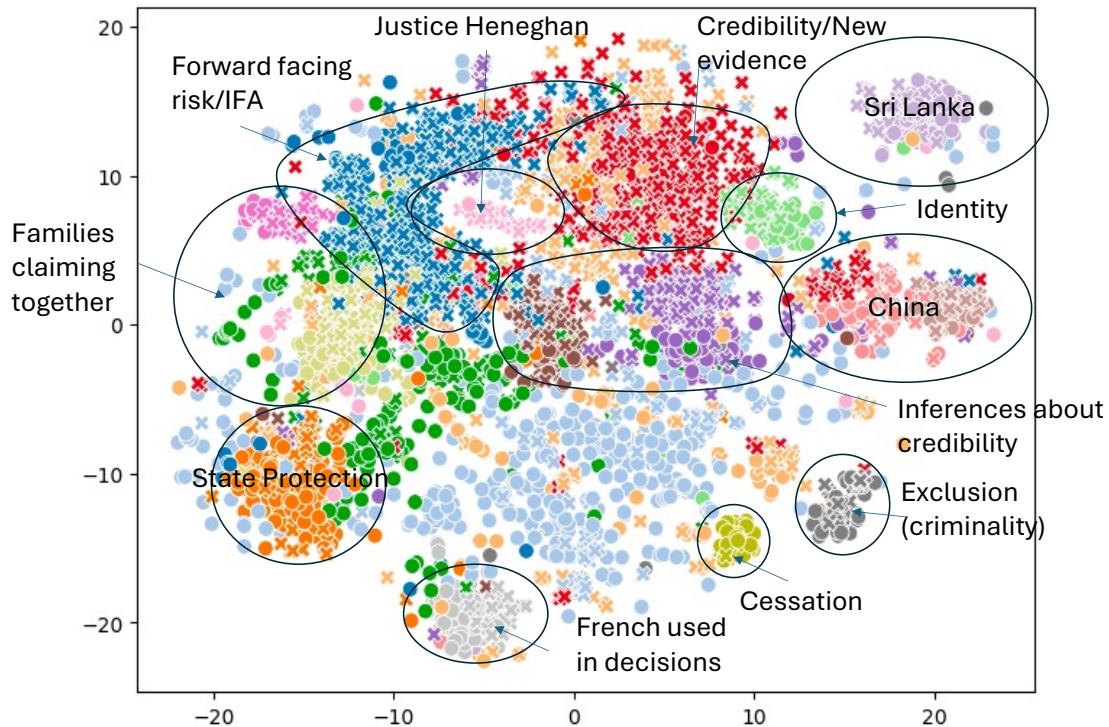


Figure 7: This map shows all refugee law decisions. Each colour corresponds to a different cluster identified by the Hierarchical Agglomerative Clustering algorithm. I then labelled different regions of the map based on my subjective analysis of the content of each region. Circles represent decisions the originated from the RPD, “x”’s represent decisions from the RAD.

I begin with the not descriptive categories. Recall that for each cluster I took an average silhouette score, a measure designed to assess whether each decision’s cluster assignment was useful. There were three clusters that returned average negative scores, indicating that many decisions occupied liminal spots and were close to another cluster assignment. Those three clusters are the light blue, light orange, and dark green clusters on the map. The decisions in these clusters span a large part of the map, with individual decisions commingled with other clusters. What is to be made of these large, decentralized clusters? At bottom, not much. Had I directed the clustering algorithm to search for more than eighteen clusters it is likely that these diffuse clusters would fracture into several discrete, and more interpretable, categories. But with classification studies, this is always the trade-off: the more particular our categories, the less useful they become. All that can really be said is that here we see that my decision to search for eighteen categories returned three residual categories of documents that evaded classification. What this ought to remind us of is some of the limitations in this approach. Sometimes we wrongly assume

that computers return strict and certain quantitative information, but these residual categories remind us that in this case we are reviewing something much more impressionistic.

Four clusters are idiosyncratic. As discussed, in some cases the algorithm was attracted to patterns in the text that, while interesting, are not necessarily responsive to the research questions I posed. For example, after sampling five decisions⁶⁷ from the light grey cluster I noted that each case shared the same linguistic feature. Sometimes the judges of the Court quote statutory sections and include both the English and French language versions of the section. To the algorithm, this writing decision looked meaningful likely because it recognized that in a small subset of decisions French language words were used. Had the research question for this project been “what proportion of decisions include bilingual quotations of statute?”, this method would have been remarkably effective (the answer is 4.11%) but, for the matters that interest me here, this cluster does not appear to have much value.

Another cluster that is intrinsically interesting, but that is non-responsive to the matters under study, is the light pink cluster. All 68 decisions in the light pink cluster were authored by Justice Heneghan. After reviewing a sample of these decisions⁶⁸ it became clear that Justice Heneghan has a distinct and parsimonious writing style. Compared to her colleagues, who tended to write at greater length and using a wider and more legalistic vocabulary, Justice Heneghan’s decisions looked to the algorithm to be linguistically distinct. The extent to which her decisions are unique can be quantitatively demonstrated. Figure 8 shows us the relationship between two metrics: the mean number of words used

⁶⁷ *Buwu v. Canada (Citizenship and Immigration)*, 2013 FC 850; *Diallo v. Canada (Citizenship and Immigration)*, 2014 FC 471; *Tamayo v. Canada (Citizenship and Immigration)*, 2015 FC 1217; *Qaddafi v. Canada (Citizenship and Immigration)*, 2016 FC 629; *Kouassi v. Canada (Citizenship and Immigration)*, 2016 FC 1398.

⁶⁸ *Madava v. Canada (Citizenship and Immigration)*, 2017 FC 149; *Glonczi v. Canada (Citizenship and Immigration)*, 2019 FC 931; *Kawekwune v. Canada (Citizenship and Immigration)*, 2020 FC 333; *Okuyemi v. Canada (Citizenship and Immigration)*, 2023 FC 808; and *Chapagain v. Canada (Citizenship and Immigration)*, 2023 FC 949.

by each judge and the mean entropy of each judge’s decision. Entropy is a measure of the unpredictability or randomness in a text, with higher scores indicating greater linguistic




Figure 8: This plot shows the mean entropy and mean word length for each judge’s decisions. Higher entropy scores indicate greater linguistic complexity.

variability and complexity.⁶⁹ As can be seen, Justice Heneghan’s decisions are an outlier in terms of length and entropy. For a future study, it would be interesting to determine whether entropy and word length are predictors of outcome or, perhaps, the likelihood of a case being cited as authority by another judge. Perhaps, too, there are different models of the judge presented in this data: do low entropy judges immediately capture the essence of the case in their writing? Do high entropy judges engage with a richer variety of ideas, sources, and materials? Questions for another day.

Two related clusters round out this description of idiosyncratic clusters. The light green and dark pink clusters are perhaps the most interesting and surprising clusters. Both clusters featured high tf-idf scores for the expression “pa” or “principal applicant.” In Canadian refugee law, this term is often used when a family comes to Canada and claims protection together but relies on the evidence and experiences of one family member. An

⁶⁹ For a discussion on the relationship between the overall length of a text and its entropy see Yaqian Shi & Lei Lei, “Lexical Richness and Text Length: An Entropy-Based Perspective” (2022) 29:1 J Quantitative Linguistics 62.

excerpt from *Ibrahim v. Canada (Citizenship and Immigration)*⁷⁰ demonstrates how this language is used:

[4] The Applicants are a family of four. The Principal Applicant, Abouzaid Mohamed Abdelhamid Ibrahim, is a 38-year-old citizen of Egypt. The Principal Applicant's spouse, Asmaa Mohamed Zein Ali (the "Associate Applicant"), is 29 years old and also a citizen of Egypt. They have two Egyptian-born children, ages 11 and 5. The Applicants have a third child, age 2, who was born after their arrival in Canada.

...

[7] On June 9, 2019, the Principal Applicant alleges that he received a threatening call from Mostafa, telling him that this was the last chance to sell the land to him, or he would "destroy" his life. The Principal Applicant filed a complaint with the police the following day. A week later, the Principal Applicant followed up with the police, but was told to drop the matter and that nothing further would be done.⁷¹

From an algorithmic perspective, 6.8% of the decisions in the dataset used similar linguistic patterns to describe the fact of a family unit claiming protection, and advancing an application for judicial review, together. Like the decisions that used French, it appears that the algorithm noticed that for a subset of decisions the judge writing the case used unique and particular language. From my perspective, this feature of the data is not terribly interesting nor useful for the study.

This leaves eleven substantive categories for analysis. These remaining categories were the largest, collectively accounting for 60% of all decisions. Beginning at the bottom right of the map the light green cluster concerned cases engaged cases flowing from a unique type of allegation the Canadian government made. In rare cases, a person's refugee status can be revoked (it will "cease") if the government can show that the reasons the person required protection in the first place have disappeared, they have returned and "reavailed" themselves of protection in their country, or they have developed an alternative form of protection.⁷² The light grey cluster concerned a form of "exclusion."

⁷⁰ 2022 FC 1663.

⁷¹ *Ibid.*

⁷² IRPA, *supra* note 1, s 108.

Some people, who because of their misconduct or engagement in criminality, are disentitled to (“excluded”) the right to seek protection.⁷³ Both clusters were quite small, accounting for 1.6% and 2.1% of all cases. Nonetheless, these clusters help us understand one key feature of the nature of Canadian refugee jurisprudence: it is, but only in a small sense, about that part of the law that declares that a person no longer requires protection or that part that concludes a person ought to be excluded from protection because of their conduct.

Moving to the top right of the map, we find several communities that are linked by, not necessarily a common legal issue, but a common risk profile. The cases in the light




Figure 9: Orange dots indicate that decisions used the word “Tamil” and, as we can see, there is a concentration of Tamil cases in the region of the map with the light purple cluster.

purple cluster are united by the fact that each concerns a claim for protection made by a Tamil citizen of Sri Lanka. Indeed, in Figure 9 I show which decisions used the word “Tamil” and we can clearly see that these decisions are focused in the light purple cluster. What are these cases about? I reviewed a random sample of five cases.⁷⁴ In each case, the

⁷³ *Ibid*, s 98.

⁷⁴ *S.A. v. Canada (Citizenship and Immigration)*, 2014 FC 146; *Sriratham v. Canada (Citizenship and Immigration)*, 2014 FC 614; *Sivagnanasundarampillai v. Canada (Citizenship and Immigration)*, 2018 FC 1109; *Jeyaredsagathas v. Canada (Citizenship and Immigration)*, 2018 FC 1238; and *Arumaithurai v. Canada (Citizenship and Immigration)*, 2022 FC 604

claimant alleged that they would be persecuted by the Sri Lankan government because of their Tamil ethnicity. Notably, cases in this cluster feature one of the highest grant rates (47%), suggesting that this is an area of the law that, at least for a time, was of interest to at least some members of the Court.

Immediately below the cluster concerning Tamil individuals from Sri Lanka, there were two clusters concerning nationals of China. As Figure 10 shows, cases that mentioned the word China were focused in this region of the map, where there were two clusters: one light red and one light brown. The light red cluster appeared to primarily concern cases where individuals said that they would be persecuted by the Chinese state




Figure 10: Orange dots indicate that that the word “China” was used in the decision text.

because of their Christian faith. To understand this cluster (and for the other clusters) I averaged all the tf-idf scores for the decisions within the cluster and identified the ten highest scores. In this case, those terms clearly speak to major themes in the cluster: “china,” “psb,” “summons,” “church,” “rad,” “applicants,” “arrest,” “chinese,” “christian,” and “fraudulent.” Of the five cases I sampled from this cluster, three concerned individuals who said that they needed protection because they were going to be persecuted because of their Christian faith,⁷⁵ one concerned a Falun Gong practitioner,⁷⁶

⁷⁵ *Chen v. Canada (Citizenship and Immigration)*, 2013 FC 928; *Zheng v. Canada (Citizenship and Immigration)*, 2015 FC 16; and *Hui v. Canada (Citizenship and Immigration)*, 2020 FC 446.

⁷⁶ *Yang v. Canada (Citizenship and Immigration)*, 2016 FC 543.

and the final involved a person who said he participated in protests against the Chinese state.⁷⁷ Each case was united by the refugee tribunal’s reasoning process: in all five cases, the tribunal concluded that the person was advancing a fraudulent claim based on incredible information.

The light brown cluster was closely related. As the top ten tf-idf terms show, and as my sample of five cases confirmed,⁷⁸ this cluster focused on Falun Gong practitioners: “falun,” “falun gong,” “gong,” “china,” “psb,” “rad,” “practice,” “summons,” “practitioner,” and “applicants.” This excerpt gives a good example of the type of case in this cluster:

The Applicant, Chenchen Mao, is a citizen of China. His story is as follows: In February 2015, while living in China, he began practicing Falun Gong because he suffered from depression and insomnia. He joined an underground Falun Gong practice group two weeks later. In September 2015, the Public Security Bureau [PSB] raided his Falun Gong practice group. The Applicant was one of the lookouts for the group on that day and fled to his aunt’s home. The PSB looked for him at his home and left a Chinese subpoena/summons [Chuanpiao] at his home that accused him of involvement in illegal Falun Gong activities and recruiting members for an illegal organization.⁷⁹

In this case, the RPD and the RAD both concluded that the claimant was not truthfully testifying and that there was not reliable and credible evidence substantiating his allegation that he would be persecuted. The Court dismissed, as it did for almost 70% of the cases in this cluster, the claimant’s application for judicial review.

Reviewing these two clusters allows us to better understand Canadian refugee law jurisprudence. Most claims made against the Chinese state that end up before the Court appear to engage one of two risk profiles: that the person will be persecuted for the Christian faith or because of their status as a Falun Gong practitioner. For both these risk

⁷⁷ *Ni v. Canada (Citizenship and Immigration)*, 2018 FC 948.

⁷⁸ *Xu v. Canada (Citizenship and Immigration)*, 2014 FC 1062; *Liang v. Canada (Citizenship and Immigration)*, 2019 FC 58; *Huang v. Canada (Citizenship and Immigration)*, 2019 FC 1366; *Mao v. Canada (Citizenship and Immigration)*, 2020 FC 542; and *Shi v. Canada (Citizenship and Immigration)*, 2022 FC 196.

⁷⁹ *Mao v. Canada (Citizenship and Immigration)*, 2020 FC 542 at para 4.

profiles, the main question the Court adjudicated was whether the refugee tribunal's assessment of the claimant's credibility was reasonable or not. Of course, these were not the only claims made by Chinese nationals that were considered by the Court because, as we can see and as we will soon discuss, there are several other regions of the map where we can—at least at this stage—confidently say that judges are referencing “China” in their decisions.

But for now, let us shift our gaze slightly to the left, and examine the light green cluster next to the two clusters concerning Chinese nationals seeking protection. Claimants are positively obliged by law to prove with documentation who they are or, in the absence of documentation, provide a reasonable explanation why they do not have documentation that can establish their identity.⁸⁰ As the top tf-idf terms suggest, this cluster is primarily concerned with the quality of identity documents proffered by claimants: “identity,” “rad,” “documents,” “birth,” “certificate,” “passport,” “birth certificate,” “card,” “identity documents,” “document,” and “establish identity.” When I examined the countries most mentioned in each decision in this cluster I found that 11.5% of cases most mentioned Eritrea, with China (11%), Nigeria (7%), Somalia (7%), Guinea (4%) and Bangladesh (4%) following. And, sure enough, in each of the five cases I sampled, the determinative issue was whether the claimant led sufficient evidence to establish that they were who they claimed to be.⁸¹ This cluster was small (accounting for 2.8% of all decisions) but, just like the cessation and exclusion clusters, speaks to a consistent concern in Canadian refugee law jurisprudence.

It strikes me as noteworthy that the identity cluster is close to both clusters concerning Chinese nationals and the next three large clusters (the dark red, dark purple, and dark brown clusters) we will examine. In each of these clusters the main issue assessed by the Court was evidentiary and concerned the claimant's credibility. Unlike the identity

⁸⁰ *IRPA*, *supra* note 1, s 106.

⁸¹ *Mishel v. Canada (Citizenship and Immigration)*, 2015 FC 226; *Lhamo v. Canada (Citizenship and Immigration)*, 2016 FC 873; *Daniel v. Canada (Citizenship and Immigration)*, 2016 FC 1049; *Shen v. Canada (Citizenship and Immigration)*, 2017 FC 885; and *Trawally v. Canada (Citizenship and Immigration)*, 2023 FC 1627.

cluster, unlike the Chinese Christian cluster, and unlike the Chinese Falun Gong cluster, these next clusters are not so clearly thematically defined. As we can see in the visualization, the borders between the dark brown, dark purple, and dark red clusters are not clearly delineated, at least when compared to the clusters we have examined so far. All three also have lower average silhouette scores, suggesting again that each cluster is more diffuse and less distinct. Nonetheless, some commonalities emerge.

Decisions in the dark red cluster, for example, appear to be primarily defined by their relationship with one procedural point of Canadian refugee law. Claimants who can appeal refused decisions to the appeal body (the RAD) are entitled to lead “new evidence” on their appeals. There are stringent legal criteria that govern the admissibility of new evidence. A claimant may “present only new evidence that arose after the rejection or was not reasonably available, or that the applicant could not reasonably have been expected in




Figure 11: Orange dots indicate that the term “new evidence” was used in the decision text.

the circumstances to have presented, at the time of the rejection...”⁸² The top three tf-idf terms for this cluster (“rad,” “new evidence,” and “letter”) imply that a major and consistent issue is the RAD’s consideration of a letter as new evidence. Indeed, when we visualize the location of the term “new evidence” we see that decisions containing the

⁸² *IRPA*, *supra* note 1, s 110.

term are focused in the region of the dark red cluster. This said, this cluster is clearly about much more. In the five decisions I sampled,⁸³ only one considered new evidence⁸⁴ and this one case turned on the value of new evidence led by the government, not the claimant. In each case, however, a major issue was the overall credibility of the claimant.

The word “credibility” was so common across the decisions that when I transformed each decision into its numerical representation, the tf-idf algorithm did not include the term as a discrete column because I had directed it to exclude commonly




Figure 12: Darker colours indicate higher tf-idf scores the word “credibility.”

occurring words. I reran the algorithm and directed it to include all terms, no matter how frequently they appeared. Figure 12 shows that high tf-idf scores for the word “credibility” are focused in a few regions. As we can see, the entire dark red zone, as well as those areas associated with claims from China, show elevated tf-idf scores for the word “credibility.” Together, this suggests that the overall thrust of the cluster concerns credibility and that the adjudication of admissibility claims regarding new evidence

⁸³ *Liu v. Canada (Citizenship and Immigration)*, 2016 FC 811; *Etik v. Canada (Citizenship and Immigration)*, 2018 FC 175; *Ciro v. Canada (Citizenship and Immigration)*, 2021 FC 1383; *Ahmed v. Canada (Citizenship and Immigration)*, 2023 FC 830; and *Singh v. Canada (Citizenship and Immigration)*, 2023 FC 1106.

⁸⁴ *Etik v. Canada (Citizenship and Immigration)*, 2018 FC 175.

represents a subset of those cases. Perhaps this means that claimants seek the admission of new evidence on appeal primarily to rebut concerns about credibility.

The next two clusters, the dark purple and dark brown, are the two that most evade easy classification. In each case, the top mean tf-idf scores do not illuminate a general thrust of the clusters. After sampling five decisions from each cluster,⁸⁵ I observed that in most cases the Court was evaluating a decision where the decision maker concluded that there was not convincing evidence that the person's allegations were true. To the extent that there was a common thread, it appeared to be stylistic. In these cases, located right in the middle of the map, the Court often reviewed the facts in detail before considering whether the adjudicator below properly reasoned from the evidence or not, before finding the evidence wanting.

The final two clusters, in contrast, turned on clear and distinct legal issues. In the bottom left corner, the dark orange clusters concerned cases where the key legal issue was state protection. For a person to establish that they require protection, they must prove that their own state will not protect them from the risk that they face.⁸⁶ The top ten tf-idf terms for this cluster suggest a clear thematic thrust: "state protection," "state," "roma," "hungary," "applicants," "police," "rad," "adequate," "discrimination," and "mr." As we can see, the dominant terms in this cluster suggest a focus on the Hungarian state and its ability to protect Roma individuals from discrimination and the police. In fact, it appears that the cluster has a broader focus. When I sampled five cases I found that each concerned a European national, two cases concerned blood feuds,⁸⁷ two

⁸⁵ For the dark brown cluster: *Mahmoud v. Canada (Citizenship and Immigration)*, 2016 FC 1020; *Pestova v. Canada (Citizenship and Immigration)*, 2016 FC 1024; *Ramrattan v. Canada (Citizenship and Immigration)*, 2016 FC 1108; *Bushra v. Canada (Citizenship and Immigration)*, 2018 FC 896; and *Aedo Arancibia v. Canada (Citizenship and Immigration)*, 2023 FC 1245. For the dark purple cluster: *Nyembua v. Canada (Citizenship and Immigration)*, 2015 FC 970; *Singh v. Canada (Citizenship and Immigration)*, 2019 FC 1375; *Khan v. Canada (Citizenship and Immigration)*, 2020 FC 438; *Tandi v. Canada (Citizenship and Immigration)*, 2021 FC 1413; and *Gonzalez de los Santos v. Canada (Citizenship and Immigration)*, 2022 FC 1281.

⁸⁶ See *Canada (Attorney General) v. Ward*, 1993 CanLII 105 (SCC), [1993] 2 SCR 689.

⁸⁷ *Celaj v. Canada (Citizenship and Immigration)*, 2014 FC 761 and *Meshveliani v. Canada (Citizenship and Immigration)*, 2019 FC 1351.

cases concerned Roma individuals (one from the Czech Republic⁸⁸ and one from Hungary⁸⁹), and a final case concerned a Ukrainian national.⁹⁰ In every case, the capacity and willingness of the state to provide protection was a key issue. Interestingly, this cluster aligns well with the map that showed the geographic distribution of cases. Almost all the cases that most referenced a European country are in the state protection cluster, suggesting that this legal question is primarily a concern that tracks with geography.

Above that cluster, we find the dark blue cluster. The top ten tf-idf scores for this cluster were “rad,” “applicants,” “ifa,” “nigeria,” “mr,” “harcourt,” “vavilov,” “port harcourt,” “police,” and “ms.” These terms suggest that a major theme of these cases is the viability of a proposed internal flight alternative (an “ifa”), specifically Port Harcourt in Nigeria. Indeed, when I sampled five decisions,⁹¹ each turned on the viability of an internal flight alternative. What is an internal flight alternative? Canadian refugee law teaches that a person who can reasonably and safely relocate to another part of their country of nationality or normal residence, where they will also be safe, they do not require refugee protection in Canada.⁹² In one sense, the internal flight question is a cousin of the state protection question because it asks whether the state might, in another part of the country, protect the individual. It appears that, here too, there is a geographic and political dimension that is significant. Each of the top three most referenced countries (Nigeria: 29%, Mexico: 12%, and India: 11%) in this cluster features a federal constitutional structure with different states.

⁸⁸ *Banom v. Canada (Citizenship and Immigration)*, 2016 FC 502.

⁸⁹ *Djubok v. Canada (Citizenship and Immigration)*, 2014 FC 497.

⁹⁰ *Tkachuk v. Canada (Citizenship and Immigration)*, 2015 FC 672.

⁹¹ *Ruiz Lopez v. Canada (Citizenship and Immigration)*, 2021 FC 390; *Bautista Montero v. Canada (Citizenship and Immigration)*, 2021 FC 1432; *Ui Haq v. Canada (Citizenship and Immigration)*, 2022 FC 95; *Shakil Ali v. Canada (Citizenship and Immigration)*, 2023 FC 156; and *Pulido v. Canada (Citizenship and Immigration)*, 2023 FC 463

⁹² For a thorough discussion see *Ohwofasa v. Canada (Citizenship and Immigration)*, 2020 FC 266 at para 17 and ff.

Figure 13, however, suggests that this cluster is about more than just internal flight alternative analyses. As we can see, the cases that refer to the expression “internal flight alternative” are a subset of the entire dark blue cluster because the bottom portion of the




Figure 13: Decisions that reference the terms “state protection” or “internal flight alternative.”

cluster does not appear to feature the term “internal flight alternative.” The key to determining what these cases may be about requires us to return to one of the first maps we considered. Recall that when we mapped the use of the word “China” most were located in the two clusters on the far right, but a small subset were densely located at the bottom of the dark blue cluster.

Upon closer examination of these cases, most appear to be claims based on the allegation that a person is at risk for breaching (or for wanting to breach) China’s one-child policy. In these cases, individuals often alleged that they would be forced by the state to undergo an abortion and/or be sterilized for breaching the policy. Before the refugee tribunal, however, the claim for protection was often rejected because of (i) the fact “in 2016 China changed its family planning policy to adopt a two-child policy” or (ii) that “the preponderance of the documentary evidence indicate[d] that the penalty for having an out of plan child would be a fine,”⁹³ which adjudicators reasoned was not sufficiently serious to be understood as persecutory treatment warranting refugee protection. In a sense, it

⁹³ *Huang v. Canada (Citizenship and Immigration)*, 2019 FC 120 at paras 16 and 18.

makes sense to group this class of case with matters that turned on the viability of an internal flight alternative. In both instances, the core issue is not whether the claimant is incredible or unreliable, but whether they actually face risk in their country of residence or nationality.




Figure 14: Orange dots indicate the term “child policy” was used in the decision text.

5.3.2.3 Summary

This part of the study began with over eleven million words spread out over 3,991 different decisions. This volume of text and decisions could imply a tremendous diversity of legal experiences and, of course, it surely does: each decision tells the unique story of a person or family but rearticulated in terms amenable to legal analysis. It is telling that this rearticulation, at least as presented here, capably reduces the categories of interest to just a few topics. By the time most cases reach the Court, and in light of the similarity of those circumstances that led people to claim protection in Canada, they have been refined and rearticulated so much that they fit fairly neatly into only a few categories. Before the Court, cases fall into a few distinct categories: those cases that turn on the quality of the tribunal’s assessment of a person’s credibility, the quality of its state protection analysis, the quality of its internal flight alternative analysis, cases from China or Sri Lanka that look similar, cases that consider the quality of identity documents, and cases that deal with either exclusion or cessation. From extraordinary diversity, to a few regular and repeating legal

categories. Could there be further sub-division and specificity? Of course, but these high-level generalizations are a necessary place to begin before diving in further to see if there are sensible and useful subdivisions within clusters.

5.3.3 Jurisprudential change over time

Now that distinct clusters have been identified, it is possible to analyze the changes in the composition of the jurisprudence over time. Figure 15 shows the proportionate cluster




Figure 15: The proportionate cluster representation in each year under study.

breakdown for each year under study. Immediately we can see that some clusters shrunk over the time under study, while others accounted for an increasingly large proportion of the jurisprudence. First, consider the light blue, light orange, and dark green clusters. These clusters were those that the algorithm identified as residual, or catch-all, categories. We can see that, as time progressed, proportionally fewer cases were classified in residual categories.

Second, the dark red cluster primarily considered cases where the claimant sought to adduce new evidence before the Refugee Appeal Division. This cluster appears to have emerged in 2014 and eventually stabilized in 2016, accounting for roughly 20% of the jurisprudence each subsequent year. This is not a surprising finding: the appeal division

only started hearing cases in 2013, and then only started hearing them at scale in 2014. It makes sense that it would only be after the tribunal began issuing many decisions that a legal issue specific to that part of the appeal process would appear in the Court’s jurisprudence.

Third, the light purple cluster concerned cases of Tamil individuals from Sri Lanka. These cases clearly accounted for a large proportion of the jurisprudence in 2013 but shrank to a very small percentage by 2023. In contrast, Chinese Christian (light red) and Falun Gong (light brown) cases accounted for a small, but relatively constant, proportion of each year’s jurisprudence.

Finally, and perhaps the most interesting finding, is that state protection (dark orange) cases began as a large portion of the jurisprudence, and then steadily shrank. Meanwhile, internal flight alternative cases barely registered in 2013 but by 2023 accounted for more than 20% of all the cases in the jurisprudence. Of all the changes we can spot here, this is surely the most noteworthy: over the time under study, there was a




Figure 16: The subregion of the world that the claimant sought protection from changed over time. By 2023, claims made against countries in Southern Asia and West Africa accounted for the largest proportionate share of the annual jurisprudence.

major and substantive shift in the content of the Court’s refugee law jurisprudence. While it is important to approach correlations with caution, there is one that jumps out that may explain the shift. If we return to the data that tells us the country or persecution by

counting the number of times a country is mentioned in a decision, geographical patterns appear to present themselves. As Figure 16 suggests, the proportion of state protection cases appears to fall as the proportion with the number of cases considered by the Court that appeared to concern nationals from Eastern European countries (light red). While the increase in internal flight protection cases, which begins to rise from 2019, appears to track with an increasing proportion of the jurisprudence that focuses on cases from Western Africa and Southern Asia. This suggests that cases from different regions of the world engage different parts of Canadian refugee law jurisprudence and, as cases from those regions appear more or less before the Court, the jurisprudence changes shape to adapt. To learn refugee law doctrine, then, means to learn where people come from.

5.3.4 Algorithmic insights: more than just a sum of different parts?

After hours staring at the maps produced in this study, I noticed something. It is not a hard and fast observation, but it is one that generally holds, and generally describes the distribution of cases. Look back at the map with the clusters and remind yourself how these diagrams were produced. Each decision was converted into numbers and then an algorithm spatialized each decision by, in approximate terms, trying to locate decisions that were close together in multi-dimensional space, close together in two-dimensional space. This process was just as much art as it was science and the x and y axes referring not to some clear, human understandable metric, but an approximation of location in space. As the programmer, I can (and did) turn the dials a bit here and a bit there to produce a visualization that presented the data well. A nudge here and a nudge there, the data looks different.

But let us return to these maps (Figure 7). Did the algorithm stumble across some deep structure in the decisions when it produced the visualization? Yes, I think so. First, look at the shape of each decision dot: if it is a circle, the decision came to the Court from the appeal tribunal; if it is an “x,” it came from the first instance tribunal. On the top, there are more appeal cases, on the bottom there are more first instance cases. But more interesting, consider the distribution of cases from right to left. On the right, each cluster turns on a core evidentiary question: is the claimant credible? Are they telling the truth?

Is their claim fraudulent? These decisions are singular, focused, and individual. As we move left, the aperture widens a bit. Yes, cases still seem focused on questions of credibility, but the assessment of the evidence is more holistic and broader than the focused inquiries on the far right. Then, on the left side, the questions are different. No longer are decisions about individual claimants and their credibility, rather the questions are large: will this state protect this person? Are there zones of safety someone could reasonably relocate to? The y-axis: a proxy for the originating tribunal? The x-axis: the continuum from individual questions to large questions about societies and states?

Like any map, the way it is made determines the information it communicates. I am sensitive to the fact that I could easily have produced another map that may have represented different patterns, and perhaps not ones I could detect. Yet here lies the great promise and potential of reading from a distance. From the perspective of the individual, each day will feel unique and singular, but in aggregate and over time, our behaviour falls into repeating patterns that we can only truly appreciate when examined at scale. So too with our texts. Sometimes there are things we think we know but that we cannot see, but—even with the rudimentary tools deployed here—come into focus when seen differently.

Perhaps this is no great observation. Any refugee lawyer could tell you the same: some parts of the refugee law test focus on the individual, others focus on the state. That is hardly an original observation, but to see it confirmed by an algorithm that knows nothing about law but a little bit about patterns suggests that there may be many other regular and repeating behaviours in ourselves and our laws ready for us to see, once we begin to see in a new way.

5.4 Discussion

One function of large textual analyses is that they can sprawl. Despite my efforts to rationalize and organize millions of words of text, and despite my general conviction that the categorizations presented here are compelling, it is apparent that the wider and more general the trend, the more difficult it is to speak about. What might those large categories that deal with credibility tell us as people interested in refugee law? Beyond the fact that

they exist (something we already knew) probably not much. So, the following sections climb a ladder: at first, the simple and narrow observations invite some easier discussion points; at the end, some of the wider and more impressionistic findings invite more open-ended remarks.

5.4.1 An economic case for intermediate appellate review?

I begin with a finding that gestures at a conclusion that, if further study confirms, could be an important empirical fact both for refugee advocates and for people interested in good institutional design. Let us return to the findings that (i) within the jurisprudence the Federal Court of Canada's judicial grant rate is declining and (ii) that cases that come to the Court from the RAD have a lower grant rate than cases that come from the RPD. As I already raised, this latter finding might well be an indicator of overall system health: if the RAD is properly finding, correcting, and not introducing errors into refugee decision making, we should expect that cases that transit through the RAD to the Court will ultimately be treated differently than cases that were not examined by the RAD. Now, there are two major caveats here. First, we do not know whether the Court is granting leave to cases that originate from the RPD and the RAD at similar rates and, second, we do not know whether settlement patterns are the same for both RPD and RAD cases. Both questions call out for further study.

But let us assume for a moment that there is no meaningful difference between how cases from the RPD and the RAD are treated before they each reach merits adjudication before the Court. If this is the case, there may well be a compelling empirical, law and economics case to be made that the RAD has, not only protected claimants against wrongful deportation, but introduced major cost efficiencies into the refugee adjudication system. Here is the argument. Each procedural step in a refugee case costs: overhead, adjudicator salary, legal aid, government resources, and so on. A cost-effective system will therefore aim to minimize procedural steps while meeting its other objectives (in this case ensuring that people who require protection get it and that refugees are not wrongly deported). To err, therefore, is to cost: when adjudicators make mistakes and those

mistakes are identified by the Court, procedural steps get repeated because matters are returned to the IRB for reconsideration, incurring expense.

When we look at the headline figures here, there is a suggestion that the system is overall more efficient when there is an intermediate appellate review step because error correction happens earlier, obviating the need for further review. And, perhaps, if the costs of each step were known and more variables could be controlled for, we would find that a system with more levels of review somewhat ironically leads to fewer procedural steps. This matters because, at least as I understand the argument, the primary reason that some claimants do not have access to appellate review is to expedite processes to realize savings. Perhaps it is more expeditious and cheaper to give all claimants the same level of rigorous review.

5.4.2 Learning the law by knowing the jurisprudence’s shape

When we look too long at trend lines, when we look at proportions, or when we look at maps that collapse years of decisions into two dimensions, we can lose sight of how




Figure 17: On the left, orange dots show indicate cases decided in 2015; on the right, orange dots indicate cases decided in 2023

dramatic small changes can be. So, two more maps. Figure 17, on the right, colours all decisions decided in 2015 (the year I first started practicing refugee law) and, on the left, all decisions decided in 2023 (the year I started this project). In 2015, we can see decisions in the state protection region and a few in the “new evidence” zone. By 2023, the law looks different: there are many more cases in total, and they are now centred around internal flight protection and new evidence. Falun Gong cases are present at a greater rate

and cessation cases, too, look more prominent. Zones that I identified as engaging credibility questions are brighter.

This presentation of Canadian refugee law says little about doctrine and offers lawyers, critics, and policy makers few new tools that can be used to tweak or develop the law as it exists in principled and black letter form, but we can learn something about how the law ought to be taught. The Court's jurisprudence today is just different, topically and thematically, from what it was less than a decade before. Courses, seminars, trainings: they ought all adapt to the new reality, a reality that soon will surely change once again.

5.4.3 What accounts for jurisprudential change? And what do lawyers do?

Lawyers and judges often speak in terms of doctrine and principle, precedent and statute, logic and error. This is, to point to a discussion I will soon pick up, the *habitus* or common sense of the profession. I do not claim here that all insights developed with the assistance of computational methods will be of no practical use to traditional legal thinking, but I do wonder whether computational methods diminish the importance of doctrine and lawyerly thinking.

Before I clustered the cases I would not, for example, have been surprised to find that clusters organized around different judges (each judge, I might have reasoned, will write so distinctively that their own decisions will natural cluster together), or that we might see clusters emerge as important new caselaw invited lawyers and judges to use new legal forms, words, and expressions. Indeed, I was relatively confident going into this project that I would be able to map the rise and fall of clusters and spot correlations with grant rates: grant rates, I thought, would go up and up until—poof—an area of dispute disappeared, now that the first instance tribunal finally got the message. Other presentations will certainly present the data differently, but this presentation shows change, but not change owing to judicial or lawyerly genius and creativity. No, new innovative legal forms develop—or are at least prompted—by energy inputs external to the world of legal thinkers.

The French theorist, Pierre Bourdieu, helps us understand what we are seeing. Bourdieu said that the legal field (or the juridical field) ought to be thought of like any other social field: education, medicine, the arts, and so on. It is, Richard Terdiman explains, “organized around a body of internal protocols and assumptions, characteristic behaviors and self-sustaining values.”⁹⁴ Most of Bourdieu’s most famous studies—*Distinction*, *Home Academicus*—consider the *habitus*, or the unspoken field-specific common sense, of different parts of social life. In one of his few in-depth considerations of law, however, Bourdieu thought more specifically about the interface between legal culture and other fields, and their respective positions in society (the field of power).

For us, this is a useful inquiry because his study can help us navigate the two great poles of legal theory: instrumentalism and autonomy. Law, for Bourdieu, is not a simple function of the arrangement of power relations in society, nor is it a distinct and separate zone governed by formal and independent laws of legal reasoning. Instead, it (like all social fields) is a site of competition where, in this case, different actors compete for the “monopoly of the right to determine the law.”⁹⁵ But the quality of the contributions different actors make is circumscribed by each actor’s position and access to different sorts of capital.

Creativity: the development of new ideas and forms that give new shape to the law. Bourdieu teaches that legal professionals have a special sort of function, but one that they may not consciously identify with. If my conclusions here are correct, the development of the most consequential new legal forms—or at least the factors that accounted for the rise and fall of legal ideas—was external to the legal field. The emergence of a cluster focused on the adjudication of admissibility claims about evidence: a function of a political decision to make a new appeals process. The collapse of the state protection cluster: the drop in cases from Europe making it to final adjudication before the Court. The rise in internal flight alternative: an increase in claimants from South Asia and West Africa advancing all

⁹⁴ Richard Terdiman, “The force of law: toward a sociology of the juridical field: translator’s introduction” (1987) 38 *Hastings LJ* 805 at 806.

⁹⁵ Pierre Bourdieu, “The force of law: toward a sociology of the juridical field” (1987) 38 *Hastings LJ* 805 at 817.

the way to final judicial review. No, there were no major cases, no brilliant new ideas, and no innovations in principle that registered on these maps.

And this is no surprise: “[g]iven the determinant role it plays in social reproduction,” Bourdieu says, “the juridical field has a smaller degree of autonomy than other fields.” This means that “[e]xternal changes are more directly reflected in the juridical field, and internal conflicts within the field are more directly decided by external forces.”⁹⁶ Refugee law looks gradual, slow, and responsive: not an autonomous zone of consequential change brought about by battling intellects. But there is legal creativity here and it is seen precisely in that smoothness of those trends and the gradualness of those changes. From this perspective, the brilliance of law is not that it causes change, but that it metabolizes it.

Recall how this paper began. More people on the move, more experiences, more text, more ideas, more confrontations: more data speaking to the variety and complexity of lives in motion. And yet, by the time these cases move through the legal system and the Court, a major levelling has occurred. Here we see legal creativity at work, formalizing and systematizing wildly divergent lives:

[Juridical labour and the law] ties the present continuously to the past. It provides the guarantee that, in the absence of a revolution which would upset the very foundation of the juridical order, the future will resemble what has gone before, that necessary transformations and adaptations will be conceived and expressed in a language that conforms to the past... The universalization effect, which one could also term the normalization effect, functions to heighten the effect of social authority already exercised by the legitimate culture and by those who control it.⁹⁷

Is this what we see here? Yes, I think so. Step back and the numbers speak for a part of the world. Something has changed in the world and for all sorts of reasons there is more movement within it. And Canadian refugee law has adapted to that reality, it has incorporated a few new forms, but only a few. And perhaps this is the place to conclude.

⁹⁶ *Ibid* at 850.

⁹⁷ *Ibid* at 845.

What is the content of Canadian refugee law jurisprudence? A universalizing synthesis that adapts to a world out of balance.

6. Law School Goes Back to the Laboratory

This paper is one part about legal education, one part about artificial intelligence, and one part about the substantive law of employment insurance in Canada. Let me explain what I mean. My view is that law schools poorly equip students, particularly students interested in public law and poverty law, with the skills and theoretical frameworks needed to confidently and critically understand the law as it plays out on-the-ground. Put differently, we want for empirically rich, inductive, accounts of law that can be used for teaching.

At least, this was my experience. When I became a lawyer, I did what I always wanted to do: I started a poverty law practice and started representing individuals at their eviction hearings, their parole hearings, their solitary confinement reviews, their refugee hearings, their deportation hearings, and their immigration detention hearings. This was thrilling work—but it was thrilling work enabled by a network of peers and mentors, not one that my law school experience directly prepared me for. No, this is not the (by now) tired complaint that law school does not prepare students for practice in the “real world,” rather it is a recognition that public law and poverty law get short shrift in law school’s core curricula.

Yes, there are good reasons for this. No one argues that contract law doctrine, as taught in first year, really teaches students all that much about contracts and contractual relationships; but it does, at least in theory, give them a language and a toolkit to approach

contract law problems. Same, too, for tort law, constitutional law, criminal law, and property law: in most law schools the doctrinal presentation is self-aware and comfortable with the idea that its theoretical vision of the law does not have much purchase on the world and social life. Rather, it is a vector for other intangibles: learning to think like a lawyer, learning to deductively reason, learning to read case law, and so on.

I am not here to litigate whether the classical vision of legal education ought or ought not to remain dominant. Instead, I say that, if nothing else, we can do a better job of supplementing it. Mass legal experience: that law which is largely delineated by statute and regulation and that intervenes at high frequencies in people's lives, this category of law can be taught differently.

This is hardly a new position. Holmes always maintained that “[t]he life of the law has not been logic; it has been experience.”¹ The realists were puzzled that law schools did not teach students about the “actual behavior of judges, sheriffs, and litigants.”² Critical race theorists remind us that something inevitably gets lost when we deduct down from preestablished categories and fail to induct up from the complexity of experience.³ The crits taught us that doctrinal teaching lies to students by implying that some fields are based on principle (contracts law), while others are based on policy (environmental law), when, in fact, it is politics all the way down.⁴

While I do not suggest that law students have not heard and responded to these points (the proliferation of experiential learning and various law and society courses show that they have), I do suggest that a core divide remains in most schools' curricula: there is the core doctrinal program on one side; and matters of policy, specialization, and criticism

¹ Oliver Wendell Holmes, Jr & Steven Alan Childress, *The Annotated Common Law* (New Orleans: Quid Pro Law Books, 2010) at 1.

² Felix Cohen, “Transcendental Nonsense and the Functional Approach” (1935) 35:6 *Colm L Rev* 809 at 831.

³ Kimberlé Crenshaw, “Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics” (1989) *U Chicago Legal F* 139.

⁴ Duncan Kennedy, “Legal Education and the Reproduction of Hierarchy” (1982) 32:4 *Journal of Legal Education* 38.

on the other.⁵ This division's staying power can, at least in part, be understood in practical terms: it is just very hard to develop a theoretical model of mass legal experience, especially when it has so many local idiosyncrasies. One of the advantages of teaching traditional contract law, tort law, or criminal law is that there is a lot of secondary material to work with. The casebooks and textbooks are national (if not international) in scope, and draw on decades' worth of cases.

But teach a large and important public law field, such as the law of residential evictions, and each state and each province will have its own statute, its own set of regulations, its own changing policy frameworks, its own jurisprudence, its own political constraints, and so on. Here is the reality: a professor teaching residential evictions in Saskatoon just cannot teach the same course as their colleague in Halifax; nor can either of them teach the cases that were important a decade or two decades ago, because these fields of mass legal experience are so dynamic and subject to so much legislative and regulatory intervention.

Do you think much about family law, first-instance criminal law, mental health law, or prison law? If so, I bet that you are nodding. The challenge of understanding and developing theoretical accounts of these fields is complicated by the fact of scale. How many eviction orders were issued in Ontario yesterday? Hundreds. How many of these will be appealed? Probably none. In administrative and poverty law fields we are swimming in first-instance law and have little appellate law to guide us—nor much reason to believe that the appellate law that we do have can tell us much about on-the-ground experience. So long as we teach law based on the idea that some decisions are worth studying because they are influential, we will poorly teach those areas where there are, quite frankly, fewer influential cases but regular and repeating patterns of thought to be found.⁶

⁵ For a discussion on this feature of law school curriculums see Linda Edwards, "The Trouble with Categories: What Theory can Teach Us about the Doctrine-Skills Divide" (2014) 64 J Legal Educ 181.

⁶ For a discussion on how different methods can produce different pictures of law, at least compared to traditional studies of influential cases, see Sean Rehaag, "Claim Types in Canada's Refugee Determination System: An Empirical Snapshot (2013–2021)" (2024) 40:1 Refuge 1.

This leaves teachers interested in mass legal experience in a difficult spot. They know that students want to learn about these fields, they know that these fields may benefit most from a focused and rigorous law school curriculum, they know that students want to develop critical analyses of the local shape of these fields, but they cannot ask students to read ten thousand cases, nor can they just teach a statute and set of regulations from the top to the bottom. What is needed is a methodology that allows for students and their teachers to rapidly develop theoretical accounts of those mass legal experiences happening in and around the jurisdictions that they study.

Hear me out, then, when I suggest that some of the most consequential impacts of new artificial intelligence technologies for legal study may be felt in public, administrative, and poverty law fields. Inductive reasoning: the process of building up a theoretical model from evidence. With its ability to develop new representations of data, artificial intelligence can help scholars, teachers, and students navigate those huge masses of data collections that collectively form the jurisprudence of mass legal experience to build robust, but locally relevant, theoretical accounts of law.

Does this make me sound like an extremist or a fabulist? I hope that it makes me sound like an empiricist and, in some strange way, a traditionalist. I am no friend of the casebook method, but I respect Christopher Langdell's contributions to legal education. Like us today, Langdell conceived of the casebook method partly as an answer to the problem of scale: this new book would help students navigate the "great and rapidly increasing number of reported cases in every department of law."⁷ But more than that, Langdell was—as I am—a believer in the law library. In a valedictory address, he explained that the law library "is the proper workshop of professors and students alike; that it is to us all that the laboratories of the university are to the chemists and physicists, the museum of natural history to the zoologists, the botanical garden to the botanists."⁸

Collect and read the law: then build an account of it. Today, our libraries and databases are by many orders of magnitude larger than they were in Langdell's day and,

⁷ Christopher Columbus Langdell, *A Selection of Cases on the Law of Contracts* (Boston: Little, Brown and Company, 1871) at vi.

⁸ Christopher Columbus Langdell, "Harvard Celebration Speech" (1887) 3:1 L Q Rev 123 at 124.

more importantly, our understanding of what counts as law (and what ought to interest legal scholars and law students) is more expansive. For these reasons and more, Langdell’s methods look outdated, but there is an original impulse to work with here.

So, because it is better to show than it is to tell, let us work together and see how modern technology can help us develop new theoretical descriptions of law, how it can enable new types of pedagogy, and how it can expand our knowledge of mass legal experience. My proposal is simple. Employment insurance is one of Canada’s most important and largest social programs. It provides some income support to some workers some of the time. When the fund administrator makes a decision that a claimant disagrees with, they can initiate an appeal before the Social Security Tribunal of Canada (the “SST”). What are these cases about? In what follows, I show how we can take thousands of legal decisions—many more than any one person could ever hope to read—and use artificial intelligence to represent these as numbers before algorithmically sorting them into clusters. By the end of this paper, we will have a new theoretical understanding of one aspect of law as it relates to mass unemployment, a comprehensive assessment of the types of disputes brought to the tribunal, and a glimpse of how these methods might inform new, intellectually rigorous, but still familiar, pedagogies.

6.1 Employment Insurance (“EI”) and the Social Security Tribunal of Canada (“SST”)

6.1.1 Origins of the EI program

Employment Insurance (“EI”) is one of Canada’s largest social support programs. Millions of people in Canada rely on the program each year to provide income support when their earnings are interrupted. In 2023, for example, the EI program received

3,069,140 initial applications for benefits, and granted 2,830,600 of these.⁹ That same year, the program paid out 20,151,572,390 dollars.¹⁰

Not only is the EI program large, it is also old. The first EI program was established during the depression in 1935, by the Conservative government led by R.B. Bennet. After the next government, led by Mackenzie King, referred a question regarding the program's constitutionality to the Supreme Court of Canada, the Court found that the insurance program was *ultra vires* Parliament's jurisdiction—a conclusion subsequently upheld by the Judicial Committee of the Privy Council.¹¹ In 1940, the government and provinces agreed to a constitutional amendment specifically granting the federal government jurisdiction over “unemployment insurance,”¹² following which Parliament passed new legislation enabling an insurance program for the unemployed.¹³

The Second World War, Alvin Finkle argues, provided both the “rationale and a political window” for implementing a constitutionally durable unemployment insurance program. The fact of low unemployment because of the war effort meant that funds for the program could be quickly raised, while ensuring that payouts would be kept to a minimum. At the same time, the fact of the fund, coupled with its authority to issue loans, would help the government finance the war. By July 1940, after reaching an accord with the provinces, King convinced the imperial Parliament to amend the British North America Act to give explicit constitutional authority to the federal government for unemployment insurance.¹⁴

⁹ Statistics Canada, Table 14-10-0005-01 Employment insurance claims received by province and territory, monthly, seasonally adjusted (Ottawa: Statistics Canada), online: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410000501>.

¹⁰ Statistics Canada, *Table 14-10-0007-01 Employment insurance benefit characteristics by class of worker, monthly, unadjusted for seasonality* (Ottawa: Statistics Canada), online: <https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410000701>.

¹¹ *Attorney General of Canada v Attorney General of Ontario et al*, [1937] UKPC 7 (28 January 1937).

¹² *The Constitution Act*, 1867, 30 & 31 Vict, c 3, s 91(2A).

¹³ See “The UI Act of 1940” in Georges Campeau, *From UI to EI: Waging a War on the Welfare State* (Vancouver: UBC Press, 2004).

¹⁴ Alvin Finkle, *Social Policy and Practice in Canada: a history* (Waterloo: Wilfred Laurier Press, 2006) at 116.

This first program was, at least by current standards, limited. Coverage only extended to individuals who worked for a wage (“Only those who work for wages or a salary, who are at liberty to quit their jobs yet may also be deprived of them by someone else, may be unemployed”)¹⁵ but workers in large swaths of the economy were exempted from coverage, with the program really only protecting male industrial workers. At the moment of the program’s origins, only 42% of the workforce was covered by the program.¹⁶ For Georges Campeau, the story of the insurance program since its origin is double-edged. On the one hand, the program steadily expanded decade after decade to include more workers, and to expand the range of circumstances in which workers were entitled to benefits. If the original EI program applied only to industrial workers who were laid off, the modern program is significantly more expansive. Now seasonal workers, self-employed fishers, and some self-employed professionals are all eligible for coverage.

Likewise, the program is no longer narrowly focused on workers who are laid off, but it also covers workers who need to stop working because of illness, to care for children, or to care for ill family members. Most of these expansions in eligibility came in the 1970s when, Campeau argues, Keynesian politics dominated, and there was an understanding that unemployment was a social phenomenon that required a social response.¹⁷ But on the other hand, Campeau says that the austerity of the 1980s and 1990s, cut at the program’s overall generosity. While the categories of people covered by the program grew, the overall eligibility criteria were tightened and the bases upon which a worker could lose coverage increased.

6.1.2 The establishment of the SST

In each of these stories, adjudication matters. When the insurance program was established, a Board of Referees was also created to hear disputes between the insured and the insurance plan (and, in some rare cases, between the plan and the employer). Later,

¹⁵ Quoted in Georges Campeau, *From UI to EI: Waging a War on the Welfare State* (Vancouver: UBC Press, 2004) at 62.

¹⁶ *Ibid.*

¹⁷ *Ibid* at vii.

additional appeal mechanisms were also included. Regardless, Campeau concludes that during the 1970s and the early 1980s, the decisions of the Board tended to “favour the jobless.”¹⁸ By the 1980s, however, Mulroney Progressive Conservatives made a series of small changes that restricted the jurisdiction of the Board. It lost its jurisdiction to review some classes of decisions made by the EI program (for example, the program can require a worker to take a course to upgrade their skills and, where once workers who disagreed with their placements could appeal to the Board, they lost this right in the 1980s).¹⁹ But more importantly, adjudication itself became more significant when accessing the program itself became more difficult: as the qualifying criteria for eligibility were tightened, Campeau points out that the number of appeals brought to the board increased by almost 100%.²⁰

In 2012, adjudication for EI appeals was substantially revamped. By 2010, the Board was a large organization, staffed by over 1,000 part-time members.²¹ To the government of the day, the tribunal looked bloated, inefficient, and “prone to [making] local, idiosyncratic decisions.”²² Legislation was passed to establish the SST, which was given jurisdiction to consider appeals for EI matters, as well as other social security programs, such as Old Age Security and Canada Pension Plan disability benefits.²³ Unlike the Board of Referees that preceded it, the SST is staffed by a much smaller complement of full-time professional adjudicators. Reaction to the new tribunal was swift and, especially from labour groups, negative. As one labour official explained, the new model “completely confiscated the process and squeezed out workers and business, despite the fact that they pay for the system... the people now involved [in making judgements] are only accountable to meeting the performance requirements of the Social Security

¹⁸ *Ibid* at 108.

¹⁹ *Ibid* at 136.

²⁰ *Ibid* at 166.

²¹ André Léonard, *The Employment Insurance Program in Canada: How It Works*, Publication number: 2010-52-E (Ottawa: Library of Parliament, 2010) at 13.

²² “What the New EI Rules Mean” (2012), Mowat Centre for Policy Innovation at 10, online: <https://mowatcentre.munkschool.utoronto.ca/wp-content/uploads/publications/52_what_the_new_ei_rules_mean.pdf>.

²³ *Department of Employment and Social Development Act*, SC 2005, c 34, Part V.

Tribunal; previously the Board of Referees were accountable to the union and business groups that nominated them.”²⁴

The new tribunal was equipped with a broad jurisdiction to decide all disputes about law. This meant that a person could appeal any decision made by the EI authorities to the tribunal if they thought the EI authorities erred in law.²⁵ Importantly, the tribunal was not equipped with any equitable jurisdiction, meaning that it had no authority to waive any rule or require the fund administrators to refrain from taking an action that, though potentially unfair, was legally prescribed.

Shortly after SST began to process cases in April 2013, these criticisms sharpened. By 2017, the government understood that there was a general level of discontent regarding the new tribunal and commissioned the accounting firm KPMG to conduct an independent audit of the tribunal. The audit found two major systemic problems: first, EI appeals were taking up to five times as long to process in the new tribunal compared to processing times before the Board of Referees; second, the new professional tribunal was too wedded to a self-conception that its mandate was to strictly apply the law, leading to confusing and unnecessarily complicated processes for workers to follow while advancing their cases.²⁶

Though the SST undertook to incorporate the recommendations made by the auditors,²⁷ within 18 months the government announced that the tribunal process would be reorganized again. In place of the SST’s general division, a new trilateral tribunal would

²⁴ Quoted in Donna Wood, *The Seventy Five Year Decline: How Government Expropriated Employment Insurance from Canadian Workers and Employers and Why This Matters*, Mowat research publication 151 (Toronto: Mowat Centre for Policy Innovation, 2017) at 22, online: < https://tspace.library.utoronto.ca/bitstream/1807/99369/1/Wood_2017_The_Seventy_Five.pdf>.

²⁵ *Department of Employment and Social Development Act*, SC 2005, c 34, s 58.

²⁶ Employment and Social Development Canada, *Review of the Social Security Tribunal of Canada for Employment and Social Development Canada – October 2017* (Ottawa: ESDC, 2017), online: < <https://www.canada.ca/en/employment-social-development/corporate/reports/evaluations/social-security-tribunal-review.html>>.

²⁷ Social Security Tribunal of Canada, *Publication of KPMG’s Review of the Social Security Tribunal* (January 2018), online: < <https://sst-tss.gc.ca/en/whats-new/publication-kpmgs-review-social-security-tribunal>>.

be created, modeled after the Board of Referees, while the SST's appeal tribunal would remain to hear appeals from the new Board.²⁸ In 2023, Parliament included legislation in a budget implementation bill officially describing this system although, at the time of writing, the new tribunal had not yet been stood up.

6.1.3 Little is known about the substance of the EI disputes brought to the SST

Here is how the program works at a high level: workers and employers make contributions into a fund. A government department (referred to here as the “fund administrator”) uses that fund to provide temporary financial assistance to workers when they lose their earnings through no fault of their own. Sometimes this may be because they need to care for a child or a sick family member, because they were laid off, because they work in a specified seasonal industry, or because they are in quarantine. The amount a worker is entitled to can depend on many factors, just as a worker's eligibility for benefits depends on many factors.

Inevitably, disputes arise between workers and the fund administrator. When a claimant for benefits believes that the fund administrator has made a legal error with respect to their application, they may bring an appeal to the Social Security Tribunal of Canada (the “SST”). The SST is composed of four divisions, of which two handle EI matters: the general division hears first level appeals, and the appeals division hears, with leave (or, in lay terms, permission to appeal), appeals from the general division.

What is the case law of the SST, particularly its EI general division, about? The short answer is that we do not know, or at least that we do not know much. To my

²⁸ Employment and Social Development Canada, “Helping middle-class Canadians with the support they need: Government of Canada announces new reforms to the Social Security Tribunal” (15 August 2019), online: <<https://www.canada.ca/en/employment-social-development/news/2019/08/helping-middle-class-canadians-with-the-support-they-need.html>>.

knowledge, while there are a handful of studies from the 1970s,²⁹ 1980s,³⁰ and 2000s³¹ that analyze aspects of a predecessor tribunal's work, there are no scholarly studies that consider the SST's decisions. The SST itself only publishes decisions and does not comprehensively organize or summarize its work product. The fund administrator, Employment and Social Development Canada, publishes an excellent twenty-six chapter manual on employment insurance law principles, but this text is not an empirical account of the SST's work so much as it is the government's position on what the law is.³² The *EI Act*³³ itself and its associated regulations are convoluted and, on their own, tell us little about the sorts of disputes about the law claimants bring to the tribunal.

6.2 Designing a computational study to understand the SST's EI jurisprudence

To understand the shape of the SST's jurisprudence, to understand the quality of the disputes claimants bring to the tribunal, I designed a qualitative computational study. The object of this study is to cluster similar decisions of the tribunal into similar categories, and then isolate the decisions that are most emblematic of that cluster. At the end of the study, the goal is to map the SST's jurisprudence and show the different categories of disputes and better understand what those disputes are about. Put differently, the goal is to be able to explain—in large macro terms—what the SST does regarding employment insurance law on a regular and recurring basis.

²⁹ Daniel Potter “Unemployment Insurance: Policies and Principles of Disqualification and Disentitlement for Benefits” (1976) 3 Dalhousie LJ 178.

³⁰ Reuben Hasson “The Cruel War: Social Security Abuse in Canada” (1981) 3 Can Tax'n 3 114 and Reuben Hasson, “Discipline and Punishment in the Law of Unemployment Insurance--A Critical View of Disqualifications and Disentitlement” (1987) 25 Osgoode Hall LJ 615.

³¹ Gaile McGregor, “Anti-claimant bias in the employment insurance appeals system: Causes, consequences, and public law remedies” (2001) 15 Canadian Journal of Administrative Law & Practice 229.

³² Employment and Social Development Canada, *Digest of Benefit Entitlement Principles* (Ottawa), online: <https://www.canada.ca/en/employment-social-development/programs/ei/ei-list/reports/digest.html>.

³³ *Employment Insurance Act*, SC 1996, c 23.

My methodology has five parts which I describe in the following sections: (1) data acquisition, (2) data transformation, (3) clustering, (4) validation, and (5) qualitative exploration and theory building.

6.2.1 Data acquisition

The SST publishes a selection of its cases on its website. It allows individuals to engage with those decisions permissively, explaining that “you may reproduce the materials [on the website] in whole or in part for non-commercial purposes, and in any format, without charge or further permission.”³⁴ As a federal tribunal, the SST is also covered by the *Reproduction of Federal Law Order* which allows anyone to “reproduce enactments and consolidations of enactments of the Government of Canada, and decisions and reasons for decisions of federally-constituted courts and administrative tribunals.”³⁵ This means that a researcher can freely access the tribunal’s decisions.

To support free access to law, Professor Sean Rehaag’s *Refugee Law Lab* has scraped—that is, computationally collected—all decisions on the SST’s website and made them freely available to researchers.³⁶ I downloaded this dataset in January 2024. At that time, the dataset contained 26,550 decisions published between March 8, 2013, and November 3, 2023.

After some examination, I determined that the dataset contains English and French versions of almost every decision. I decided to work with only the English versions of each decision, leaving me with 13,284 decisions. As discussed above, the SST has jurisdiction for more than just EI matters, but the dataset does not explicitly distinguish between different categories of SST cases. This said, within the decision texts several data types were consistently presented. For example, the type of case, the division of the SST, the final decision, and the identity of the adjudicator always appeared in the same part of

³⁴ Social Security Tribunal of Canada, “Terms and Conditions” (visited 29 July 2024), online: <<https://sst-tss.gc.ca/en/terms-and-conditions>>.

³⁵ SI/97-5.

³⁶ “SST Bulk Decisions Dataset,” Refugee Law Lab, online: <<https://refugeelab.ca/bulk-data/sst/>>.

the decision. I developed some simple code in the Python computer programming language to extract this information.




Figure 1: Number of decisions in the dataset (left) and the number of words issued by the tribunal by year (right)

With this information, I was able to isolate for all general division final decisions. This left me with 3,336 decisions in the dataset. Importantly, this dataset is partial. The SST does not publish every decision it renders. Instead, it only publishes decisions that (1) it deems representative of its body of work or (2) cases that were appealed.³⁷ This means that the subsequent findings do not necessarily speak to the totality of the tribunal’s work, but to the totality of its published first-instance jurisprudence. Put differently, it is a study of every available decision, not every decision ever rendered. Nonetheless, the dataset is large, containing over 9.6 million words. As Figure 1 shows, the SST began to publish more decisions and more words in 2019, suggesting that the dataset is somewhat chronologically skewed. Without knowing the total number of cases the SST adjudicated each year, it is not possible to assess the extent to which this trend is representative of some increase in the tribunal’s output or a function of some change in its publication protocols.

³⁷ “What decisions the Social Security Tribunal publishes,” Social Security Tribunal, online: <<https://www.sst-tss.gc.ca/en/decisions-laws-rules-and-policies/what-decisions-social-security-tribunal-publishes>>

Interestingly, I found that the mean length of decisions changed substantively, from a high of just over 5,000 words in 2017 to a low of just under 2,500 words in 2019. It appears that this change was part of an intentional effort by the tribunal. In a 2020 report, the tribunal explained that stakeholders told it that its decisions “were long, legalistic and difficult to understand” and that it had implemented a training program to help adjudicators “consistently produce more readable decisions that do not demand from the reader a higher education or knowledge of the law.”³⁸ Finally, most decisions in the dataset concerned appeals brought by claimants that were ultimately dismissed by the tribunal. As Figure 2 shows, the grant rate hovered around 30% across most years in the dataset, before steeply dropping in 2021.




Figure 2: The proportion of cases in the dataset for which the appeal was allowed or dismissed.

6.2.2 Data transformation: using start-of-the-art artificial intelligence models

Computers do not understand human text. What computers can do, and excel at, is numerical calculations. The basic intuition underlying most (perhaps all) computational

³⁸ Social Security Tribunal of Canada “An evaluation of how easy it is to read decisions of the Social Security Tribunal” (10 January 2024), online: < <https://www.sst-tss.gc.ca/en/our-work-our-people/evaluation-easy-it-read-decisions-social-security-tribunal>>.

analyses of human text is that at least some textual meaning can be re-represented as a numerical pattern. This means that texts must be transformed into numbers before computational methods can be applied to the data to uncover patterns of interest.

A traditional approach to data transformation illustrates what I mean. The term frequency-inverse document frequency algorithm (tf-idf) was developed by Karen Ida Boalsh Spärck Jones in her 1960s doctoral dissertation. So far ahead of its time, the algorithm was not published in an article by her until twenty years later.³⁹ The basic concept is simple and elegant. To understand mathematically what a document is about, we might seek to identify words in that document that make it unique. For example, if we have twenty documents that are recipes for a tomato sauce, and twenty that are recipes for a pesto, we might distinguish the two sets of documents by their relative use of words like “basil” and “tomato.”

The tf-idf algorithm mathematically represents this phenomenon by counting the frequency of each word within a document and dividing it by the inverse frequency of that word across all documents within the cluster. The effect is to yield, for each word, higher values for words that appear relatively more frequently within one document as compared to the corpus as a whole. Once applied to a set of documents, each document has a numerical, almost spreadsheet, representation where each individual word used in the corpus is a column, each row is a document, and each value is the word’s tf-idf score. Once so mathematically represented, various techniques can be applied to make inferences about the document’s thematic meaning⁴⁰ or compare the similarity between different documents.⁴¹

The numerical values of each document are called embeddings. One objective of computer scientists is to produce embeddings that capture more of the meaning within a

³⁹ Yorick Wilks, “Jones, Karen Ida Boalsh Spärck (1935–2007)” in Oxford Dictionary of National Biography, (06 January 2011, updated 25 September 2014).

⁴⁰ See for an example Juan Ramos, “Using tf-idf to determine word relevance in document queries” (2003) 242:1 Proceedings of the first instructional conference on machine learning 29.

⁴¹ Prafulla Bafna, Dhanya Pramod & Anagha Vaidya, “Document clustering: TF-IDF approach” (2016) 2016 International Conference on Electrical, Electronics, and Optimization Techniques 61.

document (or, for that matter, any type of data). Today, new classes of techniques to generate document embeddings show extraordinary promise. In particular, these new classes of techniques promise to address two major shortcomings of the tf-idf algorithm and algorithms like it. First, these older algorithms are referred to as “bag of words” algorithms pay no attention to the order in which words are used within a document. Second, these older algorithms have no method of taking into account the intrinsic meaning of different words: they simply count words and do not necessarily notice that documents that use the word “dalmatian” and “rottweiler” might both be discussing dogs.

If the tf-idf algorithm and its peers were part of the first generation of mechanisms used to generate document embeddings, some scientists argue that a fourth generation of technique is just now emerging to the fore.⁴² This fourth-generation approach, marked by the rise of “universal text embeddings,” aims to represent so much meaning in text in numerical form that they can “mimic the fundamental process of how humans understand and process text.”⁴³

The key technological distinguishing feature of the first and fourth generation techniques is that this most recent class of techniques depends on the use of a large language model to generate the embeddings. These artificial intelligence models “leverage deep learning to represent the complicated associative relationships between words as they are used in the training dataset of text-based content.”⁴⁴ Put differently, these models are shown many billions of words and slowly begin to develop a mathematical formula that accounts for what those words may mean, in light of how they are used in combination. Once such a model has been trained, it can be used to generate embeddings, or mathematical representations, of a document’s meaning.

⁴² Hongliu Cao, “Recent advances in text embedding: A Comprehensive Review of Top-Performing Methods on the MTEB Benchmark” (2024) arXiv:2406.01607, online: <<https://arxiv.org/abs/2406.01607>>.

⁴³ *Ibid* at 3.

⁴⁴ Arun James Thirunavukarasu et al, “Large language models in medicine” (2023) 29 Nature Medicine 1930 at 1930.

As I write, there has been a major development in this field. Up until several months ago, models that were designed to produce document embeddings were quite limited, perhaps capable of generating embeddings for documents that were, at most, 400 or 500 words long. In the spring and summer of 2024, however, several artificial intelligence firms released models that could generate embeddings for much longer documents. For example, Nomic AI released a model that could generate embeddings for documents up to 8024 tokens (approximately 6,000 or 7,000 words) long.⁴⁵ The novelty of this class of model and this technique, coupled with the length of the standard publication cycle, means that I am unaware of any social scientific projects that have yet exploited this technology. Computer scientists, too, are only just starting to explore how these long, universal text embeddings might be used, but report that they “excel at capturing subtleties in structured language.”⁴⁶

I was excited to see, and perhaps be the first to report, on how these embeddings may or may not perform well on legal text. Because of my commitments to open-source research and science, I decided to use models that I could run locally. To generate the embeddings, I used the nomic-embed-text-v1.5 model, referred to above, and ran it using an open-source tool that runs models on local computers called Ollama.⁴⁷ After running the code I developed, I was left with a matrix that had 3,336 rows (one for each document) and 768 columns. The values in each row were the universal text embedding for each document.

6.2.3 Clustering the data

The object of this study is to identify clusters within the data to spot themes in the dataset. Or, to put it differently, to invite the computer to find areas where different decisions clustered together. There are many different algorithms that can be used to identify

⁴⁵ Zach Nussbaum et al, “Nomic Embed: Training a Reproducible Long Context Text Embedder” (2024) arXiv, 2402.01613.

⁴⁶ At the time of writing, one of the few articles (a preprint) on the subject is Alina Petukhova, Joao P. Matos-Carvalho & Nuno Fachada, “Text Clustering with LLM Embeddings” (2024) arXiv, 2403.15112.

⁴⁷ Ollama, online: <<https://github.com/ollama/ollama>>.

clusters but, in general terms, they each work the same way. If a data point can be described numerically, in this case if a document can be represented with an embedding, that means it can be located within space. Imagine a grid with three locations on it: an “x”, an “o”, and a “b”. Because the grid is in two dimensions, we can describe each spot numerically: “x” is at location 2 across and 1 up (2,1), “o” is 5 across and 1 down (5,-1), and “b” is two to the left and four up (-2,4). Once we have these numerical locations, we can calculate the distance between the data points, figure out which are closest, and which are the furthest. The point of a clustering algorithm is to sort through the location of all the data points and see if there are clusters or communities of data points located in similar spots.

While it is easy to imagine this sort of effort in two dimensions, the documents under consideration here are considerably more complex than two dimensional maps. Each document is described by 768 dimensions (each of those columns produced when I generated the embedding), which is difficult for us to conceive of. Nonetheless, mathematically the relative locations of each data point to each other can still be calculated, allowing for clusters to be uncovered.

For this study, I decided to use the Hierarchical Agglomerative Clustering algorithm to assign the documents to different clusters. I selected this algorithm because there is a readily available and respected open-source implementation through the scikit-learn project⁴⁸ and the fact that, based on my own experimentation, this algorithm works well for text clustering tasks. I also note that the computer science literature shows that no one clustering algorithm is superior to all others for text clustering.⁴⁹ This algorithm works essentially by merging the two closest data points into a cluster and then iteratively repeating that process until all data points and clusters are finally merged into one large cluster.⁵⁰ The algorithm does not know, however, when to stop clustering, so the

⁴⁸ Fabian Pedregosa et al, “Scikit-learn: Machine Learning in Python” (2011) 12:85 Journal of Machine Learning Research 2825.

⁴⁹ In the one article I am aware of the discusses artificial intelligence text embeddings for clustering, the authors worked with several different clustering algorithms and found that no one clustering algorithm was a consistent top performer: Alina Petukhova, Joao P. Matos-Carvalho & Nuno Fachada, “Text Clustering with LLM Embeddings” (2024) arXiv, 2403.15112.

⁵⁰ “2.3.6 Hierarchical clustering” *User Guide: scikit-learn* (accessed 30 July 2024), online: <<https://scikit-learn.org/stable/modules/>>.

researcher's job is to identify the number of clusters they wish to review, and to instruct the algorithm to stop clustering once it reaches that point.

6.2.4 Quantitative validation

Perhaps the most important job, or at least a critical job, of a social science researcher using computational methods is to validate the returns of the algorithm. Computers are superior pattern detection machines and if a computer detects a pattern there is no guarantee that it is of interest to the researcher. In this study, I used a mix of quantitative and qualitative measures to validate my findings and help guide my research.

The quantitative tool I used to validate the quality of the clustering was the silhouette coefficient. The silhouette coefficient is a simple statistical measure used to identify the quality of a cluster assignment. To calculate the coefficient for an individual data point one takes (a) the average distance between the data point and all data points in the next closest cluster and from that subtracts (b) the average distance between the data point and all other data points within the cluster and then divides that sum by the larger of a or b. Higher scores indicate higher quality assignments, while lower scores (and certainly scores below zero) indicate a lower quality assignment. Again, I used the scikit-learn implementation to calculate this coefficient.⁵¹

I used the silhouette score in three different ways. First, I used it to help me identify the optimum number of clusters to search for. To do this, I directed the algorithm to make cluster assignments for 2 clusters, and then iteratively all the way to 40 clusters. Each time, I calculated the average silhouette coefficient for all the data points (this is known as the silhouette score). Once this was completed, I found that the highest silhouette score was returned when the algorithm assigned data points to 18 different clusters. I therefore reasoned that 18 clusters, at least by this measure, returned the algorithmically optimal arrangement of the data.

This is not to say, however, that all 18 clusters returned were valuable. The second way I used the silhouette score helped me identify those clusters that were too diffuse to

⁵¹ For a discussion of this tool and its use see “2.3.11.5. Silhouette Coefficient,” *Ibid.*

be of interest. After examining the 18 clusters identified by the algorithm, I calculated the silhouette scores for each cluster. I found that 3 clusters had negative average silhouette scores. This led me to conclude that these clusters were poorly defined or, to put it simply, were so spread out that most individual data points ended up closer to other clusters than to their in-cluster peers. I excluded these ill-defined clusters from my analysis.

I was curious to see whether the universal text embeddings produced algorithmically superior clusters when compared to tf-idf embeddings. This metric suggested that they did. When I ran the same experiment with tf-idf embeddings, the optimal clustering returned 12 clusters with positive mean silhouette scores.

Finally, some data points will necessarily be better representations of the essential quality of a cluster than others. This is because some cases will fall at the edge of the boundary line and might contain some topics and ideas that are part of another cluster. If I focused my analysis on decisions that were all at the periphery of the cluster, I feared that I could be misled as to the essentially unique quality of that cluster. I therefore sorted each case by the value of their silhouette score so, when it came time to qualitatively inspect each cluster, I knew which data points were algorithmically the most central within each cluster.

6.2.5 Qualitative discovery and theory building

While quantitative tools are useful for validation, for qualitative studies the most effective validation tools are human. Just because a computer has identified clusters of documents does not mean that those clusters are responsive or useful to the research question. One feature of this sort of research project—where the goal is to build a theoretical representation from a dataset—is that it is inductive.⁵² Given the richness and variety of textual data, it can ground studies that suggest “a new way of looking at the world.”⁵³ Instead of approaching the data with a preconceived notion of its arrangement, the goal of

⁵² Justin Grimmer, Margaret Roberts and Brandon Stewart, *Text as Data: A New Framework for Machine Learning and the Social Sciences* (Princeton: Princeton University Press, 2022) at 14-15.

⁵³ *Ibid* at 15.

this sort of study is to recursively, iteratively, and inductively build theoretical representations.

Put differently, this study worked on the methodological theory that a computer could identify generalizable patterns in the data and once those patterns were qualitatively investigated, new theoretical descriptions of law could be developed. This type of study departs from the scholarly norm (where “the standard practice of writing articles begins by stating a theory, its observable implications, and the measurement strategy; and then the dataset is introduced”)⁵⁴ by beginning with data and working up to a theoretical vision.

I used four qualitative strategies to help me understand and describe the clustering. The first strategy was the most important: close reading. As discussed above, I identified the decisions with the highest silhouette scores. Once I had these scores, I read the five decisions from each cluster with the highest silhouette score to develop my own subjective assessment of the subject. This was the most enlightening but labour-intensive aspect of this project. If each decision is 3,000 words long, this meant that I read 15,000 words per cluster, and 225,000 words across all fifteen clusters.

To help guide my reading, I (and this is the second tool) occasionally used a large language model to summarize the decisions I was reading. The third tool I used was word clouds. Earlier I discussed the tf-idf algorithm. I realized that this algorithm could usefully assist me in this project. I used this algorithm to generate a new set of embeddings for each document. Then for each cluster I took the mean tf-idf embedding and, using a Python library,⁵⁵ created word clouds for each cluster where words and expressions with higher mean tf-idf scores were larger.

Finally, I created a visualization, or a map, of all the data points. To make this visualization, I used a dimensionality reduction method⁵⁶ to project the 768 dimensions in each document to two dimensions. Inevitably when 768 dimensions are reduced to two,

⁵⁴ *Ibid* at 14.

⁵⁵ Andreas Mueller, “wordcloud,” online: <https://github.com/amueller/word_cloud>.

⁵⁶ For this task, I used the scikit-learn implementation of the t-sne algorithm, which is described here: “2.2.2.9. t-distributed Stochastic Neighbor Embedding (t-SNE)” *User Guide: scikit-learn* (accessed 30 July 2024), online: <<https://scikit-learn.org/stable/modules/>>.

quite a bit of information gets lost, but computer scientists and mathematicians have developed approaches that attempt to preserve as much of the structure within the data, to give some sense of how data points relate to each other. For present purposes, I consider the resulting map to be a qualitative aid largely because I am aware from experience that the visualization can look quite different if some of the algorithm's hyperparameters are adjusted, making it more of a subjective aid than a robust quantitative description of the data.

As we shall see, these methods allowed me to see that the embeddings generated by a modern artificial intelligence model, and then clustered with an agglomerative clustering algorithm, were—while not without problems—effective for supporting a study that set out to develop an empirically grounded theoretical representation of employment insurance law adjudication in Canada.

6.3 Report on the clustering

My findings in this study fall into two classes. First, I describe some overall impressions of the clustering and point to some limitations of the computational method employed here. Second, I report on my qualitative assessment of the clusters.

6.3.1 LLM text embeddings produce useful clusters for analysis, but can hallucinate relationships between decisions

In Figure 3, I present two maps. On the left, we see each decision, with a colour corresponding to its cluster assignment. On the right, I have adjusted the vividness of each data point, to correspond to the data points' silhouette coefficient. Data points with negative silhouette coefficients or silhouetted coefficients of zero have been removed, while data points with low silhouette scores are faded.



Figure 3: Two maps of the dataset. On the left, each data point's colour corresponds to its cluster assignment. On the right, the vividness of each data point corresponds to its silhouette coefficient.

Recall that part of my interest in working with the silhouette coefficients was to identify clusters that were not well defined and that, when I directed the algorithm to search for 18 clusters, 3 of the returned clusters had negative mean silhouette scores. Here we can clearly see these clusters: the dark orange, the light blue, and the dark clusters almost entirely disappear when the map accounts for each data points' silhouette score. A final cluster, the light orange cluster, has so few members that it barely even registers on the visualization.

The visualization also points to at least one area for concern. Unlike every other cluster in the dataset, the dark gray cluster gets visualized in two different locations, right next to both the light purple and dark blue clusters. This is notable because, in just about every other case, the visualization algorithm and the clustering algorithm present similar

conclusions, with data points from the same cluster finding themselves in similar regions of the map. This is a matter to which we will shortly return.

When we look at the map, we can also see more clearly why it might be advantageous to for qualitative study, focus on those decisions from each cluster with the highest silhouette scores. Examining the dark blue and dark brown clusters in the map on the left, we can see that they appear to be quite closely related, butting right up against each other. When each data point's silhouette score is accounted for on the right, we see that data points closer to the border have, as we would expect, lower silhouette scores. By focusing on the data points with higher silhouette scores, the analysis stays away from the liminal zones in the map and draws our attention to those sets of decisions that are most emblematic of the cluster's content.

Figure 4 shows each cluster's proportionate share of the jurisprudence per year. There are three striking features of this visualization. First, those clusters that have negative silhouette scores (the light blue, dark orange, and dark red clusters) accounted for a relatively small part of the jurisprudence, suggesting that the overall conclusions drawn from the clusters might tend towards being comprehensive, as opposed to narrow.




Figure 4 The proportionate representation of each cluster in each year under study.

Second, the dark blue cluster suggests that there was some sudden change in the SST's jurisprudence, beginning in 2022.

Finally, the dark red and light grey clusters look decidedly strange. Unlike every other cluster, both clusters are focused around the year 2017, with only some representation in the year immediately preceding (2016) and the year immediately after

(2018). I wondered whether this was evidence of some discrete and sudden jurisprudential change and investigated these clusters closely.

Ultimately, I was unable to detect any pattern or trend that explained the clusters, other than the fact that they grouped around the year 2017. I concluded that these clusters are, unlike most of the others I discuss here, unmoored from the meaning of the text. Scholars have, for generative large language models, investigated a phenomenon dubbed “hallucination,” where text-generating models “provide information which is incorrect and inappropriate, presented as fact.”⁵⁷ Some computer scientists suspect that hallucinations are products of the model training process, where the models notice a pattern in the training data that is idiosyncratic or peculiar to that dataset. It might be the case here, for example, that the model was trained on an unbalanced dataset where, for some reason, all documents that referenced “2017” looked substantively different than every other document in the dataset. Given “the inherent opacity of [large language model] parameters,” combined with the current state of the science on artificial intelligence model interpretation, “it is difficult to explain their behavior”⁵⁸—meaning that this supposition is only my best guess.

⁵⁷ Nick McKenna et al, “Sources of Hallucination by Large Language Models on Inference Tasks” (2023) arXiv preprint arXiv: 2305.14552, online: <<https://arxiv.org/abs/2305.14552>>.

⁵⁸ *Ibid* at 1

As I examined the clusters, I found three additional clusters that looked like artificial intelligence model hallucinations. The small dark pink and dark red clusters at the top of the map and the dark gray cluster discussed above did not appear to correlate




Figure 5 This map shows each cluster labeled after my subjective assessment of its meaning.

with any common meanings. The fact of this type of error's presence underlines the importance of human-in-the-loop validation of conclusions rendered by artificial intelligence systems. It also reminds us how vulnerable artificial intelligence systems are to absorbing and reproducing all manner of biases from their training data.

6.3.2 Qualitative analysis of the clusters

All in, this means that the model returned ten clusters that I assessed as meaningful, for the purposes of this research. I now turn to my assessment of those clusters and describe the law of employment insurance adjudication in Canada, as it appears when inductively analyzed. I begin in the top left of the map and work clockwise through the clusters.

6.3.2.1 Dark pink: Refusing extension requests for reconsideration requests

This cluster has 140 members, of which 98 have positive silhouette scores. Approximately 25% of the appeals in this cluster were allowed, which is slightly below the global



Figure 6 The key terms and words in this cluster show that the core theme relates to how the fund administrator adjudicated applications to hear late reconsideration requests.

jurisprudential mean of 29%. After examining the five cases with the highest silhouette scores,⁵⁹ I determined that this cluster concerned a procedural feature of EI appeals. Before an appeal can be filed with the tribunal, the claimant must first ask the fund administrator to reconsider the decision the claimant disagrees with.⁶⁰ For a reconsideration to be considered, it must be filed within 30 days of the date when the claimant was notified of the decision, or after “any further time that the Commission may allow.”⁶¹ The Regulations explain that the fund administrator “may allow a longer period to make a request for reconsideration of a decision if the Commission is satisfied that there

⁵⁹ *MM v. Canada Employment Insurance Commission*, 2017 SSTGDEI 49, *TC v. Canada Employment Insurance Commission*, 2018 SST 1185, *GC v. Canada Employment Insurance Commission*, 2019 SST 1523, *SR v. Canada Employment Insurance Commission*, 2018 SST 307, and *WM v. Canada Employment Insurance Commission*, 2018 SST 328.

⁶⁰ *Employment Insurance Act*, *supra* note 33, s 112.

⁶¹ *Ibid*, s 111.

is a reasonable explanation for requesting a longer period and the person has demonstrated a continuing intention to request a reconsideration.”⁶²

In this cluster, the common denominator was that claimants missed the thirty-day window to file a reconsideration request, asked the fund administrator to grant an extension of time, so they could file a request, but that motion to extend time was refused. They then appealed the refusal to extend time to SST. Before the tribunal, the question was whether the fund administrator acted “judicially” when it refused the motion to extend time. As the Member explained in *M. M. v. Canada Employment Insurance Commission*, the question to be answered was whether “the decision-maker (in this case the Commission [the fund administrator]): acted in bad faith, acted for an improper purpose or motive, took into account an irrelevant factor, ignored a relevant factor, or acted in a discriminatory manner.”⁶³

In each of the five cases surveyed, the claimant alleged that the fund administrator ignored or did not properly consider a relevant factor. In three of the cases, the tribunal found that the fund administrator acted judicially. In two cases, the tribunal allowed an appeal, finding that the administrator did not consider a relevant factor. In *S. R. v. Canada Employment Insurance Commission*, for example, the Member reasoned that the fund administrator ought to have given consideration to the facts that (i) the claimant was misled about the reconsideration process by an employee of the fund and told that there was no extension request process, (ii) the claimant was further told by staff “not to bother” with an appeal, and (iii) the claimant had no experience with the employment insurance application process.⁶⁴ In the other case where the tribunal allowed the appeal, the tribunal concluded that the administrator ought to have given more weight to the fact that the claimant mistakenly thought that he had filed a reconsideration request, and only realized his mistake a year later.⁶⁵

⁶² *MM. v. Canada Employment Insurance Commission*, 2017 SSTGDEI 49, at para 26. Also see *Reconsideration Request Regulations*, SOR/2013-63.

⁶³ 2017 SSTGDEI 49, at para 29.

⁶⁴ *SR v. Canada Employment Insurance Commission*, 2018 SST 307 at para 22.

⁶⁵ *WM v. Canada Employment Insurance Commission*, 2018 SST 328.

full EI entitlement, but the fund administrator refused, concluding that they had not established that there was “good cause” for the late application. In each case, the claimant said that the reason for their late application was that they misunderstood the claims process:

- ***IR v. Canada Employment Insurance Commission, 2021 SST 188***: The claimant wrongly assumed they needed proof of their employment from their employer before they could apply for benefits.
- ***AO v. Canada Employment Insurance Commission, 2020 SST 770***: The claimant improperly relied on advice from friends about when they should apply for benefits.
- ***SL v. Canada Employment Insurance Commission, 2016 SSTGDEI 97***: The claimant delayed applying for benefits for three years because family members were ill, she became ill, and she did not have all the paperwork that she wrongly thought she needed to apply.
- ***KB v. Canada Employment Insurance Commission, 2020 SST 700***: The claimant improperly relied on advice from their employer's human resources department about when they could apply for benefits.
- ***B.H. v. Canada Employment Insurance Commission, 2018 SST 101***: The claimant delayed applying because she did not understand that she was eligible and the fact of her pregnancy and birth of her child prevented her from researching the matter further.

The tribunal rejected each appeal after asking whether each claimant acted as a “reasonable person.” The decision makers concluded that the claimants failed to establish good cause because they failed to demonstrate that they proactively researched how to make a claim.

6.3.2.3 Light pink: Making the wrong election for parental leave benefits

The light pink cluster has 139 members, all of which have positive silhouette scores. This cluster featured a high grant rate (65% of appeals were allowed or allowed in part,

This cluster concerned appeals by claimants who disputed this aspect of the law. In each case that I reviewed, the claimant indicated that they wanted the extended option on their application, but later said that they actually wanted the standard option.⁷² After the fund administrator refused their requests to re-elect saying that their decision was irrevocable, the claimant appealed to the tribunal.

In each case, the legal question the adjudicator asked was whether the claimant could prove that they meant to elect the standard option but selected the extended option erroneously. As one member explained:

[31] I find the Claimant's choice on her application doesn't reflect what she wanted for her parental benefit type. In considering the Claimant's testimony and the inconsistent information on her application for benefits, I find that it is more likely than not that she elected standard parental benefits.

[32] The law does not allow a Claimant to change her election after she has been paid parental benefits. However, as I find the Claimant did not elect standard benefits, there is nothing to revoke. Rather, the Claimant should be put back in a position consistent with her true choice of standard parental benefits.⁷³

In these cases the tribunal allowed appeals, reasoning that it was correcting basic errors and honouring the claimant's true choice, not the downstream consequences of a typographical mistake. Notably, if we review Figure 4—the chart that shows the proportionate representation of each cluster in each year under study—we see that this cluster steadily accounted for a larger share of the tribunal's jurisprudence between 2019 and 2021, before dropping substantially in 2022. It may be that this is indicative of some change in the way the fund administrator makes decisions regarding these cases, some change in the way claimants make their election, or some change in application patterns during the height of the COVID-19 pandemic.

⁷² *LM v Canada Employment Insurance Commission*, 2021 SST 932, *ZZ v Canada Employment Insurance Commission*, 2021 SST 610, *AM v Canada Employment Insurance Commission*, 2022 SST 737, *NP v Canada Employment Insurance Commission*, 2021 SST 754, and *JN v Canada Employment Insurance Commission*, 2021 SST 716.

⁷³ *LM v Canada Employment Insurance Commission*, 2021 SST 932.

administrator to decline to write-off or to decline to reduce an overpayment liability.⁷⁸ Altogether, this means that the tribunal's jurisdiction to consider an overpayment was limited to the question of whether the fund administrator had, in a strict mathematical sense, calculated the quantum of the overpayment correctly.

The fact of the limited jurisdiction of the tribunal caused understandable frustration to claimants who felt that they were being unfairly treated. In *DH v Canada Employment Insurance Commission*,⁷⁹ a claimant who failed to report that he worked while collecting benefits (causing the administrator to remit him more funds than he was entitled to) explained that he was, at that time, challenged by significant addictions. He explained that the overpayment debt was causing him significant hardship. The tribunal dismissed his appeal, explaining:

I understand the Claimant has been going through a lengthy and stressful process to put his past behaviour behind him. I truly sympathize with the difficult financial circumstances he is experiencing. However, I do not have any discretion to waive the overpayment, no matter how compelling the Claimant's circumstances may be. The law simply does not allow me to relieve the Claimant from responsibility for the overpayment. I cannot ignore the law, even if the outcome may seem unfair.⁸⁰

The low grant rate in this cluster, combined with these sorts of statements, suggests that claimants were often disappointed to find that the tribunal could not or would not address what seemed to them to be an obvious injustice.

⁷⁸ *Employment Insurance Act*, *supra* note 33, s 112.1.

⁷⁹ 2021 SST 338.

⁸⁰ *Ibid* at 18.

6.3.2.5 Dark olive: Definition of earnings and the allocation of earnings

This cluster has 144 cases, of which 116 have positive silhouette scores. This cluster’s grant rate of 37% is similar to the overall grant rate. The five cases I reviewed⁸¹ showed that this cluster concerned claimants who thought that they ought to receive regular benefits when their earnings were interrupted, but who the fund administrator determined were entitled to no benefits or reduced benefits because of earnings they acquired during their




Figure 11 The most important terms and words from this cluster show that the critical issue concerned the correct allocation of a claimant’s earnings, and any impact that allocation had on their eligibility to collect benefits.

benefit period. These disputes were a product of a part of the law that holds that earnings a claimant receives while receiving benefits can be deducted from their weekly remittance, up to the point that they receive no benefits at all.⁸²

Before the tribunal, this aspect of the law produced two sorts of disputes. First, some claimants disputed the fund administrator’s definition of “earnings.” The *Regulations* define earnings, with a few narrow carveouts, as all “pecuniary or non-pecuniary income that is or will be received by a claimant from an employer.”⁸³ Despite

⁸¹ *OC v. Canada Employment Insurance Commission*, 2015 SSTGDEI 46, *KM v Canada Employment Insurance Commission*, 2022 SST 421, *JM v Canada Employment Insurance Commission*, 2020 SST 345, *MM v. Canada Employment Insurance Commission*, 2019 SST 235, and *JR v Canada Employment Insurance Commission*, 2021 SST 643.

⁸² *Employment Insurance Act*, *supra* note 33, s 19.

⁸³ *Employment Insurance Regulations*, *supra* note 67, s 35.

this broad definition, claimants objected to some earnings classifications made by the fund administrator. For example, in *KM v Canada Employment Insurance Commission*,⁸⁴ the claimant received some compensation in the form of gift cards, which they said should not count as income. The tribunal disagreed.

The second class of dispute, while distinct, was related. The *Regulations* explain that lump-sum payouts received from former employers will be counted as weeks of pay beginning at the point of the interruption of earnings. For example, if a claimant receives a large payout for unused vacation time or a large severance payout, that total sum will be amortized out (where the sum is divided by their average weekly pay) on a weekly basis.⁸⁵ If the lump-sum is large enough, this amortization might functionally eliminate their entitlement because it would be deducted against their earnings for each week of benefits that they may have been eligible for otherwise. In essence, this section acknowledges that some payouts received by employees are designed to protect and support a person during a period of unemployment and, if a person has access to such benefits, they do not need EI.

For these cases, the key question was whether the fund administrator correctly characterized the nature of the lump-sum payment as earnings. For example, in *DB v Canada Employment Insurance Commission*,⁸⁶ the claimant received \$22,600 from their employer as part of a resolution/settlement to a suit brought by the claimant. The fund administrator argued that this sum ought to be considered as earnings, but the claimant successfully argued that this sum was not earnings, but compensation for a waiver of a right. Here the claimant proved that he was entitled to reinstatement but accepted the sum in exchange for waiving that right. The tribunal ruled that this sum was therefore not earnings, effectively making the claimant eligible for benefits. This case, however, was an exception. In every other instance, the tribunal did not have difficulty concluding that large payments received by the claimant were, in fact, earnings and could be amortized going forward. This functionally meant that the claimant was not entitled to EI benefits.

⁸⁴ 2022 SST 421.

⁸⁵ *Ibid*, s 36.

⁸⁶ 2019 SST 566.

program.⁸⁹ People who commit misconduct can face criminal sanctions (if, for example, they are convicted of fraud) and/or administrative sanctions. There are two categories of administrative sanction: financial penalties⁹⁰ and increases in the number of hours a person must work to become eligible for EI.⁹¹

Cases in this cluster were factually similar: after some sort of investigation, the fund administrator concluded that the claimant failed to disclose all of their earnings, and the fund administrator decided to impose a sanction for knowingly misrepresenting a fact. On appeal, there were two questions the tribunal asked. First, the tribunal asked whether the claimant “knowingly” misrepresented. Here the tribunal asked whether the claimant, who submitted false information, was aware that the information they were providing was wrong. For example, in *GC v. Canada Employment Insurance Commission* the tribunal allowed an appeal because it found that the claimant who submitted erroneous information “had no reason to disbelieve his employer’s system and instructions,” which was in fact wrong.⁹²

The second major question asked in this cluster was whether the fund administrator imposed a penalty “judicially.” The decision to impose a sanction is discretionary, meaning that even if a claimant commits a form of misconduct, the fund administrator may or may not choose to levy a penalty against that person. On appeal, the tribunal may consider whether that discretion was exercised “judicially.” In cases where the tribunal asked whether the imposition of some sanction was judicial, it framed its assessment in terms of the adequateness of the fund administrator’s explanatory reasons for the imposition of the penalty. For example, in *MG v. Canada Employment Insurance Commission*, the tribunal found that the fund administrator did not act judicially because it failed to “explain its consideration of the overall impact and the seriousness of [an] additional sanction.”⁹³ Similarly, in *SV v. Canada Employment Insurance Commission*, the

⁸⁹ *Employment Insurance Act*, *supra* note 33, s 38.

⁹⁰ *Ibid*, s 38(2).

⁹¹ *Ibid*, s 7.1.

⁹² 2019 SST 596, at para 88

⁹³ 2015 SSTGDEI 28, at para 37.

tribunal found that the fund administrator failed to demonstrate that it took into account the claimant’s compromised psychological health when it imposed a sanction.⁹⁴

6.3.2.7 Light purple: Disentitlement to regular benefits

This cluster has 254 members, of which 239 have positive silhouette scores. The appeals in this cluster had an above average grant rate, with an allowance rate of just over 50%. After reviewing five cases from this cluster,⁹⁵ I determined that these cases primarily concerned the related EI law principle of disentitlement. A claimant becomes disentitled to regular benefits—meaning that they cannot collect benefits—if they fail to prove that they meet the qualifying requirements for benefits.⁹⁶ Disentitlement is fundamentally remedial: if a claimant addresses the underlying concern, they will become entitled to benefits again.

There are many potential grounds for disentitlement. My review, however, suggests that one ground looms large in the jurisprudence. The *EI Act* requires claimants who seek or who are receiving benefits to (i) remain “capable of and available for work and unable to obtain suitable employment”⁹⁷ and (ii) provide proof upon request that they are “making reasonable and customary efforts to obtain suitable employment.”⁹⁸ Together, these sections explain that claimants who are receiving regular benefits must actively try to re-enter the workforce. Three cases I reviewed explicitly turned on this question.⁹⁹ Each case also turned on a similar factual matrix: after a claimant who lost their employment enrolled in an educational program while collecting benefits, the fund administrator

⁹⁴ 2018 SST 1169.

⁹⁵ *CB v Canada Employment Insurance Commission*, 2022 SST 1505, *JP v Canada Employment Insurance Commission*, 2021 SST 320, *CE v Canada Employment Insurance Commission*, 2023 SST 634, *KB v Canada Employment Insurance Commission*, 2020 SST 683, and *PP v Canada Employment Insurance Commission*, 2023 SST 462.

⁹⁶ *Employment Insurance Act*, *supra* note 33, s 49.

⁹⁷ *Ibid*, s 18(1)(a).

⁹⁸ *Ibid*, s 50(8).

⁹⁹ *CB v Canada Employment Insurance Commission*, 2022 SST 1505, *JP v Canada Employment Insurance Commission*, 2021 SST 320, and *CE v Canada Employment Insurance Commission*, 2023 SST 634.

work because of on-the-job misconduct,¹⁰² the claimant takes a leave from their job,¹⁰³ the claimant is participating in a labour action,¹⁰⁴ the claimant is in prison,¹⁰⁵ or the claimant is outside of the country.¹⁰⁶ In one of the cases I reviewed, the tribunal upheld a finding that the claimant was disentitled because she moved from Canada to California.¹⁰⁷ In another case, the tribunal upheld a finding that a claimant was properly disentitled after she was suspended for failing to comply with her employer's mandatory COVID-19 vaccination policy.¹⁰⁸ The matter of COVID-19 and EI law is one to which we will soon return.

6.3.2.8 Dark green: Disqualification from coverage because the claimant voluntarily left their employment

This was the largest cluster in the dataset. It includes 448 decisions, of which 393 have positive silhouette scores. The allowance rate within the cluster is 27.5%, below the global jurisprudential average. After reviewing the cases with the highest silhouette scores,¹⁰⁹ I determined that this cluster engaged the EI law concept of disqualification. Disqualification concerns instances where a claimant, because of their behaviour, loses access to EI coverage. The cases I reviewed in this cluster concerned the same basis for disqualification: voluntarily leaving employment without just cause.¹¹⁰ Put differently, the law explains that people who intentionally end their employment without a reason that the law recognizes as just are not authorized to receive benefits. For a reason to be just, the

¹⁰² *Employment Insurance Act*, *supra* note 33, s 31.

¹⁰³ *Ibid*, s 32.

¹⁰⁴ *Ibid*, s 36.

¹⁰⁵ *Ibid*, s 37(a).

¹⁰⁶ *Ibid*, s 37(b).

¹⁰⁷ KB v Canada Employment Insurance Commission, 2020 SST 683.

¹⁰⁸ PP v Canada Employment Insurance Commission, 2023 SST 462.

¹⁰⁹ OA v Canada Employment Insurance Commission, 2020 SST 53, SN v. Canada Employment Insurance Commission, 2018 SST 1182, RB v Canada Employment Insurance Commission, 2020 SST 206, SL v Canada Employment Insurance Commission, 2020 SST 445, and SV v. Canada Employment Insurance Commission, 2019 SST 482.

¹¹⁰ *Employment Insurance Act*, *supra* note 33, s 30.

dismissed because “it would have been reasonable for him to continue working while he looked for other employment.”

- ***SL v. Canada Employment Insurance Commission, 2020 SST 445***: The claimant left his job as an airplane pilot without arranging for new employment. His appeal was dismissed, with the tribunal concluding that he “could have continued this employment until he received confirmation of another employment prior to quitting.”
- ***SV v. Canada Employment Insurance Commission, 2019 SST 482***: The claimant left her job because it was negatively impacting her health. The tribunal accepted her explanation and allowed the appeal, concluding that it was convinced that “staying would lead to a deterioration of her mental and physical health.” Therefore, there was no reasonable alternative available to her other than leaving her employment.

For the claimants, these disqualifications were consequential. Unlike cases where a claimant loses coverage because of a disentitlement, a disqualification cannot be remediated by a subsequent change in claimant behaviour. While the *EI Act* specifies that disqualifications can be for a determinate or an indeterminate amount of time, for instances where a person is disqualified for voluntarily leaving their employment without just cause, the disqualification is automatically indeterminate, meaning that they will not become again eligible for coverage until they find new employment and requalify for insurance by working again the minimum number of hours.¹¹²

¹¹² *Employment Insurance Act, supra* note 33, s 30(1)(a).

such as to impair the performance of the duties owed to his employer and that, as a result, dismissal was a real possibility.”¹¹³

In these cases, as with the cases of disqualification for voluntarily leaving one’s employment, the core question was factual. Notably, the nature of proof and the location of the persuasive burden matter here. As the following review shows, the tribunal routinely held that the fund administrator was unable to establish that the claimant was dismissed because of misconduct, preferring the evidence of the claimant. From a doctrinal perspective, the law here can best be described as fuzzy. In several cases, tribunal members cited Federal Court of Appeal authority for the proposition that the *onus* to prove misconduct lay with the fund administrator, but it is not clear (at least to me), that there is in fact authority for the legal proposition that the legal burden to prove misconduct on appeal lies with the responding fund administrator.¹¹⁴

Regardless, we can clearly observe in the following cases (even if the *de jure* reasons may be less than compelling) that there is a *de facto* shift of the onus: once the claimant has set out their version of events and explains why they had not engaged in misconduct, the evidentiary burden (or, the quality of the evidence that must functionally be led) on the fund administrator to disprove it necessarily increases:

- ***CJ v. Canada Employment Insurance Commission, 2015 SSTGDEI 215***: The claimant’s employment at a food processing centre was terminated because the employer believed that he “had sabotaged part of the production line by attaching a drumstick bone to a chicken.” The fund administrator disqualified him from receiving benefits. On appeal, and after hearing from the claimant, the tribunal found that the evidence presented by the fund administrator could not establish the allegation. The appeal was allowed.
- ***SM v. Canada Employment Insurance Commission, 2014 SSTGDEI 50***: The claimant’s employment was terminated after an altercation at work and the claimant was disqualified from coverage. The appeal was allowed because the tribunal found that the evidence showed the claimant was provoked and his

¹¹³ *Mishibinijima v. Canada (Attorney General)*, 2007 FCA 36 at para 14.

¹¹⁴ Members cited *Canada (Attorney General) v. Granstrom*, 2003 FCA 485 and *Lepretre v. Canada (Attorney General)*, 2011 FCA 30 for the proposition that the evidentiary burden to prove misconduct lies with the fund administrator, but I do not read either case to endorse such a position or, at least, not to endorse the position clearly.

response could be “considered [as] a natural defence response.” Therefore, the claimant could not be said to have engaged in willful misconduct.

- ***GC v. Canada Employment Insurance Commission, 2016 SSTGDEI 73***: A former director of a public municipality was terminated after it was alleged that he helped falsify timesheets. The fund administrator disqualified the claimant, saying that he was dismissed because of misconduct. The claimant disputed this version of events. The tribunal allowed the claimant’s appeal, concluding that the fund administrator had not led sufficient evidence to establish its allegations and could not, therefore, establish that the claimant was dismissed for misconduct.
- ***DL v. Canada Employment Insurance Commission, 2016 SSTGDEI 39***: the claimant’s employment was terminated after he refused to undergo a fitness-to-work evaluation, which his employer said amounted to insubordination. He was disqualified from coverage. The tribunal allowed the appeal, finding that any mistakes made by the claimant did not rise to the level of misconduct.
- ***NM v. Canada Employment Insurance Commission, 2015 SSTGDEI 170***: The claimant was dismissed for insubordination and being rude. On appeal, the tribunal found that the evidence did not show that any poor behaviour rose to the level of misconduct warranting disqualification.

At bottom, these cases show how a gap can open between an employer’s vision of a legitimate dismissal for misconduct, the fund administrator’s view, and the finding of a tribunal after reviewing the evidence as presented by the parties.

6.3.2.10 Dark brown: Disqualification from coverage because the claimant’s employment was terminated because of misconduct (non-compliance with mandatory vaccination policy)

The dark blue cluster has 430 members, of which 423 decisions have positive silhouette scores. This cluster featured a very high dismissal rate, with over 92% of cases being dismissed wholly, summarily, or in part. This cluster was also unique for suddenly emerging in 2021.




Figure 16: The key terms and words show that this cluster focused on disqualifications based on non-compliance with an employer’s mandatory vaccine policy.

My review of the five cases¹¹⁵ showed that this cluster is a close cousin to the disqualification for misconduct cluster, but with a distinct thematic focus. In the Fall of 2021, many workplaces required employees to get vaccinated against the COVID-19 virus. Employees who failed to get vaccinated or who failed to provide an acceptable reason explaining why they could not get vaccinated were often suspended or terminated for their non-compliance with a workplace protocol.

The fund administrator took the position that employees who were suspended from work, and whose income was therefore interrupted, were disentitled to employment insurance benefits during the length of the suspension, while workers whose employment was terminated were disqualified from coverage for misconduct. The tribunal found in each case that non-compliance with these policies amounted to misconduct. As one member reasoned:

[F]or EI purposes, I find that the fact that the Claimant was placed on leave without pay made her unable to fulfill her duties toward her employer. The

¹¹⁵ *CM v Canada Employment Insurance Commission*, 2022 SST 1652, *SP v Canada Employment Insurance Commission*, 2022 SST 1323, *TP v Canada Employment Insurance Commission*, 2023 SST 1217, *AP v Canada Employment Insurance Commission*, 2022 SST 1410, and *MH v Canada Employment Insurance Commission*, 2023 SST 559.

Claimant did not comply as a result of an informed decision she made to not be vaccinated despite the consequences outlined in the policy.¹¹⁶

Before the tribunal, claimants took issue with their employer's policies, claiming that they violated their rights, offended provisions of their employment contract, or were inappropriately denied religious exemptions. The tribunal declined to hear these arguments, reasoning that they were matters that the tribunal did not have jurisdiction over. Members said that claimants who disagreed with their employers' policies could raise these concerns with other tribunals, while the SST would remain squarely focused on whether the definition of misconduct as set out in the *EI Act* was met. Members often wrote directly to claimants to address their frustration, for example:

[46] However, none of the Claimant's arguments or submissions change the fact that the Commission has proven on a balance of probabilities that she was suspended and subsequently terminated because of conduct that is considered misconduct under the *EI Act*.

[47] I understand why the Claimant may not agree with my decision. I do not address the fundamental legal or factual issues that the Claimant raises, for example regarding her collective agreement, [or] religious accommodations. This does not make the decision unreasonable. The key problem with such arguments is that I am not permitted, by law, to address.¹¹⁷

In some ways, this cluster is the most interesting, at least when considered against the methodology employed here. Not only has this approach shown us different genres and themes within the law, but it can also detect how the law adapted and understood a black swan event—the COVID-19 pandemic.

¹¹⁶ *SP v Canada Employment Insurance Commission*, 2022 SST 1323 at para 31.

¹¹⁷ *MH v Canada Employment Insurance Commission*, 2023 SST 559.

6.4 A theoretical model of EI law amenable for teaching

This paper began by asking whether new computational approaches could enable new classes of pedagogy for law teachers. If the life of law is experience, I suggested the new artificial intelligence models could help us induct up from masses of raw evidence to develop a sensible model of an area of law that could, especially for students interested in public law or poverty law, better describe the field and help them develop both practical and critical insights. Or, put differently, for those fields of law defined by mass experience that are, at least in the core curriculum, neglected in law schools—areas such as prison law, tenancy law, deportation law, unemployment law, social benefits law—artificial intelligence might help us develop new theoretical accounts that could meaningfully advance our knowledge and our students’ knowledge.

As is now clear, I believe, this approach has produced a vision of employment insurance law that could assist in classroom teaching. First, consider how the model, while presenting ten thematic clusters, also suggests three overall organizing themes. At the top of the map, we see three clusters concerning procedural disputes. In the middle, we see three clusters concerning financial disputes and disputes related to the quantum of benefits. And at the bottom, four categories of qualitative substantive dispute about eligibility. Employment insurance is a massively complex program, but only some disputes make their way to final adjudication. Here, we now have a vision of what those disputes are.

For teaching purposes, we now know where to focus future lawyers’ attention. Now, too, they will have perspectives with which to read the rest of the statute and regulations to understand both how these disputes evolve and how others might, for various reasons, not finally end up in adjudication. Equally, critical questions force themselves on us. The SST explicitly does not have equitable jurisdiction, meaning that it must decide the right legal answer, not necessarily the fairest answer. But over and over, we see how the tribunal stretches itself to develop a legal jurisprudence responsive to




Figure 17 Once each cluster is labelled, we can see that the clusters naturally arrange themselves into three categories: (1) disputes concerning EI applications, (2) disputes concerning the quantum owed to (or owed by) by the claimant, and (3) substantive disputes about whether the claimant may collect EI, because of their behavior.

equitable concern. Is the law of equity, in this context, the law of procedural fairness? It would appear so. Think, for example, how students might think through the question of the parental leave application and how the tribunal found a way to develop its jurisprudence to honour a parent’s intention, despite the alleged irrevocability of the paperwork. But, too, is it not interesting that fairness concerns tend only to find expression in that first cluster of cases at the top of the map? The closer a dispute is to the tribunal’s process, it appears, the more room there is for an equitable version of fairness.

This only scratches the surface of how teachers might use this sort of model. But by now, the key point should by now be clear: “The law is a discipline that is brought to life by us and our students through its practice.”¹¹⁸ And, especially for fields of mass legal experience, methods that help us and our students conceptualize the law-in-action will

¹¹⁸ Jeremiah Ho, “Function, Form, and Strawberries: Subverting Langdell” (2015) 64:4 Journal of Legal Education 656 at 657.

immeasurably improve the quality of learning. This, then, is at least one contribution of artificial intelligence to law: by helping us orient to the great mass of legal experiences, we can develop new valuable vantage points on the law as practiced, as lived, and as experienced.

7. Conclusion: The Paths of Computational Law

I

It was not a fair thing to do. When Justice Oliver Wendell Holmes Jr. started speaking at Boston University School of Law in 1897, he broke that most important judicial rule: make sure that you tell the law students that they are special. His talk would soon be published in the *Harvard Law Review* and would become—incredibly, for a lecture from a 19th century judge about law—a prototypical example of one of those distinctly anglophone essays that we all aspire to write: a polemic with an ambivalent thesis statement that only gets around to its major point in its last few sentences.¹

But for now, let us imagine him standing at the lectern, surveying a room of men wearing hats, about to speak. Soon, the atmosphere will change as it dawns on the assembled law students, professors, and lawyers that they were not the heroes of this story. “When we study law,” he began, “we are not studying a mystery but a well-known profession.” What is a lawyer? A lawyer is a professional who trades their knowledge and analyses for a fee. Sometimes, at least in “societies like ours,” the public assigns some decision-making authority to judges, for whom “the whole power of the state will be put

¹ For a discussion on the many ways in which this essay can be conceived of a classic and a general inspiration for my own reading of the essay see Dan Priel, “Holmes’s ‘Path of the Law’ as Non-Analytic Jurisprudence” (2016) 35:1 UQLJ 57.

forth, if necessary, to carry out their judgments and decrees.” And against this power, a person: a person who wants to “know under what circumstances and how far they will run the risk of coming against what is so much stronger than themselves.” So, what is a lawyer? A lawyer is a businessperson who gets paid to predict when others will risk confrontation with an “incidence of the public force.”²

Holmes’ prediction theory of law says that for all that busy and important talk about rights and privileges, statutes and principles, precedent and propositions, lawyers really just make prophecies, prophecies that others will use to, as the expression goes, “govern themselves accordingly.” Yes, we are separated from Holmes’ talk by over one hundred years, but can you not still feel that the stomachs of the students there are tightening? We cannot really help ourselves. Those of us who teach in law schools, work in courts, or practice, tell large tales about the profession. Integrity, honour, advocacy, right and wrong, truth, justice, history: these are watchwords of the lawyer. And, Holmes argued, fundamentally misleading ones.

If the business of law is the business of prophecy, there is no percentage in morality. A client might be good or they might be bad, a legal effect might be just or it might be unjust, but what of it? Clients do not come to lawyers because their legal training has made them good moral counsellors or because it has made them thoughtful observers of the public condition, they come for help to legally maximize that which they want to be maximized and legally minimize that which they want to be minimized. It is hardly a contentious point: neither of us knows anyone who has visited a lawyer to become a better person.

This means that all of that grandiose lawyerly talk is just sanctimony or, worse, just confusing. Maybe, Holmes said, we would be better off if “every word of moral significance could be banished from the law altogether”³ so, at last, we could just speak plainly to each other. And students could focus on what they ought to focus on: “a right study and mastery of the law as a business with well understood limits,” that is, “a body

² Oliver Wendell Holmes Jr., “The Path of Law” (1897) 10:8 Harv L Rev 457 at 457. [“The Path of Law”]

³ *Ibid* at 464.

of dogma enclosed within definite lines.”⁴ Each year thousands of eager students write law school applications and explain why they want to be lawyers. How many say that their objective is the focused study of the business opportunity produced by the fact of encounters with public force? Surely, not many. It would therefore be destabilizing to listen and be told, by no less an authority than a judge (soon a judge of the American Supreme Court), that the law school engendered conviction in greater purpose is cant. “I ask you,” he implored, “for the moment to imagine yourselves indifferent to other and greater things.”

And it got worse. Holmes told the law students that there were others out there with a legitimate claim to law’s larger purposes. How should law orient to law morality, ethics, and politics? These were questions for the historians and, horror, the economists. If lawyers busy themselves with business, it was the job of social scientists to examine law as the “external deposit of our moral life”⁵ and think about what it could and should be. First, historians. In contrast to prophecy, Holmes argued that “the rational study of law is still to a large extent the study of history.”⁶ Why is it, he wondered, that our judges obey some rules and render decisions somewhat? Inertia:

Most of the things we do, we do for no better reason than that our fathers have done them or that our neighbors do them, and the same is true of a larger part than we suspect of what we think.⁷

To know why law is what it is, as opposed to just knowing how it will treat a problem, is to wonder about circumstances that led to its development and shape. To think about law critically, to treat its claims with suspicion, to get at the truth of its function, requires historical methods. It is only after one gets the “dragon out of his cave on to the plain and in the daylight” that one can decide whether this dragon is one worth keeping.⁸

Lawyers, he said, those “blackletter” practitioners: these are the people of the present; but the professionals of the future, those are the statistician and the “master of

⁴ *Ibid* at 459.

⁵ *Ibid*.

⁶ *Ibid* at 469.

⁷ *Ibid* at 468.

⁸ *Ibid* at 469.

economics.”⁹ So, second, it would be the critical thinkers, not idiosyncratic rule followers, licensed to use new empirical methodologies as an acid, revealing law’s purposes and envisioning a larger “rational generalization.”¹⁰ Facing the coming twentieth century, Holmes offered a high modernist vision of what law could be: intentional, evolved, human, and built. Aided by new methods and tools that could plumb the unconscious, that could measure the population, that could re-describe the world, that could account for our best and worst impulses, law could be made harmonious with, and catalytic of, human flourishing. Though he could not have thought of it in precisely these terms, he proposed a law that could complement skyscrapers and highways, surrealism and psychoanalysis, and (soon) an FDR ready to offer America a new deal.¹¹

Whither those assembled law students? Lawyering might, he said, make them “counsel for great corporations” or win them “an income of fifty thousand dollars,” but it would not make them happy. To satiate their intellect, to find a larger meaning for their purpose, they would need to step out from behind their law books. Do not, he said, just “become a great master in your calling but connect your subject with the universe.” For if one did, if one made common cause with those historians, those statisticians, and those economists, they might “catch an echo of the infinite, a glimpse of its unfathomable process, a hint of the universal law.”¹²

So maybe I have been the unfair one. Maybe Holmes did not break that most important judicial rule. Maybe, in his own roundabout way, he did tell all those law students that, yes, when all was said and done, they would live through times of great change, that they could be of consequence, that they were on modernity’s new *Path of Law*, and that they might just be special after all.

II

⁹ *Ibid.*

¹⁰ *Ibid* at 471.

¹¹ Daniel Hulsebosch, “The New Deal Court: Emergence of a New Reason” (1990) 90:7 Colum L R 1973.

¹² *Path of Law*, *supra* note 2 at 478.

Until lately, no one has talked much about universal law. But in 2016, Benjamin Alarie announced that, just maybe, he had, yes, caught that echo and glimpsed that infinite, unfathomable process. How? By studying tax law.¹³

Holmes' vision, Alarie hypothesized, came over a century too soon. For all the stock that Holmes put in new empirical methods of research and discovery, these remained too rudimentary and too disconnected from legal concern to enable a law of generalizable rationality. It would take the digital revolution, the proliferation of tools of cheap mass communication, the accumulation of data, and the development of machine learning technology to make deep and wide inquiry into law possible, in turn authorizing a new—and better—set of legal relationships among all people on earth: “[t]he legal singularity will arrive,” he argued, “when the accumulation of a massive amount of data and dramatically improved methods of inference make legal uncertainty obsolete.”¹⁴

Here, biography matters. Alarie is not only a Canadian law professor, he is also an entrepreneur. In 2014, he co-founded Blue J Legal, a tax law, machine learning, and legal analytics firm. Blue J's basic product is simple: it has collected a large number of data points about tax law cases and used these to train a machine learning prediction model. Show the model a new case with new data points, and it makes its best guess about how a court will decide the case.¹⁵

Take the legal problem of how workers are classified. For a corporation (and, it must be said, for the worker) it makes a difference whether someone is an employee or an independent contractor: firms owe more to employees than they do to contractors. But how to tell the difference and know whether a worker is an employee or not? “After

¹³ Benjamin Alarie, “The Path of Law: Towards Legal Singularity” (2016) 66:4 University of Toronto Law Journal 443. [Towards Legal Singularity]

¹⁴ *Ibid* at 445.

¹⁵ On the company's website, it is explained that: “Our CEO Benjamin Alarie was the Associate Dean of the Faculty of Law and was invited to judge an IBM Watson competition. He became fascinated by the possibilities of applying AI to tax law and the opportunity to use machine learning to predict outcomes and recommendations. By 2015, our first early prototype was built. In 2016, early adopters firms started piloting Blue J. Quickly thereafter we began selling our tax product commercially before expanding to employment, HR, and US law applications.” See “About Us,” online: <<https://www.bluej.com/about-us>>.

hundreds upon hundreds of cases,”¹⁶ Alarie tells us, the common law failed to articulate a clear test that individuals and firms could rely upon to govern their affairs. Yet, there was a subterranean and persistent logic in how judges made decisions “implicit in [those] hundreds of judgments,”¹⁷ that could be mathematically detected and replicated. This meant that the law could be rendered with greater clarity and precision in the language of the machine than it could be in the language of lawyers.

As Alarie contemplated this model, a new perspective on law came into view. An echo, a glimpse: as machines learn more about the law, and learn the law better than humans, a singularity dawns. The idea of the singularity is not new, rooted in ideas advanced by Ray Kurzweil¹⁸ who was in turn influenced by a 1993 essay by the mathematician Vernon Vinge, which certainly drew on breakthroughs in theoretical physics concerning black holes and the origins of the universe. It describes a moment, in Vinge’s terms, when machines become so capable, so intelligent, so smart (and maybe even so wise), that “our models must be discarded and a new reality rules.”¹⁹ For Alarie, the legal singularity, that moment when old rules would be eliminated and a new law would reign, means:

the elimination of legal uncertainty and the emergence of a seamless legal order, universally accessible in real-time. In the legal singularity, disputes over the legal significance of agreed facts will be rare. There may be disputes over facts, but once found, the facts will map on to clear legal consequences. The law will be functionally complete.²⁰

What does it mean for the law to be complete? Alarie suggests that the legal singularity will be described by three core features. First, it will be complex, but instantly knowable. Instead of relying on the judgement of a lawyer, based on their experience and their doctrinal training, to make a prediction, a machine can make a much more accurate

¹⁶ Towards Legal Singularity, *supra* note 17 at 447.

¹⁷ *Ibid* at 448.

¹⁸ Ray Kurzweil, *The Singularity Is Near: When Humans Transcend Biology* (New York: Penguin Books, 2006).

¹⁹ Vernon Vinge, “The Coming Technological Singularity: How to Survive in the Post-Human Era” (1993) online: <<https://edoras.sdsu.edu/~vinge/misc/singularity.html>>.

²⁰ Towards Legal Singularity, *supra* note 17 at 446.

prediction, based on much more information, much faster. This means that the law can account for more diversity and be definitively known before a dispute. Practically, when a machine knows what the law is, one need not wait for a trial date.

Second, the law will be without gaps. Once the machine has learned the animating objectives of a given law, it will be able to reason consistently to an appropriate outcome, no matter how surprising or novel a factual circumstance might be. If today, legislators, regulators, and decision-makers must deliberate their way through the unexpected, the machine can instantly calculate how novelty ought to be metabolized, considering the law's known preferences. For Alarie, this in turn will produce decidedly positive democratic outcomes: instead of arguing about matters around the edges, "Politics and policymaking, domestic and international, will be dominated by rich debates surrounding what values tax [and legal] systems ought to embrace and reflect."²¹

Finally, the law might achieve equilibrium. Suitably complex, suitably aware of the important objectives, the law might optimize. From one year to the next, one generation to the other, the law will remain stable because it will have developed responsively to the ideal circumstances of human justice. This, Alarie insists, is no techno-futurist idea, but a Rawlsian one: "Reflective equilibrium," the idea that once "regulative principles are brought to light" people will naturally revise their judgements and bring them into accord with morality.²² For Alarie, this is an idea that could be made vital by combining data, facts, and algorithms capable of solving disequilibrium.

All made possible by a machine. Modern algorithms develop models and representations of the world as they are shown a countless number of data points about the world. Alarie's bet is that there is now almost enough data and enough compute power that the law can be computationally modeled and, once so modeled, improved. Holmes' infinite and unfathomable process: Alarie says we can find it by scaling up legal prediction to describe the legal world. And, once so described, it will expand, grow, and improve. No,

²¹ *Ibid* at 453.

²² *Ibid* at 454. The quotation is from John Rawls, *A Theory of Justice* (Cambridge, Harvard University Press, 1971) at 49.

Holmes did not need historians, statisticians, and economists: Holmes needed data scientists.

III

I am sure that this much can be said about the future: it is characterized “by the manner in which human foresight, human plans, and their execution always diverge in the course of time.”²³ All disciplines are guilty of this charge, but I think law especially: we do not much think about our thinking. Between professors’ normative claims about what law ought to become, and lawyers claims about how their clients’ lots ought to be improved, the study of law is often the study of the future.

And the future is a strange place. Between past, present, and future, it is the future that is the least knowable and the place we are most likely to locate our desires, ambitions, and hopes. And, of course, desire is not prophecy because desire really only says something about ourselves. At least this is my claim here. When Holmes envisions an infinite law, he tells us something about his moment; when Alarie imagines a singularity, he tells us something about ours. So, what are each telling us about themselves and their moments?

Let us begin by returning. For Holmes, lawyering had a purchase on the present, not the future. The prediction theory of law, the idea that lawyers were paid to make informed prophecies about what some entity would do, implies that lawyers are uncritical historians. Their job is not to explore history’s meanings, their job is to understand how the events of yesterday portend the events of tomorrow.

In contrast, Holmes gave the new social scientists a monumental task. First, it was their job to excavate the past to uncover the true reasons for things. Second, (did Holmes read Marx?) it was to ruthlessly criticize all that existed so a new world could be made.

²³ Reinhart Koselleck, *Futures Past: On the Semantics of Historical Time* (New York, Columbia University Press: 2005) at 201. He adds: “‘history in and of itself’ always occurs in the anticipation of incompleteness and therefore possesses an open future. That is, in any case, a lesson of all previous history, and whoever wishes to argue the opposite will have to prove his case.”

Once the fact of the matter is discovered, Holmes said researchers could and should re-examine the idea and its function. Is this idea worth keeping or disposing of? Ought it be changed? Can we fashion a new legal arrangement capable of radically improving the human condition?

Modernism. In an age, such as ours, without manifestos, it is jarring to realize that one is encountering a legal *avant-garde*. An admirer of the arts, it is as if Holmes is trying to imagine a legal move that could match the expressionist move to depart from the figurative. In an age of new writing, new painting, new music, Holmes sought to develop a new vision of law.

And like the modernist arts surrounding him, Holmes' vision was a vision of the future. The historian Reinhart Koselleck explains that a distinctive feature of modernity was its new relationship with time. Medieval time, the time of the *ancien regime*, was characterized by repetition and return. The future was circumscribed by an "eschatological limit," the coming biblical apocalypse, making the experience of time cyclical: the past just kept turning up again in the present. But, Koselleck says, the rise of new forms of social relations, combined with new visions of the sacred (and the rise of the secular), authorized in modernity a new sort of future: an indeterminate and unwritten one.²⁴ Once, revolution referred to a cycle; later, it referred to a break with the past's hold on us.

The great sociologist of modernity, Perry Anderson, agrees. The *belle époque*, that late nineteenth century period that stretched to the beginning of the First World War, characterized by the tremendous artistic, technological, and sociological innovation that we now see as modernity's apogee, says that the time is best understood as a "field of force triangulated by three coordinates:" a sense of revolutionary possibility, dramatic invention (the telegraph, electricity, ocean liners, airplanes), and the fact that most of world remained semi-feudal, governed by agrarian or aristocratic powers.²⁵

²⁴ "Modernity and the Planes of Historicity" in *Ibid* at 3-20.

²⁵ Perry Anderson, *The Origins of Postmodernity* (London, Verso: 1998) at 81.

Progress, and its exponents, were radically transforming small pockets of the world, proving that time's repetitions could be disrupted. Modernity, Anderson says, "lived from the non-synchronous,"²⁶ or the fact that its vision of time was not yet totally ascendant: the trajectory of its arrow all the more obvious when backgrounded against a majority experience of repetition. And the archers wore the same uniform. The bourgeoisie, that emerging class that vied against landed interests for power, had a self-conscious identity. It is difficult to imagine today, given how enmeshed elite culture is in popular culture, but, yes, these were the times when these men wore hats.²⁷

These were dynamic times for arts, letters, and law. The fact of the new bourgeois world gave creativity an "aesthetic foil,"²⁸ a crucible, from which novelty could emerge. It was a period of artistic and political innovation, made possible by how unsettled and dynamic social relations were. This helps us understand why Holmes' vision of the law (rational, planned, intentional) can look not so dissimilar from soviet law: each took as given the idea that the world would be remade, and law would have something to do with that remaking. Holmes and modernity writ large each countenanced a type of utopia, and the teleology of those times said that a path might just be followed to it.

In comparison, the legal singularity looks different. Where Holmes wanted to build an infinite law in and for the future, Alarie proposes to expose it today. Between the two, Alarie has more respect for the idea that lawyering is part of some important, larger, collective project. Once viewed at scale and in aggregate, Alarie says, the sum of lawyerly predictions well exceeds the sum of their parts. Viewed together, we can divine the deep and true rules that govern legal decision, though these rules may not be legible to individual practitioners. But armed with new algorithms and new datasets, the collective wisdom of the common law will be surfaced and, perhaps after some tinkering, made available to all.

Algorithms assume the places of Holmes' historian and his lawyer. For Holmes, prediction was the work of the businessperson who happened to be a lawyer but for Alarie it is the whole ball game. It is not that human critical inquiry needs to be focused on the

²⁶ *Ibid* at 82.

²⁷ *Ibid* at 82.

²⁸ *Ibid* at 86.

past or that human ingenuity ought to re-imagine our future: it is finally exposing those things about us that have been true all along, but that we have not been able to well articulate.

Holmes' modernist future this is not. Instead, the singularity's future evinces a relationship with time familiar to critics of postmodernity. Where antiquity was dominated by a repeating past, where modernity was dominated by its relationship with a future to come, is our time dominated by the present? The cultural critic Fredric Jameson argues that, at least since the early 1980s, time's scope has narrowed to the here and now. There has been, he says, "a dissolution of past and future alike, a kind of contemporary imprisonment in the present."²⁹ The future, as a time unlike our own, becomes unimaginable; the past, a collection of "dusty images"³⁰ that, we think, has little bearing on our lives.

Examples best illuminate what is meant here. Let us consider three raised by Jameson. First, the action film. Consider how,

nowadays they are reduced to a series of explosive presents of time, with the ostensible plot now little more than an excuse and a filler, a string on which to thread these pearls which are the exclusive centre of our interest: at that point the trailer or preview is often enough, as it offers the high points of films which are essentially nothing but high points.³¹

Where the art and thought of modernity was focused on deep time and the relationship between past and future, the action film points to a cultural moment fixated on the moment.

Or, second, think of the presentation of visual art. Today, galleries constantly rotate in and out new shows, each warranting its own advertising campaign and opening. Consider how "[t]he installation and its kindred productions are made, not for posterity, nor even for the permanent collection, but rather for the *now* and for a temporality that

²⁹ Fredric Jameson, "The Aesthetics of Singularity" (2015) 92 *New Left Review* 101 at 120.

³⁰ *Ibid.*

³¹ *Ibid* at 105.

may be rather different from the old modernist kind.”³² This is not the age of the permanent collection, it is the age of the artistic event.

Or, third, something that approaches law: the derivative, a strange “mutation of traditional insurance investment.”³³ Today, the derivative—common throughout the market—is more than just a hedge on a future, it is intrinsic to and enables relations between firms and people. Think about how you might build a phone. You would purchase one component from one country, a second from another country, a third from a third country, all of which get assembled in a fourth. Here we have “different currencies, their exchange rates in perpetual flux” and the “risk of unforeseen variation between these exchange rates will then be underwritten by a kind of insurance.” Together, this entire package makes the financial instrument: an entirely unique legal device, that will never be replicated or reproduced. Indeed, it is more than a contract, which is “something with a stable structure and a juridical status.”³⁴ Put plainly, a global economy enabled by the immediate and singular production of financial instruments, each of which has never existed before and never will exist again.

Are there are resonances here with Alarie’s legal singularity? Like the action film and the installation, the singularity speaks to events, not histories. Individual people and firms come to the singularity seeking immediate answers about particular and singular factual problems. And does not, when we step back, the derivative look like the economic analogue of the legal singularity? Indeed, it occurs to me that we best glimpse the legal singularity where we peer into the high velocity world of modern finance. Trades actualized without human oversight, algorithms that move untold sums from here to there, derivatives that coordinate singular and singularly unique economic relations.

The cultural critic Anna Kornbluh would recognize the legal singularity—pace Jameson—as part of a larger cultural trend that defines our moment. Immediacy, a “pulsing effulgence purveys itself as spontaneous and free, pure vibe. Let it flow, let it

³² *Ibid* at 111.

³³ *Ibid* at 117.

³⁴ *Ibid* at 118.

flow!”³⁵ that eschews “organizations and institutions in favor of organic horizontalism.”³⁶ It is a just-in-time way of seeing the world that, and this is the critical point, arrays itself against mediation. This, for Kornbluh, is critical. As our society favours immediacy more, we devalue those institutions where we traditionally came together to make sense of them. What is lost? The active relating and “making available in language and image and rhythm the super-valent abstractions otherwise unavailable to our sensuous perception—like ‘justice’ or ‘value.’”³⁷ With no past and no future, only the desire for immediate cultural, economic, and (maybe) legal gratification there is flow, not conscious and collective attention.

And with immediacy, with a new relation to time, our vision of how society might change necessarily also changes. In an introductory essay for an innovative series on the history of futures, published by *History and Theory*, Zoltán Boldizsár Simon and Marek Tamm also say that something “has changed immensely since the heyday of Western modernity, and that something is the future.”³⁸ Two characteristics of the new ways of conceiving the future that they identify are relevant here, and dovetail with our discussion so far.

First, think of how people in Holmes’ time as opposed to ours imagined a transition from the present to the future. During the moment of modernity, transition was often tied to the sociopolitical domain. One could spot the arrival of the future by spotting the revolution, the upheaval, the transfer of power, or the new deal. In contrast, transition to the future today is often located in the technical, the scientific, or the ecological domains. It is runaway growth, tipping points, or—and this interests us here—the exponential increase in the availability in data and availability of compute power, that anticipate the future. The point: for at least some classes of thinkers in the height of modernity, and at least for some classes of people today, the future looks qualitatively different. In one case,

³⁵ Anna Kornbluh, *Immediacy: or, The Style of Too Late Capitalism* (London: Verso: 2024) at 6.

³⁶ *Ibid* at 7.

³⁷ *Ibid* at 5.

³⁸ Zoltán Boldizsár Simon & Marek Tamm, “Historical Futures” (2021) 60:1 *History and Theory* 3 at 4.

occasioned by political rupture; in the other, quantitative changes in things beyond our reach.³⁹

Second, think of how close the future may seem. At the height of modernity, the future looked distant and uncertain. What would actually come to pass looked wildly unpredictable. But, today, especially for those who are algorithmically inclined, the future has been “defuturized,” where what will come to pass is made theoretically knowable and predictable. Where modernity held utopia at a distance, an algorithmic future holds it close: for we have, we think, the tools and data at hand to know it. Recall what the legal singularity offers: a complete law, extrapolated from all that is already known about law.⁴⁰

Two visions of an infinite and complete law, yet two very different visions. Each is marked by a different relationship to the future, which is in turn a signal of the conditions of our own time. For one: an intensely mediated and built world; for the other, a flow state, extrapolated from our present.

IV

So what? Here is one thing (soon we will discuss another): there is every reason to think that the idea of the legal singularity rests on junk science. Despite the fact that in a follow up text to his original article, Alarie and his co-author, Abdi Aidid, confidently open the book declaring that “We are on the path to the legal singularity”⁴¹ the evidence for the claim is thin on-the-ground. To be sure, there are firms like Alarie’s exploring what algorithms might do for law, but are there really firms that are on a path to algorithmically redescribing legal relations? It is a large claim and, to be fair to Alarie and Aidid, they say that while the singularity is relatively close at hand, it remains at least a few decades away.

Nonetheless, the singularity’s ledger boasts few success stories. If there is a single industry that best augurs the world of prediction to come, it must surely be online advertising. Large firms, we are told, know us better than we know ourselves, hoovering

³⁹ *Ibid* at 15-17.

⁴⁰ *Ibid* at 20-22.

⁴¹ Abdi Aidid & Benjamin Alarie, *The Legal Singularity: How Artificial Intelligence Will Make Law Radically Better* (Toronto: University of Toronto Press, 2023) at 3. [Legal Singularity Book]

up data about us to such a degree that we are fully described. But even in this prediction industry *par excellence*, things are not going well. Online ads, it turns out, perform worse than traditional ads and cost more, and have perhaps produced to an extraordinary bubble that risks taking down the tech and content producing sectors of the economy.⁴²

Alarie has heard this point but he bet on that which will come. Aidid and he explain that “[i]n our view, current systems barely scratch the surface of legal technology’s potential. Critiques that are hyper-focused on today’s capabilities, then, risk becoming outdated by the next technological improvement.”⁴³

What are the consequences of this form of argumentation? Futurism necessarily means making predictions and being wrong is the cost of doing business. But is it a larger sleight of hand? In “Predictions Without Futures,” Sun-Ha Hong wonders whether the idea that we are about to be able to make algorithmically correct predictions actually just “persists as an eternally deferred and recycled horizon.”⁴⁴ To hold out for the next technological improvement is to hold out for some breakthrough—and breakthroughs have been promised for decades. All the while:

the present and its teething problems are somewhat diluted of reality: there is less need to worry so much about concrete, existing patterns of inequality or inefficiency, the idea goes, since technological breakthroughs will soon render them irrelevant.⁴⁵

And the future gets closed off as a source for alternate collective imagination. This of what the promise of automation has done:

in organizations ranging from mid-twentieth-century machine shops and twenty-first-century self-driving carmakers to Cold War military research, the expectation of comprehensive automation in the near future drove seismic changes—changes that persisted regardless of the fact that total automation never arrived.⁴⁶

⁴² See Tim Hwang, *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet* (New York, Farrar, Straus and Giroux: 2020).

⁴³ Legal Singularity Book, *supra* note 41 at 107.

⁴⁴ Sun-Ha Hong, “Predictions Without Futures” (2022) 61:3 *History and Theory* 371 at 374.

⁴⁵ *Ibid.*

⁴⁶ *Ibid* at 378.

If we are to plan for a singularity today, does it not make it likely that “whatever these technologies cannot account for [ought to be] explicitly or implicitly cast aside as irrelevant at base and a threatening source of irrationality at worst?”⁴⁷

Is that wool over our eyes? Indeed, a curiosity. Survey the terrain, the most enthusiastic evangelists of the power of prediction are almost invariably associated with firms looking to raise capital to make predictions. A claim that a legal flow state is at hand; a confident assertion that the singularity is near; an admonition to plan for it; a warning not to doubt the technology: if the world is rearranged in anticipation of something, but that something does not come, was it worth rearranging the world for?

Meanwhile, the statisticians, mathematicians, and computer scientists lodged in out of the way university departments, those people who discovered and popularized the algorithmic ways of seeing the world, urge us to recognize a different predictive credo: all models are wrong, but some are useful.⁴⁸

V

In its day, Word2Vec was the world’s most impressive machine learning modelling technique for language. Developed by researchers at Google in 2013, the idea was simple.⁴⁹ If an algorithm was shown a massive amount of text in context, it could notice the relationships between words and intuit their meaning. After training, every word had its own mathematical representation (called a word embedding) that could be used to locate it next to other words. You could have fun playing around with these words: take the values for “Paris,” then subtract the values for “France” and add the values for “Italy” and you would end with a number extremely close to “Rome.” The approach was wildly exciting

⁴⁷ *Ibid* at 375.

⁴⁸ For a discussion on this concept, particularly as it relates to the digital humanities, see Richard Jean So, “All models are wrong” (2017) 132:3 PMLA 668.

⁴⁹ Tomas Mikolov, “Efficient estimation of word representations in vector space” (2013) arXiv, 1301.378.

for artificial intelligence researchers. Indeed, as of the time of writing, the original Word2Vec paper has been cited over 43,000 times.⁵⁰

Was it then strange that it took until 2015 for anyone to notice something obvious? Brian Christian tells the story of when PhD student Tolga Bolukbasi met his supervisor for Friday happy hour drinks. As computer scientists are wont to do, they brought their laptops. One of them spun up Word2Vec and started playing around: “Doctor” – “Man” + “Woman” = “Nurse.” They tried another: “Shopkeeper” – “Man” + “Woman” = “Housewife.” Finally: “Computer programmer” – “Man” + “Woman” = “Homemaker.”⁵¹ Within a year they published an article decrying the fact that the computer science community had missed how “blatantly sexist the embeddings are.”⁵²

Around the same time, ProPublica published what must be one of the most consequential pieces of writing on machine learning. In *Machine Bias*, journalists Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner explained how they analyzed an algorithm used to predict recidivism amongst prisoners and found that it “proved remarkably unreliable in forecasting violent crime” and would “falsely flag black defendants as future criminals, wrongly labeling them this way at almost twice the rate as white defendants.”⁵³

In what ways were and are these models wrong? Quickly the machine learning and artificial intelligence research communities organized into camps that had different answers to this question. For researchers like Ruha Benjamin,⁵⁴ Timnit Gebru,⁵⁵ Cathy

⁵⁰ This is according to statistics presented on “Google Scholar,” online: <https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&q=Efficient+Estimation+of+Word+Representations+in+Vector+Space&btnG=>. (accessed 22 August 2024)

⁵¹ See “Introduction” in Brian Christian, *The Alignment Problem: Machine Learning and Human Values* (New York: WM Norton Company, 2020).

⁵² Tolga Bolukbasi et al, “Man is to computer programmer as woman is to homemaker? Debiasing word embeddings” (2016) 29 *Advances in Neural Information Processing Systems* 29.

⁵³ Julia Angwin et al, “Machine Bias: There’s software used across the country to predict future criminals. And it’s biased against blacks” (23 May 2016), online: <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>>.

⁵⁴ Ruha Benjamin, *Race After Technology: Abolitionist Tools for the New Jim Code* (Cambridge: Polity Books, 2019). [Race After Technology]

⁵⁵ Karen Hao, “We read the paper that forced Timnit Gebru out of Google. Here’s what it says.” (4 December 2020), online:

O’Neil,⁵⁶ Meredith Broussard,⁵⁷ Joy Buolamwini,⁵⁸ and Safiya Umoja Noble,⁵⁹ a core problem is that some firms believe (or want to believe) that algorithms can transcend or bypass the inequitable and prejudicial circumstances of their production. Ruha Benjamin, for example, powerfully argues that we are watching the emergence of a “New Jim Code.” This is a reality marked by “the employment of new technologies that reflect and reproduce existing inequities but that are promoted and perceived as more objective or progressive than the discriminatory systems of a previous era.”⁶⁰

In a similar vein, Timnit Gebru warns that scientists’ self-conception as “‘objective’ clouds them from being self-critical and analyzing what predominant discriminatory view of the day they could be encoding, or what goal they are helping advance.”⁶¹ In a sense, these scholars advance an argument about hubris: surely history has not been waiting for today’s computer scientists to arrive, given that they possess the unique skillset required to develop decision making processes to sidestep prejudice and address unjust distributions of power.

Alarie and Aidid agree, to a point. When firms and computer scientists design algorithms uncritically, they run the risk of relying on incomplete data, where “in the process of representing real-world events, some context is inevitably lost.”⁶² For example, an algorithm that is given apparently race-neutral data about a racist system, “risks projecting those inequities into the future.”⁶³ For this reason, it is incumbent upon

<<https://www.technologyreview.com/2020/12/04/1013294/google-ai-ethics-research-paper-forced-out-timnit-gebru/>>.

⁵⁶ Cathy O’Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Crown, 2016).

⁵⁷ Meredith Broussard, *Artificial Unintelligence: How Computers Misunderstand the World* (Cambridge, MIT Press: 2018).

⁵⁸ Joy Buolamwini, *Unmasking AI: My Mission to Protect What Is Human in a World of Machines* (New York, Random House, 2023).

⁵⁹ Safiya Umoja Noble, *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: New York University Press, 2018).

⁶⁰ Race After Technology, *supra* note 54 at 6.

⁶¹ Timnit Gebru, “Race and Gender” in Markus Dubber, Frank Pasquale & Sunit Das, eds, *The Oxford Handbook of Ethics of AI* (Oxford: Oxford University Press, 2020) 253 at 253.

⁶² Legal Singularity Book, *supra* note 41 at 188.

⁶³ *Ibid* at 190.

designers to account for prejudice, so that their algorithms can be developed without uncritical reliance on incomplete data. But, they suggest we ought not focus all of our attention on algorithm designers because they “worry about taking society ‘off the hook,’ so to speak, for issues that predate algorithmic technologies.”⁶⁴ For this reason, they say that:

The path to equitable outcomes along categories like race and gender involves the collective and concerted heavy-lifting of reducing systemic barriers, critically interrogating biases (both conscious and unconscious), and making substantial social, political, and economic investments. If – and only if – these important measures are also being taken, then a concurrent program of technological safeguards and interventions will likely yield more equitable results.⁶⁵

This is a curious and revealing paragraph. Despite their confidence that the legal singularity is at hand, here they acknowledge that absent a massive reordering of society, the singularity would lead to horrifying results. If a just singularity requires “heavy-lifting” and “important measures” as a precondition, ought not defining and securing those be the first order of business? And, we might also wonder, might all this heavy-lifting in turn produce a society that no longer needs, desires, or could produce a singularity?

At bottom, the project rests on two extraordinary leaps of faith: trust that the science will work and trust that society will iron out its inequities before the singularity arrives. Commerce expects nothing less. In 2021, a press release announced that Blue J Legal had secured \$9 million in series B funding. “This is a big win for the clients of law and accounting firms,” Alarie is quoted as saying, later adding “[w]ith the ongoing support and confidence of our investors, we at Blue J continue to serve our customers by expanding our ability to bring absolute clarity to the law, everywhere and on-demand.”⁶⁶

VI

⁶⁴ *Ibid* at 195.

⁶⁵ *Ibid* at 196.

⁶⁶ “Blue J Raises \$9M in Series B Funding, Adds Key Personnel, Announces New Tax Diagramming Solution” *businesswire* (10 August 2021), online: <<https://www.businesswire.com/news/home/20210810005199/en/Blue-J-Raises-9M-in-Series-B-Funding-Adds-Key-Personnel-Announces-New-Tax-Diagramming-Solution>>.

In 1976, Raymond Williams published his groundbreaking text, *Keywords: A Vocabulary of Culture and Society*. In it, he explored how the meanings of words shifted or changed as social relations changed or social contests sharpened or faded. Our language, he reasoned, is:

not a *tradition* to be learned, nor a *consensus* to be accepted, nor a set of meanings which, because it is ‘our language’ has a natural authority; but as a shaping and reshaping, in real circumstances and from profoundly different and important points of view: a vocabulary to use, to find our own ways in, to change as we find it necessary to change it, as we go on making our own language and history.⁶⁷

“Doctor” – “Man” + “Woman” = “Nurse.” Words and ideas: these exist in relation to each other, in relation to historical contests, in relation to current contests. Is it a bug that statistical models represent patterns of patriarchal thought in our language, culture, and society? Or might it be a model of language that could, yes, be useful?

Holmes and Alarie are both wrong: history—those evolutionary, forward and backward, changes in our relations with each other—will not end. There is no infinite, semi-divine law to be had because, just as humanity will not be completed, neither will our laws. But there is something in Holmes’ thinking that ought to be recommended to us again. For all the forward-looking self-confidence of the modernists, Jameson reminds us that their thought remained devilled by anxiety and pathos. God was dead, but the modernists remained uncertain and upset by what that meant.⁶⁸ We, or at least I, am therefore unfair to stress the idea that Holmes’ vision was unwaveringly teleological: the future was not at hand, but it was worth working towards. Unlike Alarie, who tells us that the singularity is at hand, Holmes only promised access to glimpses and echoes. And, pace Kornbluh, it was a project thick with mediation.

The singularity is a project of commerce; Holmes’ infinite law was a project for the new social scientists. The singularity offers absolute truth immediately; infinite law promises answers based on intense human consideration and evaluation. The singularity

⁶⁷ Raymond Williams, *Keywords: A Vocabulary of Culture and Society* (New York: Oxford University Press, 1983) at 24.

⁶⁸ Fredric Jameson, “The Aesthetics of Singularity” (2015) 92 *New Left Review* 101 at 125.

gives all access to an unchanging truth; Holmes urged us to study those dragons of the generations that preceded us to see their purpose.

Both do err in the same way. They make the classic mistake: to assume that their subject is the subject. From the perspective of law schools, courtrooms, and lawyers' offices, the problems of law just look so clear. It is a problem how long it takes people to get adjudication, it is a problem that legal texts are hard to understand, it is a problem that different adjudicators reach different conclusions regarding similar cases, it is a problem that lawyers are so expensive. But to solve these problems (and, yes, many of these problems ought to be solved) is not to solve law. Indeed, some things that lawyers conceive of as problems may, in actual fact, serve some other important social purpose. Sometimes lengthy trial delays might be a good thing. When I was a refugee lawyer, for example, I noticed that my clients' cases only seemed to get stronger, more well-supported, and clearer with time.

It is exciting to imagine how new methods can transform one's subject, but excitement is not a case against humility. Law is a pretentious subject and can encourage pretension among its members. But it seems to me that empirical discoveries about law tend not to raise law's importance, but diminish it. If culture is a text, then law is a cultural text. And law's vocabulary is, in its own way, an index into the conflicts, relations, and thoughts that shape our society. So long as our society remains out of balance, history will march on. New conflict and new energies will be metabolized by culture and social institutions into new, adapted, and more complex texts. Volcanoes will erupt, pathogens will multiply, inequity will deny many rich and full lives, new ideas will proliferate and make new ways of living possible. Conflict, tragedy, and triumph are not about to go away—and each will, in some way, be encoded in culture and text. And now there are new ways to see that code which, once seen, maybe we can do something with. It is always worth knowing what that dragon is up to so we may all think more clearly about its function. With modernity's passing, we now know that we will never find, tame, or kill all those dragons, and we know that new ones are always being born, but this ought not stop the trying.

So, Holmes was right about something: law is too important, about too much, contains too much history and contradiction, to be primarily about that lawyerly business of prediction.

VII

But this is a dissertation, and it is time to end. So how do I appraise my own efforts? The confusing thing about computational approaches to the study of law is that they are confused. While there are many legal scholars thinking about what law should say about artificial intelligence, there are comparatively few using computers to say something about law. Yet for those that do, those two paths of computational law are a clear and present concern: can computers help us make better predictions about law, making a business-like and commercial contribution to the practice of law? Or can computational approaches ground new critical inquiries about law's shape, changes, and drifts?

In this dissertation's introduction, I argued that legal scholars could draw inspiration from a generation of digital humanists who set out to describe—in different times and in different places—the shape of creative production, in turn revealing some deep truths about the relationship between the material and the cultural. Bourdieu's field of law is embedded in a larger field of power, which in turn orders the dialogues and conflicts between fields. Some one hundred thousand words in, I am more convinced than ever of the general truth of this proposition. The legal realists of the 1920s and the 1930s believed that the new social sciences, psychoanalysis, statistics, and a broad view of legality would give new perspectives on law's function and form. Since then, critical legal scholars, new legal empiricists, pluralists, law and society scholars, and sociologists of law have each invited us to see how much larger, contradictory, and nebulous legality is than even the realists imagined. New computational methods promise to help us develop even richer and more comprehensive grasps of legality because, unlike the generations that preceded ours, we now have new ways of reckoning with scale.

I am mindful, however, that there is a dissonance between the studies undertaken here and the largeness of the theoretical vision that guides me. Yes, we now have convincing proof that judges who are men speak more than judges who are women. Yes,

we now have convincing proof that in this officially bilingual country, English is the language of appellate law. Yes, we now have convincing proof that the law of terrorism is unevenly applied. And, yes, this proof derives from new computational methods that were not available to us just a few years ago. But I readily acknowledge that the questions and answers do not come from a whole new cloth and that they speak in the same tenor of preexisting empirical research agendas. You could fairly say that all this talk of sociology and French theory and deep tectonic change and culture and materiality does not much assist: these are focused empirical questions, answered by new focused methodologies. Maybe the true use of these sorts of studies, and this sort of methodology, is strictly practical. Administrators, policy experts, lawyers, voters: each wants to know how the law is applied to ensure its better administration and computational tools could help with that project. As Jenna McGill and Amy Salyzyn observe, the practice of law and the practice of legal institutions will be changed (for better and potentially also for worse) by new computational technologies of analysis.⁶⁹ So, yes, institutionalists ought to be excited by critical computational methodologies because they will allow for more hands-on management and evaluation of institutions.

But there is more here too. The sum of empirical knowledge exceeds the total of its parts. For those of us who want to understand the habitus of law and better understand its function, the only way in is incremental, evidence-driven exploration. Once we know with better precision who talks more, we can ask secondary empirical questions about why they talk more. Once we know that some classes of terrorism cases are more likely to be decided differently, we can ask more probing questions about the variables that lead to uneven decision making. Second steps must follow first steps, and the state of computational legal studies is currently such that this is a time of first steps. It says something that every computationally informed legal scholar in Canada that I know of self-taught themselves how to code. There are not yet robust curricula, readily developed

⁶⁹ Jenna McGill and Amy Salyzyn, “Judging by the Numbers: Judicial Analytics, the Justice System Judging by the Numbers: Judicial Analytics, the Justice System and its Stakeholders” (2021) 44:1 Dal LJ 249.

methods, and easily accessible categories of data that would allow for a more general flourishing in the field.

The pressing next step for this scholarly agenda is volume and scale: we must ask more of our methods, we must ask more questions, we must explore more areas, and we must work with more data. Put differently, to live up to the ambition of the distant readers (to see if big data could “lead us back to big questions”⁷⁰), we must scale up these methods of scale. And in the most general sense, that is my plan now. But I also hear those notes of caution sounded by skeptical scholars. The plain reality is that, to date, the most impressive big data achievements are neither scholarly nor critical. Ours is an age of surveillance, and if borders are becoming more effective, if corporations are getting better at shaping our thought, and if police and militaries are better able to track groups that know oppression, it is in no small part due to the power of big data analysis.⁷¹ It gives me pause, for example, to now know that a lawyer with a laptop (and here I refer to myself) can now code a program that can recognize faces and voices and transcribe their utterances.

This genie will not go back in the bottle and future critical scholarship on the worst manifestations and uses of big data analysis will require scholars to become fluent with computational methodologies. While this means that we (and I) must redouble our learning to better understand that these methodologies can do, and therefore do and may do, we must also guard against unwittingly becoming the research centres for big surveillance, or a critical veneer that is used to warrant larger and problematic uses of computational methods. Equally, however, we must explore how computational tools can enable what Özgün Topak refers to as human rights surveillance. This form of countersurveillance cuts against the power of the powerful by using technology in “oppositional ways,” to challenge inequity and injustice.⁷²

⁷⁰ Franco Moretti, *Literature, Measured*, Pamphlets of the Stanford Literary Lab (Stanford, 2016) at 7.

⁷¹ See, for example, Petra Molnar, “Technology on the margins: AI and global migration management from a human rights perspective” (2019) 8:2 Cambridge International Law Journal 8.2 305.

⁷² Özgün Topak, “Humanitarian and Human Rights Surveillance: The Challenge to Border Surveillance and Invisibility?” (2019) 17:3/4 Surveillance and Society 382 at 388.

At bottom, this is my overall object: to critically redescribe the world in the open, in ways that can be useful. The artefacts produced by this dissertation that I am most proud of are its maps. I am by turns both excited by what these maps suggest and frustrated by the difficulty I have deciphering them. Against those traditional ways of teaching and thinking about law, the maps authorize decidedly inductive, evidence-first, modes of theory building. Almost a century ago, Roscoe Pound chastised the young realists for forgetting the “significance of generalizations and conceptions as instruments towards the ends of the legal order.”⁷³ We should not, he said, pretend that the law was universal in the way that, for example, Langdell conceived it, but it is an error to believe that the resolution of every case is governed exclusively by the idiosyncratic coming together of idiosyncratic variables. Pound would, I think, look at my maps and nod: huge masses of text, the grouping together of many single-instances, and the presentation of repeating mental patterns. We see law, not as it is arranged in the casebook or textbook, not as an idealized formation, but as some material force that brings people to some places, gets them to say some things, and invites others to interpret their experiences in some repeating ways.

The challenge for the empiricist is tracking the change and movement of these forms. My maps offer snapshots, but (so far) only snapshots. And here also lies frustration. First, because of my limited access to data, each of my studies was limited to just over a decade or two. The evolution of legal forms, I now think, requires a wider optic to be properly grasped (or, alternatively and perhaps likely, we may look back and see that this period of Canadian legal history was marked by comparative stasis). Second, maps need to be put in conversation with each other. In each study here, I focused on a single domain: the Parliament and the Executive that issued federal law, the Federal Court of Canada deciding just refugee law judicial reviews, and the Social Security Tribunal of Canada just deciding employment insurance matters. To develop a richer analysis of law and its habits, more domains need to be brought together: how are decisions shaped by the language used by lawyers in their arguments? How is the Federal Court’s refugee law jurisprudence

⁷³ Roscoe Pound, “The Call for a Realist Jurisprudence” (1931) 44:5 Harv LR 697 at 710.

related to the reasoning of the refugee tribunal below? How do changes in one institution impact others? These are the complex questions that the maps invite us to ask, in turn inviting us to develop more complex maps.

This is the direction I am most excited to explore because it promises to unlock new categories of critical insight. While the main impediment is access to new stores of data, with new data and more ambitious projects it will become possible to map more than just the development of legal forms in one tribunal, but the relationships between all sorts of different forms and grammars—some of which never find direct expression in the language of decision making. Think back to this dissertation’s preface. My original motivation for returning to graduate school was to empirically explore what struck me as an obvious racialized dynamic of law. To borrow from Holmes, I believe that I am starting to catch glimpses of how my methods of scale might be deployed to surface and describe some of those subterranean patterns, buried deep in law’s forms. Now the maps show how some forms on the surface get invoked, develop, and exist in relation: soon, maybe, we can use more advanced methods to better uncover deeper patterns in law.

So, two future directions. First: it is time for more studies. Second: it is time for more maps, but much more complex maps. Maps that integrate different institutions and bodies of text, different visions and different vantages, and different experiences and different backgrounds. Then the larger promise of reading from a distance might be realized and we will have some new ways of charting those habits of mind that show how we relate to each other.

8. Bibliography

Legislation

Access to Information Act, RSC 1985, c A-1.

Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment, UNGA, 39th Sess, UN Doc A/RES/39/51 (1984) GA Res 39/46

Criminal Code, RSC 1985, c C-46.

Employment Insurance Act, SC 1996, c 23.

Employment Insurance Regulations, SOR/96-332.

Immigration Appeal Board Act, SC 1966, c 90.

Immigration Division Rules, SOR/2002-229.

Immigration and Refugee Protection Act, SC 2001, c 27.

Legislation Revision and Consolidation Act, RSC 1985, c S-20.

Fees for Court Transcripts, O Reg 145/22

Official Languages Act, RSC 1985, c 31.

Reconsideration Request Regulations, SOR/2013-63.

Refugee Protection Division Rules, SOR/2012-256.

Rules of the Supreme Court of Canada, SOR/2002-156.

The Constitution Act, 1867, 30 & 31 Vict, c 3.

Jurisprudence

Aedo Arancibia v. Canada (Citizenship and Immigration), 2023 FC 1245.

Afanasyev v Canada (Citizenship and Immigration), 2012 FC 1270.

Agraira v Canada (Public Safety and Emergency Preparedness), 2013 SCC 36.

Ahmed v. Canada (Citizenship and Immigration), 2023 FC 830.

AK v Canada (Citizenship and Immigration), 2018 FC 236.

Alam v Canada (Citizenship and Immigration), 2018 FC 922.

Algonquin College of Applied Arts and Technology (Re), 2022 CanLII 122328 (ON IPC).

AM v Canada Employment Insurance Commission, 2022 SST 737.

AO v Canada Employment Insurance Commission, 2020 SST 770.

AP v Canada Employment Insurance Commission, 2022 SST 1410.

Appellant v Northern Regional Health Authority, 2023 CanLII 56246 (MB HAB).

Arumaiturai v. Canada (Citizenship and Immigration), 2022 FC 604.

Attorney General of Canada v Attorney General of Ontario et al, [1937] UKPC 7 (28 January 1937).

Bains v Canada (Minister of Employment and Immigration), [1990] 47 Admin LR 317, 109 NR 239 (FCA).

Banom v. Canada (Citizenship and Immigration), 2016 FC 502.

Bautista Montero v. Canada (Citizenship and Immigration), 2021 FC 1432.

BH v Canada Employment Insurance Commission, 2018 SST 101.

Blois v. Butcher, 2023 NSSM 26.

BR v Canada Employment Insurance Commission, 2021 SST 463.

Bruzzese v Canada (Public Safety and Emergency Preparedness), 2016 FC 1119.

Bushra v. Canada (Citizenship and Immigration), 2018 FC 896.

Buwu v. Canada (Citizenship and Immigration), 2013 FC 850.

Canada (Attorney General) v Cyrenne, 2010 FCA 349.

Canada (Attorney General) v. Granstrom, 2003 FCA 485.

Canada (Attorney General) v. Ward, 1993 CanLII 105 (SCC), [1993] 2 SCR 689.

Canada (Attorney General) v. White, 2011 FCA 190.

Canada (Citizenship and Immigration) v Harkat, 2014 SCC 37.

Canada (Citizenship and Immigration) v Tennant, 2018 FCA 132.

Canada (Minister of Citizenship and Immigration) v. Singh, 1998 CanLII 8281 (FC).

Canada (Minister of Citizenship and Immigration) v. Vavilov, 2019 SCC 65 (CanLII), [2019] 4 SCR 653.

CB v Canada Employment Insurance Commission, 2022 SST 1505.

CE v Canada Employment Insurance Commission, 2023 SST 634.
Celaj v. Canada (Citizenship and Immigration), 2014 FC 761.
Chapagain v. Canada (Citizenship and Immigration), 2023 FC 949.
Charkaoui v Canada (Citizenship and Immigration), 2007 SCC 9.
Chen v. Canada (Citizenship and Immigration), 2013 FC 928.
Chowdhury v Canada (Citizenship and Immigration), 2017 FC 189.
Ciro v. Canada (Citizenship and Immigration), 2021 FC 1383.
CM v Canada Employment Insurance Commission, 2022 SST 1652.
Daniel v. Canada (Citizenship and Immigration), 2016 FC 1049.
DH v Canada Employment Insurance Commission, 2021 SST 338.
Diallo v. Canada (Citizenship and Immigration), 2014 FC 471.
Djubok v. Canada (Citizenship and Immigration), 2014 FC 497.
Dleiw v Canada (Citizenship and Immigration), 2020 FC 59.
Domtar Inc v Quebec (Commission d'appel en matière de lésions professionnelles), [1993] 2 SCR 756 (SCC).
El (Re), 2015 CanLII 94649 (ON CCB).
El Werfalli v Canada (Public Safety and Emergency Preparedness), 2013 FC 612.
ES v Canada Employment Insurance Commission, 2020 SST 598.
Etik v. Canada (Citizenship and Immigration), 2018 FC 175.
Ferdous v Canada (Public Safety and Emergency Preparedness), 2019 FC 1115.
Foisal v. Canada (Citizenship and Immigration), 2021 FC 404.
GA v Canada Employment Insurance Commission, 2021 SST 677.
Gazi v Canada (Citizenship and Immigration), 2017 FC 94.
GC v. Canada Employment Insurance Commission, 2019 SST 1523.
GC v. Canada Employment Insurance Commission, 2019 SST 596.
Glenczi v. Canada (Citizenship and Immigration), 2019 FC 931.
Gonzalez de los Santos v. Canada (Citizenship and Immigration), 2022 FC 1281.
Gupta v. The King, 2023 TCC 82.
Haben Abrham Weldekidan v. His Majesty the King, 2023 CanLII 57199 (SCC).
Huang v. Canada (Citizenship and Immigration), 2019 FC 120.
Huang v. Canada (Citizenship and Immigration), 2019 FC 1366.
Hui v. Canada (Citizenship and Immigration), 2020 FC 446.

Intisar v Canada (Citizenship and Immigration), 2018 FC 1128.
IR v Canada Employment Insurance Commission, 2021 SST 188.
Islam v Canada (Public Safety and Emergency Preparedness), 2019 FC 912.
Islam v Canada (Public Safety and Emergency Preparedness), 2021 FC 108.
J.M.A. v. L.A.M., 2023 MBKB 110.
JB v. Canada Employment Insurance Commission, 2019 SST 595.
JD v. Canada Employment Insurance Commission, 2019 SST 599.
Jeyaredsagathas v. Canada (Citizenship and Immigration), 2018 FC 1238.
JK v. Canada Employment Insurance Commission, 2014 SSTGDEI 42.
JM v Canada Employment Insurance Commission, 2020 SST 345.
JM v Canada Employment Insurance Commission, 2023 SST 1099.
JN v Canada Employment Insurance Commission, 2021 SST 716.
JP v Canada Employment Insurance Commission, 2021 SST 320.
JR v Canada Employment Insurance Commission, 2021 SST 643.
Kamal v Canada (Immigration, Refugees and Citizenship), 2018 FC 480.
Kawekwune v. Canada (Citizenship and Immigration), 2020 FC 333.
KB v Canada Employment Insurance Commission, 2020 SST 683.
KB v Canada Employment Insurance Commission, 2020 SST 700.
Khan v Canada (Citizenship and Immigration), 2019 FC 899.
Khan v. Canada (Citizenship and Immigration), 2020 FC 438.
KM v Canada Employment Insurance Commission, 2022 SST 421.
Kouassi v. Canada (Citizenship and Immigration), 2016 FC 1398.
Layzell v. Ontario (Human Rights Commission), 2003 CanLII 23954 (ON SCDC).
Lepretre v. Canada (Attorney General), 2011 FCA 30.
Lhamo v. Canada (Citizenship and Immigration), 2016 FC 873.
Liang v. Canada (Citizenship and Immigration), 2019 FC 58.
Liu v. Canada (Citizenship and Immigration), 2016 FC 811.
LM v Canada Employment Insurance Commission, 2021 SST 932.
Madava v. Canada (Citizenship and Immigration), 2017 FC 149.
Mahmoud v. Canada (Citizenship and Immigration), 2016 FC 1020.
Mao v. Canada (Citizenship and Immigration), 2020 FC 542.

Mason v Canada (Citizenship and Immigration), 2023 SCC 21.

McNally Contractors (2011) Ltd. v. Director, Alberta Environment and Parks, 2023 ABEAB 6.

Meshveliani v. Canada (Citizenship and Immigration), 2019 FC 1351.

MG v. Canada Employment Insurance Commission, 2015 SSTGDEI 28.

MH v Canada Employment Insurance Commission, 2023 SST 559.

Miah v Canada (Public Safety and Emergency Preparedness), 2020 FC 38.

Mishel v. Canada (Citizenship and Immigration), 2015 FC 226.

Mishibinijima v. Canada (Attorney General), 2007 FCA 36.

MM v. Canada Employment Insurance Commission, 2017 SSTGDEI 49.

MM v. Canada Employment Insurance Commission, 2019 SST 235.

MM v. Canada Employment Insurance Commission, 2017 SSTGDEI 49.

MN v Canada (Public Safety and Emergency Preparedness), 2019 FC 796.

Mugesera v Canada (Minister of Citizenship and Immigration), 2005 SCC 40.

MV v Canada Employment Insurance Commission, 2023 SST 867.

Najafi v Canada (Public Safety and Emergency Preparedness), 2014 FCA 262.

Ni v. Canada (Citizenship and Immigration), 2018 FC 948.

NP v Canada Employment Insurance Commission, 2021 SST 754.

Nyembua v. Canada (Citizenship and Immigration), 2015 FC 970.

OA v Canada Employment Insurance Commission, 2020 SST 53.

OC v. Canada Employment Insurance Commission, 2015 SSTGDEI 46.

Ohwofasa v. Canada (Citizenship and Immigration), 2020 FC 266.

Okuyemi v. Canada (Citizenship and Immigration), 2023 FC 808.

Oremade v Canada (Minister of Citizenship and Immigration), 2005 FC 1077.

PB v Canada Employment Insurance Commission, 2023 SST 648.

Peer v Canada, 2010 FC 752.

Peer v Canada, 2011 FCA 91.

Pestova v. Canada (Citizenship and Immigration), 2016 FC 1024.

PP v Canada Employment Insurance Commission, 2023 SST 462.

Pulido v. Canada (Citizenship and Immigration), 2023 FC 463.

Qaddafi v. Canada (Citizenship and Immigration), 2016 FC 629.

Rahman v Canada (Public Safety and Emergency Preparedness), 2019 FC 807.

Ramrattan v. Canada (Citizenship and Immigration), 2016 FC 1108.

Rana v Canada (Public Safety and Emergency Preparedness), 2018 FC 1080.

RB v Canada Employment Insurance Commission, 2020 SST 206.

RS v. Canada Employment Insurance Commission, 2018 SST 1169.

Ruiz Lopez v. Canada (Citizenship and Immigration), 2021 FC 390.

S.A. v. Canada (Citizenship and Immigration), 2014 FC 146.

SA v Canada (Public Safety and Emergency Preparedness), 2017 FC 494.

Saleheen v Canada (Public Safety and Emergency Preparedness), 2019 FC 145.

Shakil Ali v. Canada (Citizenship and Immigration), 2023 FC 156.

Shandi (Re), [1991] FCJ no 1319, 17 Imm LR (2d) 54.

Shen v. Canada (Citizenship and Immigration), 2017 FC 885.

Shi v. Canada (Citizenship and Immigration), 2022 FC 196.

Singh v. Canada (Citizenship and Immigration), 2019 FC 1375.

Singh v. Canada (Citizenship and Immigration), 2023 FC 1106.

Sivagnanasundarampillai v. Canada (Citizenship and Immigration), 2018 FC 1109.

SL v Canada Employment Insurance Commission, 2016 SSTGDEI 97.

SL v Canada Employment Insurance Commission, 2020 SST 445.

SM v Canada Employment Insurance Commission, 2019 SST 598.

SM v. Canada Employment Insurance Commission, 2019 SST 597.

SN v Canada Employment Insurance Commission, 2018 SST 1182.

SP v Canada Employment Insurance Commission, 2022 SST 1323.

SR v Canada (Public Safety and Emergency Preparedness), 2019 FC 1118.

SR v. Canada Employment Insurance Commission, 2018 SST 307.

Sriratham v. Canada (Citizenship and Immigration), 2014 FC 614.

Suleyman v Canada (Minister of Citizenship and Immigration), 2008 FC 780.

Suresh v Canada (Minister of Citizenship and Immigration), 2002 SCC 1.

SV v Canada Employment Insurance Commission, 2019 SST 482.

Tamayo v. Canada (Citizenship and Immigration), 2015 FC 1217.

Tandi v. Canada (Citizenship and Immigration), 2021 FC 1413.

TC v. Canada Employment Insurance Commission, 2018 SST 1185.

The Highlands Residential Association v. Welland (City), 2021 CanLII 48340 (ON LPAT).

Tkachuk v. Canada (Citizenship and Immigration), 2015 FC 672.

Toronto (City) v CUPE, Local 79, 2003 SCC 63.

Toronto Coalition to Stop the War v Canada (Public Safety and Emergency Preparedness), 2010 FC 957.

TP v Canada Employment Insurance Commission, 2023 SST 1217.

Trawally v. Canada (Citizenship and Immigration), 2023 FC 1627.

Turoczi v. Canada (Citizenship and Immigration), 2012 FC 1423.

Ui Haq v. Canada (Citizenship and Immigration), 2022 FC 95.

Vemb v. The Samuel T, 1971 CanLII 1976 (CA EXC).

Weldemariam v Canada (Public Safety and Emergency Preparedness), 2020 FC 631.

WM v. Canada Employment Insurance Commission, 2018 SST 328.

Xu v. Canada (Citizenship and Immigration), 2014 FC 1062.

Yang v. Canada (Citizenship and Immigration), 2016 FC 543.

Zheng v. Canada (Citizenship and Immigration), 2015 FC 16.

ZZ v Canada Employment Insurance Commission, 2021 SST 610.

Secondary material

Monographs and collections

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

Anderson, Perry. *The Origins of Postmodernity* (London, Verso: 1998).

Armstrong, Nancy. *Desire and Domestic Fiction: A Political History of the Novel* (Oxford: Oxford University Press, 1987).

Bench-Capon, TJM, ed. *Knowledge-Based Systems and Legal Applications* (San Diego: Academic Press, 1991).

Benjamin, Ruha. *Race After Technology: Abolitionist Tools for the New Jim Code* (Cambridge: Polity Books, 2019).

Black, Ryan C, Johnson, Timothy R & Wedeking, Justin. *Oral Arguments and Coalition Formation on the U.S. Supreme Court: A Deliberate Dialogue* (University of Michigan Press, 2012).

Bourdieu, Pierre. *Outline of a Theory of Practice* (Cambridge: Cambridge University Press, 1977).

- Brooks, Cleanth. *The Well Wrought Urn: Studies in the Structure of Poetry* (New York: Harcourt, Brace, and World, 1947).
- Broussard, Meredith. *Artificial Unintelligence: How Computers Misunderstand the World* (Cambridge, MIT Press: 2018).
- Buolamwini, Joy. *Unmasking AI: My Mission to Protect What Is Human in a World of Machines* (New York, Random House, 2023).
- Campeau, Georges. *From UI to EI: Waging a War on the Welfare State* (Vancouver: UBC Press, 2004).
- Christian, Brian. *The Alignment Problem: Machine Learning and Human Values* (New York: WM Norton Company, 2020).
- Cohen, Margaret. *The Sentimental Education of the Novel* (Princeton: Princeton University Press, 2002).
- Dworkin, Ronald. *Law's Empire* (Cambridge: Belknap Press, 1997).
- Eagleton, Terry. *Literary Theory: An Introduction* (Minneapolis: University of Minnesota Press, 2008).
- Eve, Martin Paul. *The Digital Humanities and Literary Studies* (Oxford: Oxford University Press, 2022).
- Finkle, Alvin. *Social Policy and Practice in Canada: A History* (Waterloo: Wilfred Laurier Press, 2006).
- Grant, Angus. *Confronting (In)Security: Forging Legitimate Approaches to Security and Exclusion in Migration Law* (PhD dissertation, Osgoode Hall Law School, 2016), online: <digitalcommons.osgoode.yorku.ca/phd/24/> [perma.cc/7RG7-YAKF].
- Grimmer, Justin, Roberts, Margaret & Stewart, Brandon. *Text as Data: A New Framework for Machine Learning and the Social Sciences* (Princeton: Princeton University Press, 2022).
- Hardt, Michael & Weeks, Kathi, eds. *The Jameson Reader* (Oxford: Blackwell, 2000).
- Harker, Richard, Cheleen Mahar & Chris Wilkes, eds. *An Introduction to the Work of Pierre Bourdieu: The Practice of Theory* (New York: Springer, 2016).
- Hidalgo, César. *Why Information Grows: The Evolution of Order, from Atoms to Economics* (New York: Basic Books, 2015).
- Holmes, Oliver Wendell, Jr. & Steven Alan Childress. *The Annotated Common Law* (New Orleans: Quid Pro Law Books, 2010).
- Hudon, Marie-Ève. *Bilingualism in the Federal Courts* (Library of Parliament, 2011).
- Hwang, Tim. *Subprime Attention Crisis: Advertising and the Time Bomb at the Heart of the Internet* (New York, Farrar, Straus and Giroux: 2020).
- Immigration Appeal Board. *Immigration Appeal Cases: Selected Judgments* (Ottawa, 1969).

- Jameson, Fredric. *Marxism and Form: Twentieth-Century Dialectical Theories of Literature* (Princeton: Princeton University Press, 1972).
- Jameson, Fredric. *The Political Unconscious: Narrative as a Socially Symbolic Act* (Ithaca: Cornell University Press, 1981).
- Jockers, Matthew. *Macroanalysis: Digital Methods and Literary History* (Urbana: University of Illinois Press, 2013).
- Johnson, Timothy R. *Oral Arguments and Decision Making on the United States Supreme Court* (SUNY Press, 2004).
- Kornbluh, Anna. *Immediacy: Or, The Style of Too Late Capitalism* (London: Verso, 2024).
- Koselleck, Reinhart. *Futures Past: On the Semantics of Historical Time* (New York, Columbia University Press: 2005).
- Kurzweil, Ray. *The Singularity Is Near: When Humans Transcend Biology* (New York: Penguin Books, 2006).
- Langdell, Christopher Columbus. *A Selection of Cases on the Law of Contracts* (Boston: Little, Brown and Company, 1871).
- Liew, Jamie Chai Yun & Donald Galloway. *Immigration Law*, 2nd ed. (Toronto: Irwin Law, 2015).
- Livermore, Michael & Daniel Rockmore, eds. *Law as Data: Computation, Text, and the Future of Legal Analysis* (Santa Fe: Santa Fe Institute, 2019).
- Léonard, André. *The Employment Insurance Program in Canada: How It Works*, Publication number: 2010-52-E (Ottawa: Library of Parliament, 2010).
- Macfarlane, Emmett. *Governing from the Bench: The Supreme Court of Canada and the Judicial Role* (UBC Press, 2013).
- Marx, Karl. *A Contribution to the Critique of Political Economy* (Progress Publishers: Moscow, 1977).
- McLevy, John. *Doing Computational Social Science: A Practical Introduction* (London: Sage, 2022).
- Medvedev, Pavel & Mikhail Bakhtin. *The Formal Method in Literary Scholarship: A Critical Introduction to Sociological Poetics* (Cambridge: Harvard University Press, 1985).
- Moretti, Franco. *Atlas of the European Novel, 1800-1900* (London: Verso, 1999).
- . *The Way of the World: The Bildungsroman in European Culture* (London: Verso, 2000).
- . *Graphs, Maps, Trees: Abstract Models for Literary History* (London: Verso, 2007).
- . *Distant Reading* (New York and London: Verso, 2013).

- . *Literature, Measured*, Pamphlets of the Stanford Literary Lab (Stanford, 2016).
- Noble, Safiya Umoja. *Algorithms of Oppression: How Search Engines Reinforce Racism* (New York: New York University Press, 2018).
- O’Neil, Cathy. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (New York: Crown, 2016).
- Piper, Andrew. *Can We Be Wrong? The Problem of Textual Evidence in a Time of Data* (Cambridge: Cambridge University Press, 2020).
- Piper, Andrew. *Enumerations* (Chicago: University of Chicago Press, 2019).
- Rawls, John. *A Theory of Justice* (Cambridge: Harvard University Press, 1971).
- Rogerson, John. *Old Testament Criticism in the Nineteenth Century: England and Germany* (Philadelphia: Fortress, 1985).
- Sheehy, Elizabeth A. *Defending Battered Women on Trial* (Vancouver, BC: UBC Press, 2014).
- Trotsky, Leon. *Literature and Revolution* (Chicago: Haymarket Books, 2005).
- Underwood, Ted. *Distant Horizons: Digital Evidence and Literary Change* (Chicago: The University of Chicago Press, 2019).
- Wasielewski, Amanda. *Computational Formalism: Art History and Machine Learning* (Cambridge: MIT Press, 2023).
- Wiggins, Chris & Jones, Matthew. *How Data Happened: A History from the Age of Reason to the Age of Algorithms*(Norton and Company, New York, 2023).
- Williams, Raymond. *Keywords: A Vocabulary of Culture and Society* (New York: Oxford University Press, 1983).
- . *The Long Revolution* (Peterborough: Broadview Press, 2001).
- Wood, Donna. The Seventy-Five Year Decline: How Government Expropriated Employment Insurance from Canadian Workers and Employers and Why This Matters, Mowat research publication 151 (Toronto: Mowat Centre for Policy Innovation, 2017),
online: https://tspace.library.utoronto.ca/bitstream/1807/99369/1/Wood_2017_The_Seventy_Five.pdf.
- Wrightsmann, Lawrence S. *Oral Arguments Before the Supreme Court: An Empirical Approach* (New York: Oxford University Press, 2008).

Scholarly articles and chapters

- Alarie, Benjamin. “The Path of Law: Towards Legal Singularity” (2016) 66:4 University of Toronto Law Journal 443.

- Allen, Laymen. “Some uses of Symbolic Logic in Legal Practice” (1962) 3:2 MULL: Modern Uses of Law in Logic 119.
- Bafna, Prafulla, Dhanya Pramod & Anagha Vaidya. “Document clustering: TF-IDF approach” (2016) 2016 International Conference on Electrical, Electronics, and Optimization Techniques 61.
- Baker, Paul & Tony McEnery. “A corpus-based approach to discourses of refugees and asylum seekers in UN and newspaper texts” (2005) 4:2 Journal of Language and Politics 197.
- Baldick, Chris. “Literature in the Academy” in Patricia Waugh, ed. *Literary Theory and Criticism: An Oxford Guide* (Oxford: Oxford University Press, 2006).
- Bedard-Rubin, Jean-Christophe & Rubin, Tiago. “Assessing the Impact of Unilingualism at the Supreme Court of Canada: Panel Composition, Assertiveness, Caseload, and Deference” (2018) 55 Osgoode Hall L J 713.
- . “The Elusive Quest for French on the Bench: Bilingualism Scores for Canadian Supreme Court Justices, 1985–2013” (2022) 37:2 CJLS / LRCDS 249.
- Bender, Emily M et al. “On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🐦” (2021) FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency 610.
- Bergam, Noah, Emily Allaway & Kathleen McKeown. “Legal and Political Stance Detection of SCOTUS Language” (2022) arXiv, 2211.11724.
- Bi, Qifang et al. “What is Machine Learning? A Primer for the Epidemiologist” (2019) 188:12 American Journal of Epidemiology 2222.
- Black, Ryan & James Spriggs. “An Empirical Analysis of the Length of U.S. Supreme Court Opinions” (2008) 45 Hous L Rev 622.
- Bolukbasi, Tolga et al. “Man is to computer programmer as woman is to homemaker? Debiasing word embeddings” (2016) 29 Advances in Neural Information Processing Systems 29.
- Bommarito, Michael J II & Daniel Martin Katz. “A Mathematical Approach to the Study of the United States Code” (2010) 389 Physica A: Statistical Mechanics and its Applications 4195.
- Bourdieu, Pierre. “The Force of Law: Toward a Sociology of the Juridical Field” (1987) 38 Hastings LJ 805.
- Bredin, Hervé et al. “pyannote.audio: Neural building blocks for speaker diarization” (2019) arXiv, 1911.01255.
- Brügger, Niels. “Digital humanities and web archives: Possible new paths for combining datasets” (2021) 2 International Journal of Digital Humanities 145.

- Burton, Dolores. "Review of Index Thomisticus: *Sancti Thomae Aquinatis operum omnium indices et concordantiae; Sancti Thomae Aquinatis opera omnia*" (1984) 18:2 Computers and the Humanities 109.
- Cain, William E. "British and American New Criticism" in David Richter, ed, A Companion to Literary Theory (Hoboken: John Wiley & Sons, 2018).
- Cao, Hongliu. "Recent advances in text embedding: A Comprehensive Review of Top-Performing Methods on the MTEB Benchmark" (2024) arXiv:2406.01607, online: <<https://arxiv.org/abs/2406.01607>>
- Chapman, Carl & Stolee, Kathryn T. "Exploring regular expression usage and context in Python" (2016) Proceedings of the 25th International Symposium on Software Testing and Analysis.
- Cohen, Felix. "Transcendental Nonsense and the Functional Approach" (1935) 35:6 Colom L Rev 809.
- Crenshaw, Kimberlé. "Demarginalizing the Intersection of Race and Sex: A Black Feminist Critique of Antidiscrimination Doctrine, Feminist Theory and Antiracist Politics" (1989) U Chicago Legal F 139.
- Dezalay, Yves & Mikael Rask Madsen. "The Force of Law and Lawyers: Pierre Bourdieu and the Reflexive Sociology of Law" (2012) 8:1 Annu Rev Law Soc Sci 433.
- Duvall, Michael. "When Is Oral Argument Important - A Judicial Clerk's View of the Debate" (2007) 9 J App Prac & Process 121.
- Eberendu, Adanma Cecilia. "Unstructured Data: an overview of the data of Big Data" (2016) 38:1 International Journal of Computer Trends and Technology 46.
- Edwards, Linda. "The Trouble with Categories: What Theory can Teach Us about the Doctrine-Skills Divide" (2014) 64 J Legal Educ 181.
- English, James F. "Everywhere and Nowhere: The Sociology of Literature After 'the Sociology of Literature'" (2010) 41:2 New Literary History v.
- . "Literary Studies" in Tony Bennett & John Frow, eds, *The SAGE Handbook of Cultural Analysis* (London: SAGE Publications, 2008) 126.
- Epstein, Lee, William M Landes & Richard A Posner. "Inferring the Winning Party in the Supreme Court from the Pattern of Questioning at Oral Argument" (2010) 39:2 The Journal of Legal Studies 433.
- Feldman, Adam & Rebecca D Gill. "Power Dynamics in Supreme Court Oral Arguments: The Relationship between Gender and Justice-to-Justice Interruptions" (2019) 40:3 Justice System Journal 173.
- Frankenreiter, Jens & Michal A Livermore. "Computational Methods in Legal Analysis" (2020) 16 Annual Rev L & Soc Science 39.

- Gebru, Timnit. "Race and Gender" in Markus Dubber, Frank Pasquale & Sunit Das, eds, *The Oxford Handbook of Ethics of AI* (Oxford: Oxford University Press, 2020).
- Goldfarb, Avi & Catherine Tucker. "Digital Economics" (2019) 57:1 *Journal of Economic Literature* 3.
- Gorman, David. "Russian Formalism" in Richter, David, ed, *A Companion to Literary Theory* (Hoboken: John Wiley & Sons, 2018) 36.
- Grammond, Sébastien and Mark Power. "Should Supreme Court Judges be Bilingual?" in Nadia Verrelli ed., *Democratic Dilemma: Reforming Canada's Supreme Court* (McGill-Queen's Press - MQUP, 2013).
- Green, Brandon & Antoine Pécoud. "Talking about Migration in Times of Crisis: A Textual Analysis of Narratives by IOM and UNHCR on Migrants and Refugees" (2023) *American Behavioral Scientist*.
- Guillory, John. "Literary Study and the Modern System of the Disciplines" in Amanda Anderson & Joseph Valente, eds, *Disciplinary at the Fin de Siècle* (Princeton: Princeton University Press, 2021).
- Hasson, Reuben. "Discipline and Punishment in the Law of Unemployment Insurance-- A Critical View of Disqualifications and Disentitlement" (1987) 25 *Osgoode Hall LJ* 615.
- Hasson, Reuben. "The Cruel War: Social Security Abuse in Canada" (1981) 3 *Can Tax'n* 3 114.
- Ho, Jeremiah. "Function, Form, and Strawberries: Subverting Langdell" (2015) 64:4 *Journal of Legal Education* 656.
- Holmes Jr., Oliver Wendell. "The Path of Law" (1897) 10:8 *Harv L Rev* 457.
- Hong, Sun-Ha. "Predictions Without Futures" (2022) 61:3 *History and Theory* 371.
- Horwitz, Morton J. "The Rise of Legal Formalism" (1975) 19:4 *Am J Leg His* 251.
- Howard, Jeremy and Sylvain Gugger. "Fastai: A Layered API for Deep Learning" (2020) 11:2 *Information* 108.
- Hulsebosch, Daniel. "The New Deal Court: Emergence of a New Reason" (1990) 90:7 *Colum L R* 1973.
- Huttner, Liane & Denis Merigoux. "Catala: moving towards the future of legal expert systems" (2022) 25 *AI & L* 1.
- Jacobi, Tonja & Matthew Sag. "Taking Laughter Seriously at the Supreme Court" (2019) 72 *V and L Rev* 1423.
- Jacomy, Mathieu et al. "ForceAtlas2, a continuous graph layout algorithm for handy network visualization designed for the Gephi software." (2014) 9:6 *PloS one*.
- Jameson, Fredric. "The Aesthetics of Singularity" (2015) 92 *New Left Review* 101.

- Ji, Ziwei et al. "Survey of Hallucination in Natural Language Generation" (2023) 55:12 ACM Comput Surv 248.
- Johnson, Timothy R., Paul J Wahlbeck & James F Spriggs. "The Influence of Oral Arguments on the U.S. Supreme Court" (2006) 100:1 American Political Science Review 99.
- Katz, Daniel et al. "Complex societies and the growth of the law" (2020) 10 Scientific Reports.
- Kelling, Claire & Burt Monroe. "Analysing Community Reaction to Refugees through Text Analysis of Social Media Data" (2023) 49:2 Journal of Ethnic and Migration Studies 492.
- Kennedy, Duncan. "Legal Education and the Reproduction of Hierarchy" (1982) 32:4 Journal of Legal Education 38.
- Khan, Jon & Sean Rehaag. "Promoting Privacy, Fairness and the Open Court Principle in Immigration and Refugee Proceedings" (2023) 54:2 Ottawa L Rev 357.
- Kirschenbaum, Matthew G. "What Is Digital Humanities and What's It Doing in English Departments?" (2010) 150 ADE Bulletin.
- Kugel, James. "The emergence of biblical interpretation in antiquity" in Mordechai Cohen and Adele Berlin (eds) *Interpreting Scriptures in Judaism, Christianity, and Islam: Overlapping inquiries* (Cambridge: Cambridge University Press, 2016).
- Kuperman, Victor et al. "A lingering question addressed: Reading rate and most efficient listening rate are highly similar" (2021) 47:8 Journal of Experimental Psychology: Human Perception and Performance 1103.
- Langdell, Christopher Columbus. "Harvard Celebration Speech" (1887) 3:1 L Q Rev 123.
- Larsen, Norman. "Statute Revision and Consolidation: History, Process, and Problems" (1987) 19:2 Ottawa L Rev 321.
- Lax-Martinez, Gema et al. "Expanding the World Gender-Name Dictionary: WGND 2.0" (2023) World Intellectual Property Organization (WIPO) Economic Research Working Paper Series No. 64.
- Liew, Jamie Chai Yun et al. "Not Just the Luck of the Draw? Exploring Competency of Counsel and Other Qualitative Factors in Federal Court Refugee Leave Determinations (2005-2010)" (2021) 37:1 Refuge 61.
- Luscombe, Alex, Kevin Dick & Kevin Walby. "Algorithmic thinking in the public interest: navigating technical, legal, and ethical hurdles to web scraping in the social sciences" (2022) 56:3 Quality & Quantity 1023.
- Martineau, Robert J. "The Value of Appellate Oral Argument: A Challenge to the Conventional Wisdom" (1986) 72 Iowa L Rev 1.

- McGill Jenna, and Amy Salyzyn. “Judging by the Numbers: Judicial Analytics, the Justice System Judging by the Numbers: Judicial Analytics, the Justice System and its Stakeholders” (2021) 44:1 Dal LJ 249.
- McGrath, Laura, Richard Jean So & Chad Wellmon, “Culture, Theory, Data: An Introduction” (2022) 54:1 New Literary History 519.
- McGregor, Gaile. “Anti-claimant bias in the employment insurance appeals system: Causes, consequences, and public law remedies” (2001) 15 Canadian Journal of Administrative Law & Practice 229.
- McKenna, Nick, et al. “Sources of Hallucination by Large Language Models on Inference Tasks” (2023) arXiv preprint arXiv: 2305.14552, online: <<https://arxiv.org/abs/2305.14552>>
- Metzler, Jack. “Most Tweetable Justice: An Empirical Study” (2015) 18 Green Bag's Micro-Symposium on Supreme Court Rankings.
- Mikolov, Tomas. “Efficient estimation of word representations in vector space” (2013) arXiv, 1301.378.
- Molnar, Petra. “Technology on the margins: AI and global migration management from a human rights perspective” (2019) 8:2 Cambridge International Law Journal 8.2 305
- Moretti, Franco & Dominique Pestre. “Bankspeak: The Language of World Bank Reports” (2015) 92 New Left Review 75.
- Moretti, Franco & Oleg Sobchuk. “Hidden in Plain Sight” (2019) 119 New Left Review 86
- Moretti, Franco. “The Roads to Rome” (2020) 124 New Left Review 125.
- . “Conjectures on World Literature” in Franco Moretti (ed.), *Distant Reading* (London: Verso, 2013).
- . “The Slaughterhouse of World Literature” in Franco Moretti (ed.), *Distant Reading* (London: Verso, 2013).
- . “Network Theory, Plot Analysis” in Franco Moretti (ed.), *Distant Reading* (London: Verso, 2013).
- Murray, Simone. “Varieties of Digital Literary Studies: Micro, Macro, Meso” (2022) 16:2 DHQ.
- Nasager, Alexandra. “The Supreme Court, Functional Bilingualism, and the Indigenous Candidate: Reconciling the Bench” (2019) 57 Alta L Rev 797
- Nerghes, Adina & Ju-Sung Lee. “Narratives of the Refugee Crisis: A Comparative Study of Mainstream-Media and Twitter” (2019) 7:2 Media and Communication 275.
- Nesbitt, Michael. “An Empirical Study of Terrorism Charges and Terrorism Trials in Canada between September 2001 and September 2018” (2019) 67:1/2 Crim LQ 95.

- Nussbaum, Zach et al. “Nomic Embed: Training a Reproducible Long Context Text Embedder” (2024) arXiv, 2402.01613.
- Olmedilla, M., MR Martínez-Torres. & Toral, S.L. “Harvesting Big Data in social science: A methodological approach for collecting online user-generated content” (2016) 46 *Computer Standards & Interfaces* 79.
- Olsen, Stein Haugom. “The Discipline of Literary Studies” in Andrea Selleri & Philip Gaydon, eds, *Literary Studies and the Philosophy of Literature: New Interdisciplinary Directions* (New York: Springer, 2016).
- Öztürk, Nazan & Serkan Ayvaz. “Sentiment analysis on Twitter: A text mining approach to the Syrian refugee crisis” (2018) 35:1 *Telematics and Informatics* 136.
- Pedregosa, F et al, “Scikit-learn: Machine Learning in Python” (2011) 12 *Journal of Machine Learning Research* 2825.
- Peltomaa, Arthur. “Gender and Racial Diversity of Counsel at the Supreme Court of Canada: An Empirical Study” (2021), online:
https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3930191
- Petukhova, Alina, Matos-Carvalho, Joao P. & Fachada, Nuno. “Text Clustering with LLM Embeddings” (2024) arXiv, 2403.15112
- Porter, Jared. “No Rebels Allowed: The Subversion Bar in Canada’s Immigration Legislation” (2018) 8:1 *Sask L Rev* 25.
- Potter, Daniel. “Unemployment Insurance: Policies and Principles of Disqualification and Disentitlement for Benefits” (1976) 3 *Dalhousie LJ* 178.
- Pound, Roscoe. “Law in Books and Law in Action” (1910) 44:1 *Am L Rev* 12.
 ----- . “The Call for a Realist Jurisprudence” (1931) 44:5 *Harv LR* 697
- Priel, Dan. “Holmes’s ‘Path of the Law’ as Non-Analytic Jurisprudence” (2016) 35:1 *UQLJ* 57.
- Radford, Alec et al. “Robust Speech Recognition via Large-Scale Weak Supervision” (2022) Arxiv, 2212.04356.
- Ramos, Juan. “Using tf-idf to determine word relevance in document queries” (2003) 242:1 *Proceedings of the First Instructional Conference on Machine Learning* 29.
- Ravanelli, Mirco et al. “SpeechBrain: A General-Purpose Speech Toolkit” (2021) arXiv, 2106.0462.
- Rehaag, Sean & Pierre-André Thériault. “Judgments v Reasons in Federal Court Refugee Claim Judicial Reviews: A Bad Precedent” (2022) 45:1 *Dal LJ* 185.
- Rehaag, Sean. “The Role of Counsel in Canada’s Refugee Determination System: An Empirical Assessment” (2011) 49:1 *Osgoode Hall LJ* 71.

- . “Judicial Review of Refugee Determinations: The Luck of the Draw?” (2012) 38:1 Queen's LJ 1.
- . “Judicial Review of Refugee Determinations (II): Revisiting the Luck of the Draw” (2019) 45:1 Queen's LJ 1.
- . “Claim Types in Canada’s Refugee Determination System: An Empirical Snapshot (2013–2021)” (2024) 40:1 Refuge 1.
- . “Luck of the Draw III: Using AI to Extract Data About Decision-Making in Federal Court Stays of Removal” (2024) 49:2 Queen’s LJ 73
- Rennie, Jason & Ryan Rifkin. “Improving multiclass text classification with the support vector machine” Massachusetts Institute of Technology, Department of Brain & Cognitive Sciences and Artificial Intelligence Laboratory, 2001.
- Rhody, Lisa Marie. “Beyond Darwinian Distance: Situating Distant Reading in a Feminist Ut Pictura Poesis Tradition” (2017) 132:3 PMLA 659.
- Ringsmuth, Eve M et al. “SCOTUS in the time of COVID: The evolution of justice dynamics during Oral arguments” (2023) 45:1 Law & Policy 66.
- Rissland, Edwina & David Skalak. “CABARET: rule interpretation in a hybrid architecture” (1991) 34 Int J Man-Machine Studies 839.
- Roberts, John G. “Oral Advocacy and the Re-emergence of a Supreme Court Bar” (2005) 30:1 Journal of Supreme Court History 68
- Ruhl, JB & Daniel Katz. “Measuring, Monitoring, and Managing Legal Complexity” (2015) 101 Iowa L Rev 191.
- Sergot, MJ et al. “The British Nationality Act as a logic program” (1986) 29 Comm ACM 370.
- Shi, Yaqian & Lei Lei. “Lexical Richness and Text Length: An Entropy-Based Perspective” (2022) 29:1 J Quantitative Linguistics 62.
- Shoemaker, Matthew. “Bilingualism and Bijuralism at the Supreme Court of Canada” (2012) 35:2 Canadian Parliamentary Review 30.
- Simon, Zoltán Boldizsár & Marek Tamm. “Historical Futures” (2021) 60:1 History and Theory 3
- So, Richard Jean. “All models are wrong” (2017) 132:3 PMLA 668.
- Stefanowski, Jerzy. “Overlapping, Rare Examples and Class Decomposition in Learning Classifiers from Imbalanced Data” in Zachary A. Pardos & Tiffany Barnes, eds, *Emerging Paradigms in Machine Learning* (Berlin: Springer Berlin Heidelberg, 2013) 277.
- Terdiman, Richard. “The force of law: toward a sociology of the juridical field: translator’s introduction” (1987) 38 Hastings LJ 805.

- Thirunavukarasu, Arun James et al. “Large language models in medicine” (2023) 29 Nature Medicine 1930.
- Topak, Özgün. “Humanitarian and Human Rights Surveillance: The Challenge to Border Surveillance and Invisibility?” (2019) 17:3/4 Surveillance and Society 382.
- Traag, Vincent et al. “From Louvain to Leiden: guaranteeing well-connected communities” (2019) 9:1 Scientific Reports 5233.
- Vaswani, Ashish, et al. “Attention Is All You Need” (2017) arXiv, 1706.03762.
- Vivo, Pierpaolo, Daniel Katz & Ruhl, JB. “A complexity science approach to law and governance” (2024) 382:2270 Phil Trans R Soc A 1.
- Wallace, Simon. “The New Canadian Law of Refugee Exclusion: An Empirical Analysis of International Criminal Law Deportation Orders, January 2018 to July 2020” (2022) 22:4 Intl Crim L Rev 721.
- . “‘Police Authority is Necessary’: The Canadian Origins of the Legal Powers to Detain and Deport, 1893-1902” (2023) 48:2 Queens LJ 101.
- . “The next jailor: An empirical study of danger to the public immigration detentions in Canada (summer 2021)” (2024) 26:3 Punishment & Society 566.
- Watson, Juda. “Gender on the Billboard Hot Country Songs Chart: 1996-2016” (2019) 42 Popular Music and Society 528.
- Weinrib, Ernest. “Legal Formalism” in Dennis M Patterson, ed, *A Companion to Philosophy of Law and Legal Theory*, 2nd ed (Chichester: Wiley-Blackwell, 2010).
- Whalen, Ryan. “Judicial Gobbledygook: the readability of Supreme Court writing” (2015) 125 Yale LJ 200.
- Wilks, Yorick. “Jones, Karen Ida Boalth Spärck (1935–2007)” in Oxford Dictionary of National Biography, (06 January 2011, updated 25 September 2014)
- Williams, Raymond. “Marxism, Structuralism and Literary Analysis” (1981) 129 New Left Review 51.
- Williams, Raymond. “The Uses of Cultural Theory” (1986) 158 New Left Review 19.
- Wolfson, Warren D. “Oral Argument: Does It Matter” (2001) 35 Ind L Rev 451
- Ying, X. “An Overview of Overfitting and Its Solutions” (2019) 1168 J Phys: Conf Ser 022022.

Other

- “Blue J Raises \$9M in Series B Funding, Adds Key Personnel, Announces New Tax Diagramming Solution” businesswire (10 August 2021), online:
 <<https://www.businesswire.com/news/home/20210810005199/en/Blue-J->

Raises-9M-in-Series-B-Funding-Adds-Key-Personnel-Announces-New-Tax-Diagramming-Solution>.

Amnesty International, “Bangladesh. Human Rights in Asia-Pacific: Review of 2019” (30 January 2020), online: <[amnesty.org/en/latest/news/2020/01/2019-in-review-bangladesh/](https://www.amnesty.org/en/latest/news/2020/01/2019-in-review-bangladesh/)>.

Angwin, Julia et al. “Machine Bias: There’s software used across the country to predict future criminals. And it’s biased against blacks” (23 May 2016), online: <<https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>>.

Australian Government, DFAT Country Information Report: Bangladesh (Department of Foreign Affairs and Trade, 2019), online: <dfat.gov.au/sites/default/files/country-information-report-bangladesh.pdf>.

Central Intelligence Agency, “Bangladesh” (last updated 16 November 2023), online: <cia.gov/the-world-factbook/countries/bangladesh/> [perma.cc/5W98-JZSW].

Department of Justice, “Consolidated federal Acts and regulations –Bulk XML (PIT (Point in time)),” online: <<https://open.canada.ca/data/en/dataset/4d277b82-c4af-4431-b31d-48fcece6dd8a/resource/dd28f08b-0b48-485a-a9db-b6176ff3d1f7>>.

Employment and Social Development Canada, “Helping middle-class Canadians with the support they need: Government of Canada announces new reforms to the Social Security Tribunal” (15 August 2019), online: <<https://www.canada.ca/en/employment-social-development/news/2019/08/helping-middle-class-canadians-with-the-support-they-need.html>>.

Employment and Social Development Canada, Digest of Benefit Entitlement Principles (Ottawa), online: <https://www.canada.ca/en/employment-social-development/programs/ei/ei-list/reports/digest.html>>.

Employment and Social Development Canada, Review of the Social Security Tribunal of Canada for Employment and Social Development Canada – October 2017 (Ottawa: ESDC, 2017), online: <<https://www.canada.ca/en/employment-social-development/corporate/reports/evaluations/social-security-tribunal-review.html>>.

European Asylum Support Office, Country of Origin Information Report: Bangladesh Country Overview (December 2017), online: <coi.europa.eu/administration/easo/Plib/Bangladesh_Country_Overview_December_2017.pdf> [perma.cc/SK8C-J7TS]

Federal Court of Canada, “Notice: IMM Settlement Pilot (Toronto)” (4 July 2019), online: <[https://www.fct-cf.gc.ca/content/assets/pdf/base/Notice%20IMM%20settlement%20pilot%20\(Toronto\)%20-%20AMENDED%20July%204-2019%20FINAL%20ENG.pdf](https://www.fct-cf.gc.ca/content/assets/pdf/base/Notice%20IMM%20settlement%20pilot%20(Toronto)%20-%20AMENDED%20July%204-2019%20FINAL%20ENG.pdf)>.

- Federal Court, “Policy on Public and Media Access” (last modified 29 March 2023), online: www.fct-cf.gc.ca/en/pages/media/policy-on-public-and-media-access [perma.cc/NM6U-X2TS].
- Gorsuch, Neil & Janie Nitze. “America Has Too Many Laws” (5 August 2024) The Atlantic, online: <<https://www.theatlantic.com/ideas/archive/2024/08/america-has-too-many-laws-neil-gorsuch/679237/>>.
- Hao, Karen. “We read the paper that forced Timnit Gebru out of Google. Here’s what it says.” (4 December 2020), online: <<https://www.technologyreview.com/2020/12/04/1013294/google-ai-ethics-research-paper-forced-out-timnit-gebru/>>.
- Honnibal, Matthew & Ines Montani. “spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing” (2017).
- HuggingFace, “Hugging Face – The AI community building the future,” online: <<https://huggingface.co/>>.
- Human Rights Watch, “Bangladesh: Events of 2019” (2020), online: <hrw.org/world-report/2020/country-chapters/bangladesh>.
- Immigration and Refugee Board of Canada, “Claims by Country of Alleged Persecution - 2022”, online: <<https://irb.gc.ca:443/en/statistics/protection/Pages/RPDStat2022.aspx>>.
- Immigration and Refugee Board of Canada, “Refugee appeals statistics” (15 May 2024), online: Immigration and Refugee Board of Canada <<https://www.irb-cisr.gc.ca/en/statistics/appeals/Pages/index.aspx>>.
- Immigration and Refugee Board, “National Documentation Packages” (23 Jan 2024) online: <<https://www.irb-cisr.gc.ca/en/country-information/ndp/Pages/index.aspx>>.
- Immigration and Refugee Board, Quality Assurance Framework for Decision-Making (Ottawa, 2021), online (pdf): <irb.gc.ca/en/transparency/qa-aq/Documents/Quality_Assurance_Framework_for_Decision_Making_2021.pdf> [perma.cc/J7JP-B8PP].
- Immigration, Refugees, and Citizenship Canada, “Asylum claims by year – 2023” (22 April 2024) online: <<https://www.canada.ca/en/immigration-refugees-citizenship/services/refugees/asylum-claims/asylum-claims-2023.html>>.
- Jarlnar, Asta Sofie Stage. “Using topic modelling to explore case summaries” (28 March 2023), Nordic Asylum Law & Data Lab (blog), online: <<https://asylumdata.ku.dk>>.
- Liptak, Adam. “Clarence Thomas Breaks 10 Years of Silence at Supreme Court”, The New York Times (29 February 2016), online:

- Meurrens, Steven. “The Bangladesh Nationalist Party” (13 August 2019) online (blog): meurrensonimmigration.com/the-bangladesh-nationalist-party/ [perma.cc/HKP5-FUJS]
- Minister of Citizenship and Immigration, “Intake of claims for refugee protection at ports of entry” (24 December 2023) online: <<https://www.canada.ca/en/immigration-refugees-citizenship/corporate/publications-manuals/operational-bulletins-manuals/refugee-protection/canada/intake-claims-refugee-protection-ports-entry.html#>>>.
- Mowat Centre for Policy Innovation, “What the New EI Rules Mean” (2012), online: <https://mowatcentre.munkschool.utoronto.ca/wp-content/uploads/publications/52_what_the_new_ei_rules_mean.pdf>.
- Mueller, Andreas. “wordcloud,” online: <https://github.com/amueller/word_cloud>
- Office of the Commissioner for Federal Judicial Affairs Canada, “Qualifications and Assessment Criteria”, online: <<https://www.fja-cmf.gc.ca/scc-csc/2023/qualifications-eng.html>>.
- Ollama, online: <<https://github.com/ollama/ollama>>.
- Phelim Kine, “My journey down the rabbit hole of every journalist’s favorite app”, (16 February 2022), online: POLITICO <<https://www.politico.com/news/2022/02/16/my-journey-down-the-rabbit-hole-of-every-journalists-favorite-app-00009216>>.
- Public Safety Canada, “Security Certificates” (last modified 1 December 2015), online: <publicsafety.gc.ca/cnt/ntnl-scrt/cntr-trrrsm/scrt-crtfcts-en.aspx> [perma.cc/39G7-RCEK].
- Refugee Law Lab Portal, “Federal Court” (visited 13 June 2024), online: Refugee Law Lab <<https://rllp.ca/federal-court>>.
- Refugee Law Lab, “Luck of the Draw III Dataset” (2023), online: Hugging Face <<https://huggingface.co/datasets/refugee-law-lab/luck-of-the-draw-iii>>.
- Rehaag, Sean. “2013 Refugee Claim Data and IRB Member Recognition Rates” (14 April 2014), online: <ccrweb.ca/en/2013-refugee-claim-data>.
- Rehaag, Sean. “2023 Refugee Claim Data and IRB Member Recognition Rates” (29 April 2024), online: <<https://refugeelab.ca/refugee-claim-data-2023>>.
- Robertson, Joseph. “Administrative Deference: The Canadian Doctrine that Continues to Disappoint,” (18 April 2018) online: <canlii.ca/t/stvr>. [perma.cc/VKC6-XEHJ]
- Sean Rehaag, “Federal Court Bulk Decisions Dataset” (2023), online: Refugee Law Laboratory <https://refugeelab.ca/bulk-data/fc>
- Sean Rehaag, “SST Bulk Decisions Dataset” (2023), online: Refugee Law Laboratory <https://refugeelab.ca/bulk-data/sst>

- Sharma, Raj. “The Bangladesh National Party (BNP) is (Not) a Terrorist Organization” (18 April 2017) online (blog): <sshlaw.ca/the-bangladesh-national-party-bnp-is-not-a-terrorist-organization/>. [perma.cc/F72N-ML2F]
- Sherman, Jo. Guidelines For Canadian Courts: Management of Requests for Bulk Access to Court Information by Commercial Entities (Canadian Judicial Council: Ottawa, 2021).
- Social Security Tribunal of Canada “What decisions the Social Security Tribunal publishes,” online: <<https://www.sst-tss.gc.ca/en/decisions-laws-rules-and-policies/what-decisions-social-security-tribunal-publishes>>.
- Social Security Tribunal of Canada, “An evaluation of how easy it is to read decisions of the Social Security Tribunal” (10 January 2024), online: <<https://www.sst-tss.gc.ca/en/our-work-our-people/evaluation-easy-it-read-decisions-social-security-tribunal>>.
- Social Security Tribunal of Canada, “Terms and Conditions” (visited 29 July 2024), online: <<https://sst-tss.gc.ca/en/terms-and-conditions>>.
- Social Security Tribunal of Canada, Publication of KPMG’s Review of the Social Security Tribunal (January 2018), online: <<https://sst-tss.gc.ca/en/whats-new/publication-kpmgs-review-social-security-tribunal>>.
- Statistics Canada, Table 14-10-0005-01 Employment insurance claims received by province and territory, monthly, seasonally adjusted (Ottawa: Statistics Canada), online: <<https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410000501>>.
- Statistics Canada, Table 14-10-0007-01 Employment insurance benefit characteristics by class of worker, monthly, unadjusted for seasonality (Ottawa: Statistics Canada), online: <<https://www150.statcan.gc.ca/t1/tbl1/en/tv.action?pid=1410000701>>.
- Supreme Court of Canada, “Supreme Court of Canada - Archived Webcasts” online: <<https://www.scc-csc.ca/case-dossier/info/webcasts-webdiffusions-eng.aspx>>.
- Supreme Court of Canada, “Supreme Court of Canada - Judges of the Court”, online: <<https://www.scc-csc.ca/judges-juges/index-eng.aspx>>.
- United States of America, Bureau of Democracy, Human Rights, and Labor, Country Reports on Human Rights Practices for 2020: Bangladesh (Department of State, 2021), online: <state.gov/reports/2020-country-reports-on-human-rights-practices/Bangladesh>.
- University of Copenhagen, “Nordic Asylum Law & Data Lab” (last accessed 25 June 2024), online: <<https://asylumdata.ku.dk>>.
- UNSW Sydney, “Kaldor Centre Data Lab” (last accessed 25 June 2024), online: <<https://www.unsw.edu.au/kaldor-centre/our-resources/kaldor-centre-data-lab>>

Vinge, Vernon. "The Coming Technological Singularity: How to Survive in the Post-Human Era" (1993) online:
<<https://edoras.sdsu.edu/~vinge/misc/singularity.html>>

W3C, "Extensible Markup Language (XML)," online: <<https://www.w3.org/XML/>>.