

An Empirical Study on Judgments Written by the Supreme Court Judges of Singapore

Lim How Khang¹

Abstract

Judges appear to be writing increasingly longer judgments. To uncover this curious trend, this paper presents empirical findings of the writing styles and choices of Singapore judges from 7,807 of their written judgments spanning the two decades past (2003 to 2023).

1 Introduction

The recent decision *Public Prosecutor v Soh Chee Wen* [2023] SGHC 299 set a new record for the longest Supreme Court judgment at 896 pages. This was the fifth time the judgment length record has been broken in the past eight years.² Have we arrived at the era of epic-judgments?

Prior research show that judgment lengths in Singapore have increased over time. It was observed by Goh and Tan,³ and later by Soh,⁴ that judgments reported in the Singapore Law Reports grew, on average, more than twice as long since the early 2000s.

If judgments are becoming longer at record setting pace, this raises a few implications. Firstly, for lawyers and law students, keeping up with the law

¹ Assistant Professor of Law and Computer Science (Practice), Yong Pung How School of Law, Singapore Management University. This research was partially funded by the Singapore Judicial College under its Empirical Judicial Research initiative.

² The previous record holders were: *GTMS Construction Pte Ltd v Ser Kim Koi (Chan Sau Yan and Chan Sau Yan Associates, third parties)* [2021] SGHC 9; *Aljunied-Hougang Town Council v Lim Swee Lian Sylvia* [2019] SGHC 241; *Public Prosecutor v Lam Leng Hung* [2017] SGHC 71; and *Nava Bharat (Singapore) Pte Ltd v Straits Law Practice LLC* [2015] SGHC 146

³ Goh Yihan and Paul Tan. "An Empirical Study on the Development of Singapore Law". In: *Singapore Academy of Law Journal* 23 (2011), pp. 176–226.

⁴ Jerrold Soh. "A Network Analysis of the Singapore Court of Appeal's Citations to Precedent". In: *Singapore Academy of Law Journal* 31 (2019), pp. 246–284.

will require more time and effort. Secondly, longer judgments are likely to make it harder for the general public to understand judges' decisions.

Imposing word or page limits on judgments⁵ is probably too drastic as there may be valid reasons for writing longer judgments. Especially if the result of a longer judgment is legal clarity.

2 Aims of the study

Therefore, this paper has two broad goals. First, to understand what Singapore judges are writing in their judgments and whether there are any identifiable trends. Second, to evaluate whether these judgments are accessible to the general public.

With this, I considered the following questions:

1. Which judgments are becoming longer?
2. Are judges writing more on certain subjects?
3. Are judges writing more because they are doing more legal reasoning?
4. Are judges writing more or are they quoting others?
5. Are judgments difficult to read?

3 Methodology

3.1 The corpus

To draw valid conclusions about judgment writing, the size of the corpus is important. Personal experience and research elsewhere⁶ indicate that judgments vary widely on features that might be relevant to this study such as length and number of citations. Therefore the goal was to use a comprehensive corpus that covered as many judgments as possible.

The corpus used for this study is a collection of 7,807 Singapore judgments written in English by judges of the Supreme Court. As I was interested in judgment accessibility, I felt that it made more sense to use the

⁵ Neil H. Andrews. "The Supreme Court of the United Kingdom and English Court Judgements". doi: <http://dx.doi.org/10.2139/ssrn.2399103>, p. 3.

⁶ Soh, "A Network Analysis of the Singapore Court of Appeal's Citations to Precedent".

Singapore Courts' unreported judgments for this study because the general public is more likely to visit the Singapore Courts' website than rely on reported judgments behind a paywall.

The source judgments were downloaded from the publicly accessible Singapore Courts website⁷ in both HTML and PDF formats.⁸⁹

Judgment count	7,807
Token count	69 million
Earliest judgment	02 Jan 2003
Latest judgment	29 Dec 2023

Table 1: Corpus statistics

A point to note is that the judgments came in two different formats and both the HTML and PDF versions were similarly affected. Those before 2016 occasionally had inconsistent markup and formatting whereas those since 2016 were more consistent and better formatted (but still not perfect). This had an impact on the data quality, especially the older judgments. Although basic sanity checks were done to detect obvious errors, the sheer amount of data meant that it was not possible to comb through every single detail.

3.2 Text processing and mapping

The judgments were processed in two stages. First, the HTML files and PDF files were parsed to extract the source text, formatting and layout into the following data points: citation, court level, coram, decision date, catchwords,⁸ headings, paragraphs, annotations, tables and counsel. For judgments since 2016, I also extracted the hearing dates.

After a round of preliminary data exploration (which included reviewing unique values and counts), I designed a few mapping functions. Values for court level were mapped so that High Court judgments were grouped with judgments by the General and Family Divisions of the High Court. The values

⁷ <https://www.judiciary.gov.sg/judgments/judgments-case-summaries>

⁸ Two judgments were left out of the corpus due to practical constraints: [2016] SGHC 46 (missing the HTML version) and [2016] SGHC 239 (a short supplemental judgment). Earlier judgments from 2000 to 2002 were also available for download from the Singapore Courts website but they were deemed to be of insufficient quality and consistency to be included in the corpus.

⁹ cases were missing catchwords.

for author and catchwords were also mapped to handle inconsistent naming or spelling conventions in the source judgments.

Next, I used a natural language processing pipeline that I developed using the open source spaCy library¹⁰ to generate the additional data points needed for computing the measures described in the following section.

3.3 Measuring what judges are writing

To find out what judges are writing, I needed to decide what to measure in a meaningful way. I decided to focus on whether judges are spending more time on legal reasoning.

Number of legal authority mentions. To measure legal reasoning, I counted the number of times a legal authority is mentioned in a judgment (including the first mention). I assume that frequent mentions of legal authorities within a judgment means that the judge engaged in a discussion about the law or its application to the facts. Prior research studied the number of unique legal citations in judgments but did not take the number of mentions into account. While legal citation counts can show how many and which authorities were cited by a judge, they do not tell us whether that authority was discussed or merely cited for its legal rule.

To identify legal authorities I used spaCy's custom pipeline components and their EntityRuler. I designed token-based rules to capture citation formats for case law, books, articles and periodicals commonly cited in Singapore judgments. For US legal authorities, I relied on the open source eyecite¹¹ library.

Once a full citation was found in the judgment, I used a custom component pipeline to locate all subsequent mentions of that legal authority. The first step was to locate the legal authority's abbreviated title as defined by the judge immediately after the full citation. Then the next steps were to find all subsequent mentions of the abbreviated title. Following conventional citation practice, the title of a book or the case name is italicised and a journal article title is placed within double quotation marks. To avoid detecting false

¹⁰ Matthew Honnibal et al. "spaCy: Industrial-strength Natural Language Processing in Python". 2020. doi: 10.5281/zenodo.1212303.

¹¹ Jack Cushman, Matthew Dahl, and Michael Lissner. "eyecite: A Tool for Parsing Legal Citations". In: *Journal of Open Source Software* 6.66 (2021), p. 3617. url: <https://doi.org/10.21105/joss.03617>.

positives, I only selected the mentions that followed this style. However, I later learned that this might be an issue for the earliest judgments in the corpus (2003 and 2004), as they were not formatted consistently. Due to time constraints, I was unable to verify this issue, hence the numbers for the years 2003 and 2004 on this data point should be ignored for now as they are likely to be underestimated.

3.4 Measuring judgment accessibility

To measure judgment accessibility, we need to define what this means. Accessibility to judgments should go beyond hosting them on a website. Downloadable judgments are not truly accessible if readers find it difficult to understand what they are reading.

Osbeck suggests that good legal writing should be clear, concise and engaging.¹² In defining “clarity”, Osbeck focuses on whether the writing helps the reader in legal decision making. Paying attention to the reader’s needs gives the writer flexibility to treat all prescriptions on good writing as suggestions to be considered responsibly for the reader’s sake. Seen this way, clear writing is a posture the writer adopts in service to the reader.

It is in this spirit that I suggest the following measurements to evaluate the accessibility of judgments. Many of them have obvious limitations on their own but contain some

The results are thus intended to reflect conventional wisdom and not meant as a performance indicator of sorts.

3.4.1 Word and sentence lengths

Word lengths are associated with word difficulty and long sentences take longer to process than short sentences.

To split the judgment text into sentences, I used a combination of the rules-based Sentencizer from spaCy and the `nltk.tokenize.punkt` module from the open source nltk¹³ library. I found that rules-based algorithms worked better on judgment texts as some of the sentences run very long and

¹² Mark K. Osbeck. “What is “Good Legal Writing” and Why Does it Matter?” In: *Drexel Law Review* 4 (2012), pp. 417–467.

¹³ Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. O’Reilly Media, Inc., 2009. url: <https://www.nltk.org/book/>.

dependency parser based models might not detect sentence boundaries as accurately on them. To split the text into words, I used spaCy's default tokeniser but added a few rules to handle unusual character sequences that are unique to legal writing. I also count sentence fragments that are found in lists by treating each line item as its own sentence.

3.4.2 Readability scores

The widely used Flesch-Kincaid tests are built on the idea that it is possible to find a formula that combines word-level and sentence-level scores into a more useful number. For example, the Flesch-Kincaid Grade Level formula computes a score (representing the US school grade level) from the syllable counts and sentence lengths as inputs:

$$0.39(\text{words/sentence}) + 11.8(\text{syllables/word}) - 15.59$$

However, their simplicity means that they do not account for other factors that affect readability or reading comprehension. Experts thus caution against using such tests other than as a rough guide. Besides, different software uses different rules for syllable counting and sentence boundary detection, which makes it impossible to directly compare readability scores computed by different software programs.¹⁴¹⁵ Yet despite their limitations, these simple tests can often provide a quick sense of text difficulty when used appropriately with those limitations in mind and in conjunction with measures that consider other factors.

To compute the Flesch-Kincaid Grade Level scores for each judgment, I needed to split the text into sentences and words and assign a syllable count to each word.

I obtained the syllable counts by looking up each word in the CMU Pronouncing Dictionary¹⁶. If the word was not in the dictionary, its syllable count was estimated based on the open source code by readable.com¹⁷ which uses a set of rules and exceptions based on character combinations.

¹⁴ Shixiang Zhou, Heejin Jeong, and Paul A. Green. "How Consistent Are the BestKnown Readability Equations in Estimating the Readability of Design Standards?" In: *IEEE Transactions on Professional Communication* 60.1 (2017), pp. 97–111. doi: 10.

¹⁵ /TPC.2016.2635720.

¹⁶ <http://www.speech.cs.cmu.edu/cgi-bin/cmudict>

¹⁷ <https://github.com/DaveChild/Text-Statistics/blob/master/src/DaveChild/TextStatistics/Syllables.php>

I then implemented the algorithm and ran the program against the 18 passages in the Flesch-Kincaid paper¹⁸ and verified that except for passage 12, the scores were consistent or within a single grade difference. As I was not looking for precise grade levels and simply wanted to use the scores as a rough guide, I went ahead and computed the Flesch-Kincaid Grade Level data points for each judgment in the corpus.

3.4.3 Words frequently encountered by the average person

Another way to estimate word difficulty is to consider its frequency in word lists based on everyday language, since rarely encountered words are likely to be more difficult to understand.

I use the SUBTLEX-UK word frequency list¹⁹ to classify whether a word is difficult. The SUBTLEX-UK word list contains all the words (and their frequencies) taken from subtitles of selected British television broadcasts between 2010 and 2012. A subtitle-based word list is closer to everyday language than a word list based on a written corpus. Admittedly, some accuracy is lost when using a British word list, as English word usage in Singapore is not entirely the same.

I then estimated text passage difficulty by counting the number of running words within it and calculating the proportion of high frequency words over non-high frequency words.²⁰ A word is classified as 'high frequency' if its Zipf value in the SUBTLEX-UK list is 4 or higher. The higher the coverage of high frequency words within a passage, the easier it is to read, as shown in Figure 1.

¹⁸ J. Peter Kincaid et al. *Derivation Of New Readability Formulas (Automated Readability Index, Fog Count And Flesch Reading Ease Formula) For Navy Enlisted Personnel*. Tech. rep. Institute for Simulation and Training, 1975. url: <https://stars.library.ucf.edu/istlibrary/56>.

¹⁹ Walter J. B. van Heuven et al. "Subtlex-UK: A New and Improved Word Frequency Database for British English". In: *Quarterly Journal of Experimental Psychology* 67.6 (2014), pp. 1176–1190. doi: 10.1080/17470218.2013.850521.

²⁰ Proper nouns, symbols and numbers are excluded and not used in this calculation.

The Family Court heard an application by a wife for a personal protection order against her husband . Both parties appeared in person at the hearing before a District Judge. After hearing them , the judge issued the order sought . The husband 's appeal against the order came up before me .

The parties were married in 1990 and have three children . The marriage ran into difficulties and the wife Law Fay Yuen left the matrimonial home on 6 December 2001.

However , she would return to the house from time to time . The husband Teng Cheng Sin disapproves of this and believes that she should not do that without his consent or an order of court . She on the other hand feels that she is entitled to return to the matrimonial house .

Figure 1: Highlighted words are high frequency words in SUBTLEX-UK.

3.4.4 Legal jargon and archaic words

Once a mainstay in legal proceedings, legal phrases in Latin are experiencing a downgrade in status in Singapore. Since Latin phrase were replaced with updated language in the latest Rules of Court 2021, legal Latin has effectively been declared *persona non grata*. This change was a necessary one as the courts risk isolating themselves from the wider public if they continue to coat their words in a cryptic language that is utterly alien to the average person in Singapore.

To detect their presence in the judgments, I used the list of Latin legal terms from Wikipedia²¹ and spaCy's Matcher. Other old-fashioned words or expressions that commonly appear in legal writings but are almost never used elsewhere anymore (for example, "the said", "hereby", and "thereunder") were also added to the list and identified the same way.

3.4.5 Explicit Connectives

For long passages of text, fluency is an important factor for readability and engagement. To enhance reading fluency, legal writing experts suggest using connecting or transition words to guide the reader between sentences and paragraphs.

²¹ https://en.wikipedia.org/wiki/List_of_Latin_legal_terms

Bailin and Grafstein (2015) point to a number of studies that suggest that appropriate use of explicit connectives may sometimes increase reading comprehension.²² They caution, however, that simply counting connectives is unlikely to be helpful, as they may be useful only if they contribute to coherence, or by making it easier for the reader “to construct conceptual models”,²¹

This does not mean that counting connectives in judgments is necessarily unfruitful. In the right hands, the use of connectives can contribute to the overall coherence of a passage.

To identify connectives, I used spaCy’s PhraseMatcher to match whether each judgment’s paragraph and sentence beginnings contained an explicit connective from the list of 173 connectives in the The Penn Discourse Treebank 3.0 Annotation Manual.

3.4.6 Headings

Legal writing experts generally recommend headings for long documents. If for no other reason, headings help to break up long intimidating walls of text, like welcome rest stops along a lonely highway. Headings also help the reader navigate between different parts of a document. And if they are written to be informative about the case and organised in a logical sequence and hierarchy, headings can help the reader understand the writer’s thought process and logic.

To detect headings in the source judgments, I resorted to developing a fairly complex algorithm because the HTML markup was not sufficiently consistent to provide the data quality I required across the different years. To minimise errors, the algorithm would parse the HTML and PDF files and check for consistency along other factors such as font style and size. If the PDF version had a Microsoft Word generated Table of Contents, I used that as a starting point.²³ To parse the HTML, I used the open source beautifulsoup4²⁴ library and to parse the PDF, I used the open source pdfplumber library.²⁵

²² Alan Bailin and Ann Grafstein. *Readability: Text and Context*. Palgrave Macmillan London, 2015. doi: <https://doi.org/10.1057/9781137388773>, p. 160. ²¹Ibid., p. 161.

²³ Even then there were the occasional inconsistencies, likely to have been caused by the Microsoft Word document author/editor forgetting to update the generated Table of Contents after they had removed or edited a heading in the document body.

²⁴ <https://beautiful-soup-4.readthedocs.io/en/latest/>

²⁵ Jeremy Singer Vine, URL: <https://github.com/jsvine/pdfplumber>

Using the above method, I eventually collected enough examples to train a machine learning binary classifier using fasttext²⁶ which was then used to improve detection accuracy.

3.4.7 First and second person pronouns

The idea of measuring first person pronouns such as “I” and “we” in judgments is rooted in the idea that their presence makes the writing feel more personal for the reader. Williams²⁷²⁸ found in his study on English Court of Appeal judgments that there was a drop in the number of first person pronouns between 1970 and 2020.

In contrast, the presence of second person pronouns like “you” is highly unusual in judgments, as it signals to the reader that the judge’s words were intended for them. This could potentially make for uncomfortable reading. In any event, I decided to include this measure out of curiosity.

To detect personal pronouns, I rely on the part-of-speech tags identified by spaCy’s English model `en_web_core_sm`.

3.4.8 Interesting openings

In a typical judgment, the judge starts by setting the stage to provide context to the dispute, like a recap sequence from last week’s episode of a TV show. This is where the judge finds the most freedom to express their individual styles of writing, and some may use the judgment opening as an opportunity to engage the reader early on.

Osbeck boils down engaging writing to: variety, an authentic voice and sometimes, even humour. Since judgments tend to be highly formulaic, it does not take much to inject some variety for an interesting read.

To discover unconventional opening paragraphs in the corpus, I used an instruction tuned large language model Mistral-7B-Instruct-v0.2.²⁹ Large language models are text generative models that have been trained on vast

²⁶ Armand Joulin et al. “Bag of Tricks for Efficient Text Classification”. In: *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, Apr. 2017, pp. 427–431.

²⁷ C. Williams. *The Impact of Plain Language on Legal English in the United Kingdom*. Routledge studies in language and intercultural communication. Routledge, 2022. doi:

²⁸ .4324/9781003025009.

²⁹ <https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.2>

amounts of text to model language. Further instruction tuning on these large language models enable them to perform new tasks when prompted (zero-shot learning).³⁰

I gave the model the first paragraph of each judgment with this prompt: “What do you make of the writing style of this judge’s opening paragraph for their decision in a legal judgment? Explain in detail by giving reasons and examples.” In response, the model gave me its generated output for each paragraph.

I then took the model’s generated outputs and fed each one back to the model with a new prompt this time: “Based on the above, would you consider the writing style (A) more conventional like a typical judgment, or (B) more engaging than a typical legal judgment? Answer with A or B and then explain.” This gave me a list of candidate openings for me to review.

3.5 Benchmarking against newspaper articles

To provide a point of comparison, I downloaded 1,000 Straits Times articles from the Factiva database. I tried to find overlapping topics for a fairer comparison and I ended up with the following search term: ‘rst=stimes AND re=singp AND (ns=GCRIM OR ns=C12 OR ns=GHOME OR in=i835)’ which covers the following topics:

- Subject-Corporate/Industrial News-Corporate Crime/Legal Action (ns=c12)
- Subject-Political/General News-Crime/Legal Action (ns=gcrim)
- Subject-Political/General News-National/Public Security-Law Enforcement (ns=ghome)
- Industry-Business/Consumer Services-Legal Services (in=i835)
- Region-Singapore (re=singp)

³⁰ Jason Wei et al. “Finetuned Language Models are Zero-Shot Learners”. In: *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. url: <https://openreview.net/forum?id=gEZrGCozdqR>.

I did some minor preprocessing of the text before putting them through the same spaCy pipeline to compute the readability related statistics.

4 Results and Discussion

4.1 What are Singapore judges writing about?

We begin by investigating the subject areas that judges are writing about.

Figure 2 shows the 25 most common subject areas in all judgments and they come from a wide range.

The number of judgments written on criminal matters stands out. If we add the numbers for “Criminal Procedure and Sentencing” and “Criminal Law”, the resulting merged topic leaps to the top of the list, above “Civil Procedure” and “Contract”. This is a remarkable statistic because the Supreme Court hears far more civil matters than criminal ones. To put things into perspective, 12,238 civil matters were filed in the Supreme Court in 2022 while only 488 criminal matters were filed.³¹ In the same year, 108 judgments on criminal matters were written, meaning there was one judgment for every 4 or 5 criminal matters.

There are a few possible ways to interpret this. Firstly, this could be evidence of judges fulfilling their judicial duty to give reasoned decisions corresponding to the requirements of the case.³² Criminal offences that are tried in the Supreme Court carry the most serious punishments. It is likely that judges are being more detailed and transparent about their decisions

³¹ <https://www.judiciary.gov.sg/who-we-are/statistics/caseload-statistics-2022>

³² *Thong Ah Fat v Public Prosecutor* [2012] 1 SLR 676

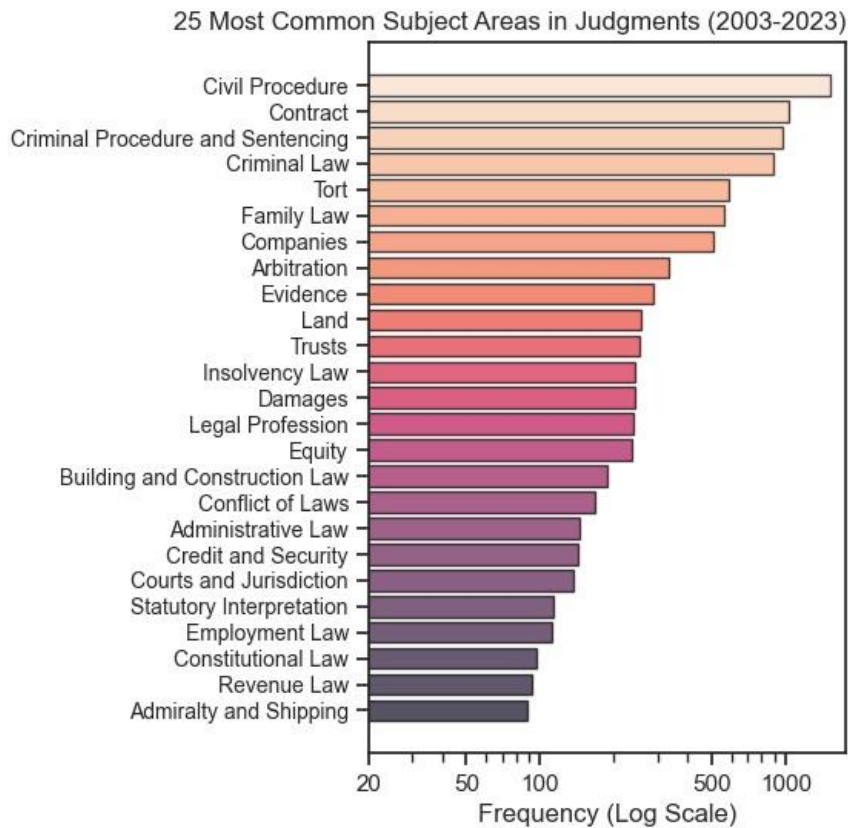


Figure 2: The 25 most common subject matters.

in criminal cases because they recognise the power they wield in such cases and the concomitant responsibility to ensure that justice is seen to be done. The second reason could be that judges who hear Magistrate’s Appeals are writing judgments to set precedent and provide sentencing guidance to the lower courts. Thirdly, judges may be more likely to write judgments on criminal matters because it is in the public interest to do so.

4.2 Are Singapore judges writing more judgments?

Figure 3 shows the different court levels and the number of judgments written. Two trends stand out: judgments by the High Court and the present General and Family Divisions of the High Court have been steadily increasing but the number of Court of Appeal judgments have seen a decline since 2021.

One possible explanation for the drop in Court of Appeal judgments could be the establishment of the Appellate Division of the High Court, which now




Figure 3: Number of judgments by court level.

hears certain civil appeals from the General Division of the High Court.

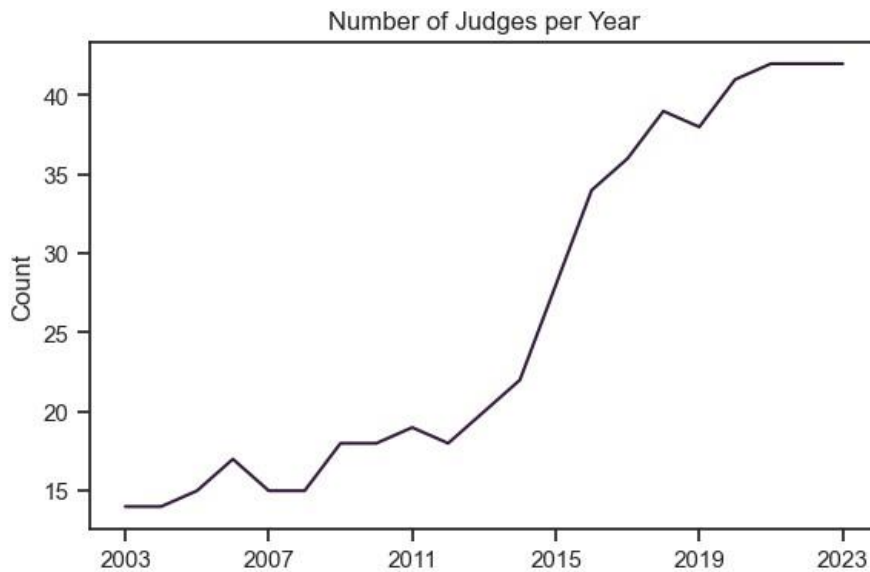


Figure 4: Number of Judges by year (incl. International Judges).

As for the increase in High Court decisions over the past 10 years, this seems to have coincided with the expansion of the Supreme Court bench after Chief Justice Sundaresh Menon’s appointment (see Figure 4 below). As the

new appointees would have begun their tenure in the High Court (and later the General or Family Divisions of the High Court), this could explain the increase in judgments for that group.

4.3 Are Singapore judges writing longer judgments?




Figure 5: Median judgment length by court level.

Median length. Turning to judgment lengths, we see in Figure 5 that the median judgment length of the General Division of the High Court has more than doubled since 2003. The median judgment length of the Court of Appeal saw a rapid rise from 2007 and appears to have peaked between 2016 and 2018.

Corresponding with the earlier observation of a decline in judgment numbers, the median judgment length of the Court of Appeal has also dropped. In fact, the median judgment length of the Court of Appeal appears to have converged again with the median judgment length of the High Court after a breakaway run of around 15 years (2006 to 2021).

Length distributions. Measures of central tendency like the median do not shed any light about the extreme cases that have been observed in

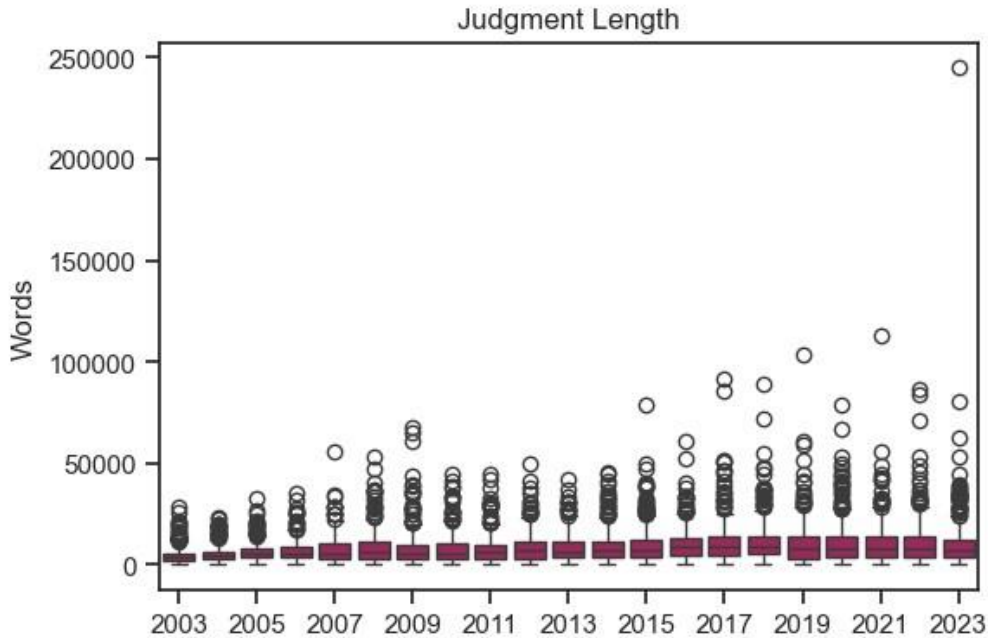


Figure 6: Boxplots showing distribution of judgment lengths with outliers.

recent years. For this, we plot the distributions of judgment lengths over time to identify any trends. Figure 6 shows that these record-breaking judgments lengths are clearly outliers, although there is also an obvious trend of more judgments crossing the 50,000 word barrier since 2015.

Looking at just the record-breaking judgments, there are some common features. Most of them were long-running disputes with multiple skirmishes in between. The number of hearing or trial days tended to be very high as well. In other words, these were fiercely contested legal battles. It is therefore unsurprising that the judges in these cases felt it necessary to lay out an extraordinary amount of material in their judgments to explain their decisions.

Therefore, we need not worry that 800-page judgments will become the norm at this stage unless there is an increase in protracted disputes. But there still remains the question why judges are, on average, writing longer judgments over time.

Quotes. Are judgments becoming longer on average because judges are quoting others more? Apparently not. Plotting the total number of words attributed to the judge and the total number of quoted words over time, it seems like the average number of quoted words has not changed by much over the years even though the average number of words attributed to judges have increased significantly.

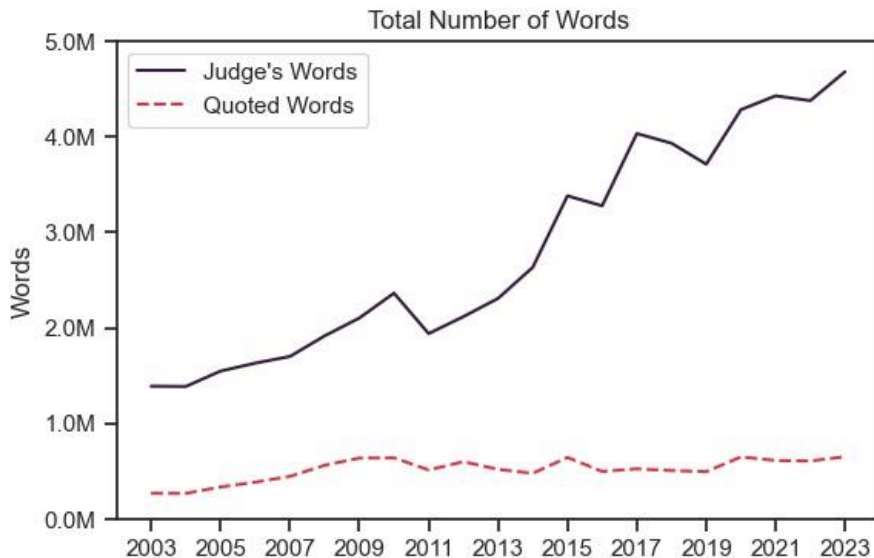


Figure 7: Number of Words in All Judgments by Year.

Based on these statistics alone, it is not possible to tell whether this means that judges' quoting practices have remained unchanged without some form of content analysis. I leave this for future research to address.

Legal reasoning Are judgments becoming longer because judges spend more time discussing legal authorities and doing legal reasoning? On this point, the evidence seems to suggest so. Figure 8 shows the average number of times a legal authority is mentioned in a judgment. Ignoring the years 2003 and 2004 (due to the data quality issues I mention in [XX] above), there is a clear trend of increasing discussion of legal authorities, and this is not limited to judgments of the Court of Appeal.

4.4 Do some judges write longer judgments?

This is a difficult question to answer because there may be a temptation to use the results to infer a causal relationship when there are likely to be other factors at play which are not accounted for in the collected data points. And where there are few writing samples for some individuals, any purported relationship becomes even more tenuous.

There are, however, individual writing styles and they certainly have an influence on overall lengths. For instance, I can say with confidence that




Figure 8: Average number of legal authority mentions per judgment over time

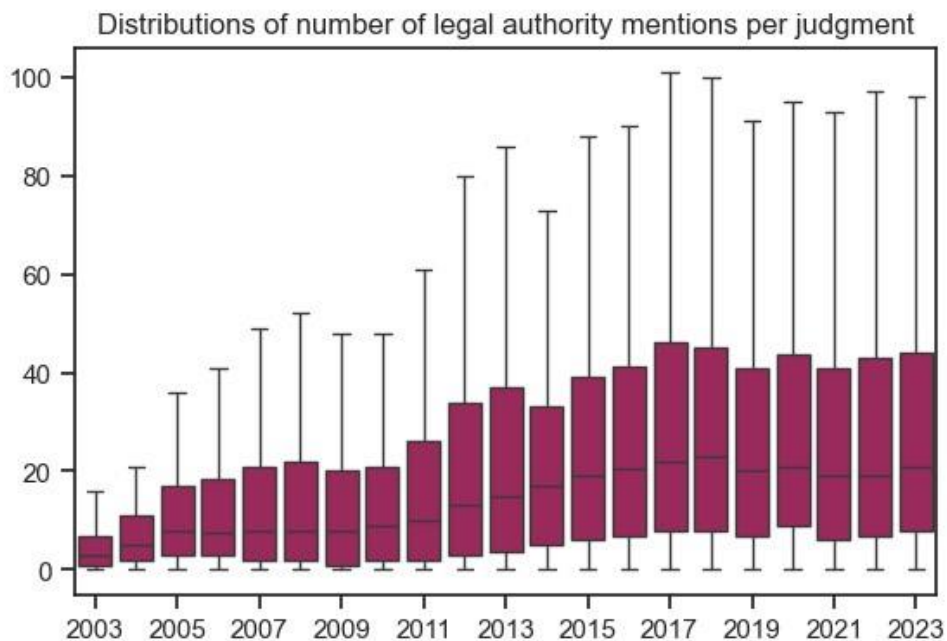


Figure 9: Distributions of legal authority mentions

one of the distinctive features of Justice Choo Han Teck's writing style is that he tends to write short judgments. The median judgment length for his 840 judgments from 2003 to 2023 is a mere 1,751 words, and he has written judgments that range many subject areas throughout his career on the Supreme Court bench. But even then, Justice Choo's longest judgment weighs in at a hefty 21,450 words, which tends to suggest that the Supreme Court judges write according to what is required.

I will present instead one related result based on the judges' data, which can be seen at Figure 10. I measured the annual rate at which judges produced judgments and plotted those against the median judgment lengths for every judge in a scatterplot. The plot clearly shows a negative correlation between them: when judges write longer judgments, they also write fewer judgments. The reverse is also true, when judges write more judgments, they also write shorter judgments. This may point to an upper limit on how many words a judge can produce in writing over a given period.

I point this out because there may be implications that could bear on the discussion in the next section. A judge who hears back-to-back cases that were bitterly-disputed and highly complex may be working under a lot of pressure when they not only have to decide difficult questions but are also expected to produce written judgments of a certain level of detail for these

cases. Under such constraints, it may not always be possible for judges to edit their written work to reach a wider audience. With this in mind, I will be discussing the readability measurements next.

4.5 Are Singapore judgments difficult to read?

Word statistics I begin by discussing the word-related statistics because the results suggest that word usage in the judgments has been relatively stable and unchanged in the past 20 years, with the exception of some special cases: such as 'plaintiff' being replaced with 'claimant'. Perhaps 20 years is a too short a period to observe word usage changes in a specialised form of writing such as judgments. In any event, these are the findings.

Word length The average word length in judgments (excluding quotations) was consistently between 4.6 and 4.7 characters over all years. This is comparable with the average word length of 4.7 for the Straits Times articles I downloaded.

Text coverage The ratio of high frequency words in judgments is consistently around 85-86% over all years. If medium frequency words are

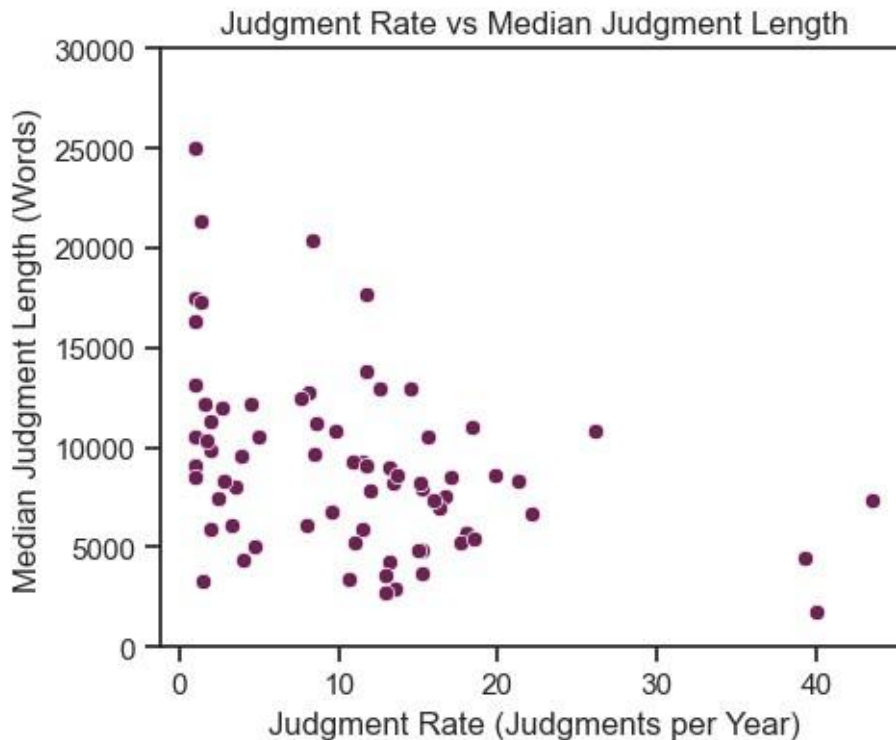


Figure 10: Caption

counted (Zipf value above 3), text coverage rises to around 91%. In comparison, the Straits Times articles have a text coverage of 86% for high frequency words and 93% if medium frequency words are included.

Legal jargon Table 2 shows the most commonly used Latin terms and low frequency words³³ in judgments. Most of the Latin terms have ready replacements in plain English without sacrificing any meaning. In fact, the Rules of Court 2021 has already replaced *ex parte* with “without notice”. On this note, the changes introduced in the Rules of Court 2021 appears to have had an effect as the word “plaintiff” has decreased in use while “claimant” has enjoyed a recent increase in numbers: see Figure 11.

Sentence lengths The average sentence length is 30.85 for Court of

³³ I removed the following words because there is likely to be sufficient context in the judgment for the reader to infer its meaning: defendant, plaintiff, appellant, respondent, applicant, counsel, claimant, filed.

Appeal judgments and 27.54 for High Court judgments. The plotted lines

Low Frequency	Count	Latin Phrase	Count	Archaic Word	Count
submit*	80,656	<i>ie</i>	23,271	thereafter	7,689
jurisdiction	20,113	<i>supra</i>	7,788	the said	5,012
arbitration	20,011	<i>inter alia</i>	7,511	thereby	3,953
transaction	19,827	<i>prima facie</i>	5,226	foregoing	2,847
imprisonment	19,208	<i>viz</i>	5,072	therein	2,682
pursuant	18,785	<i>per</i>	4,354	thereof	2,302
affidavit	17,621	<i>eg</i>	3,313	whereby	1,181
cl	16,902	<i>bona fide</i>	1,991	aforesaid	804
accordingly	15,659	<i>res</i>	1,802	thereon	757
matrimonial	14,317	<i>ex parte</i>	1,610	hereafter	752
creditor	13,467	<i>sic</i>	1,576	thereto	670
contractual	12,270	<i>res judicata</i>	1,480	wherein	667
liable	12,134	<i>per se</i>	1,386	thereunder	645
testified	11,709	<i>mens rea</i>	1,301	therefrom	593
disputed	11,369	<i>locus</i>	1,205	hereinafter	529
oral	11,082	<i>locus standi</i>	1,065	whereupon	430
provision	10,869	<i>in rem</i>	1,059	herein	401
obligations	10,434	<i>forum non conveniens</i>	725	whereabouts	327
commenced	10,243	<i>ibid</i>	722	hereof	123
cited	9,821	<i>id</i>	704	therefor	100
applicable	9,262	<i>ad hoc</i>	595	thereupon	95
contended	8,617	<i>in pari materia</i>	594	hereby	87
subsequent	8,468	<i>de facto</i>	591	therewith	76
incurred	8,467	<i>qua</i>	562	wherewithal	44
paragraph	8,333	<i>ipso facto</i>	525	hereto	35

Table 2: Top 25 Low Frequency Words, Latin Phrases and Archaic Words

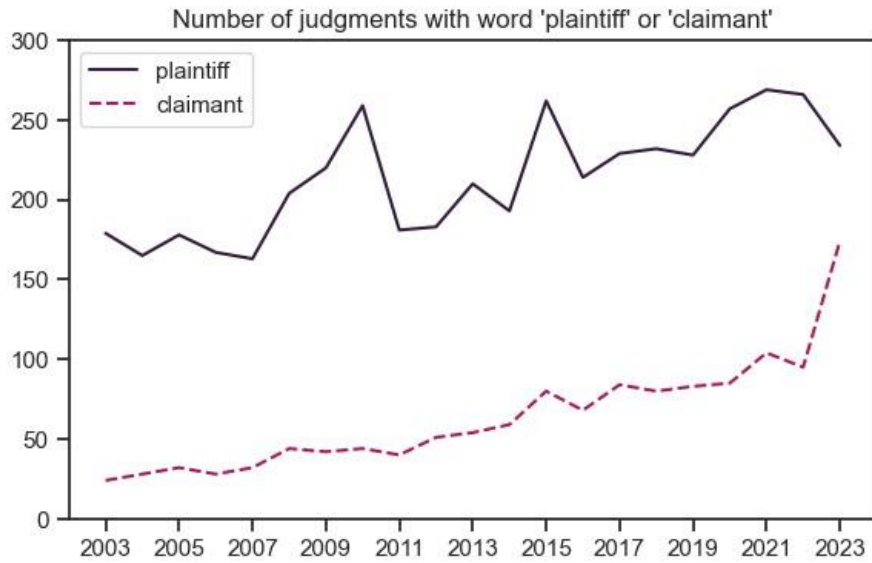


Figure 11: Number of judgments with the word “plaintiff” or “claimant”.

in Figure 12 and Figure 13 appear highly correlated because the Flesch-Kincaid score is essentially a function of word length and sentence length. Since average word length has remained stable, the key contributor to the fluctuations in Figure 13 is sentence length.

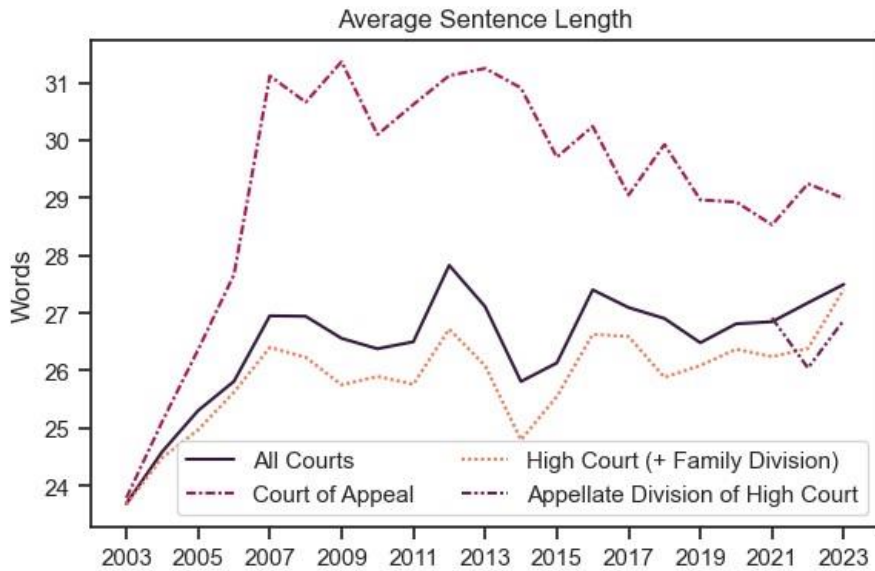


Figure 12: Average sentence lengths by court level.

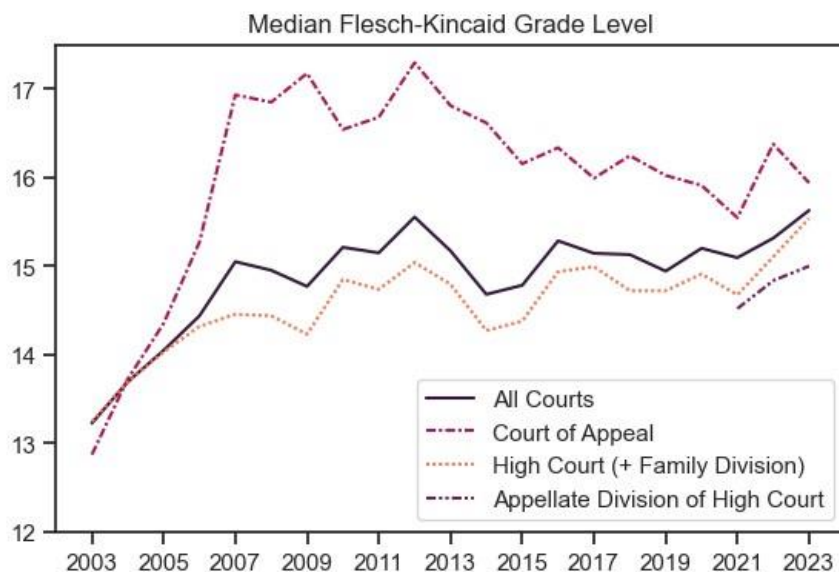


Figure 13: Median Flesch-Kincaid Grade Levels by court level.

Headings It is now customary to find headings in written judgments unless it is a relatively short one (see Figure 14).

Most of the judgments have a broadly consistent outline.

Firstly, the *Introduction* section may or may not be explicitly labelled as such. This is commonly followed by a *Facts* or *Background* heading. If the judgment concerns an appeal, there may be a section heading on *Procedural history* or the *Lower court's decision*.

The headings for the later sections in the judgment are not as consistent. This is mainly because it is increasingly common for judges to explicitly organise their decision or discussion by the issues, with case-specific headings for each issue or as sub-headings under a *Decision* heading.

The common headings can be broadly grouped as: *Parties' Cases*, *Issues*, the *Analysis* or *Decision*, followed by the *Conclusion*, which would contain a summary of the court's orders.

Explicit connectives Figure 15 shows the ratio of paragraphs and sentences that begin with a word or phrase in the list of explicit connectives. The use of connectives to mark transitions in between sentences and paragraphs have both increased.

First person pronouns. The use of first person pronouns is fairly common: they appear at a rate of slightly more than once every 200 words

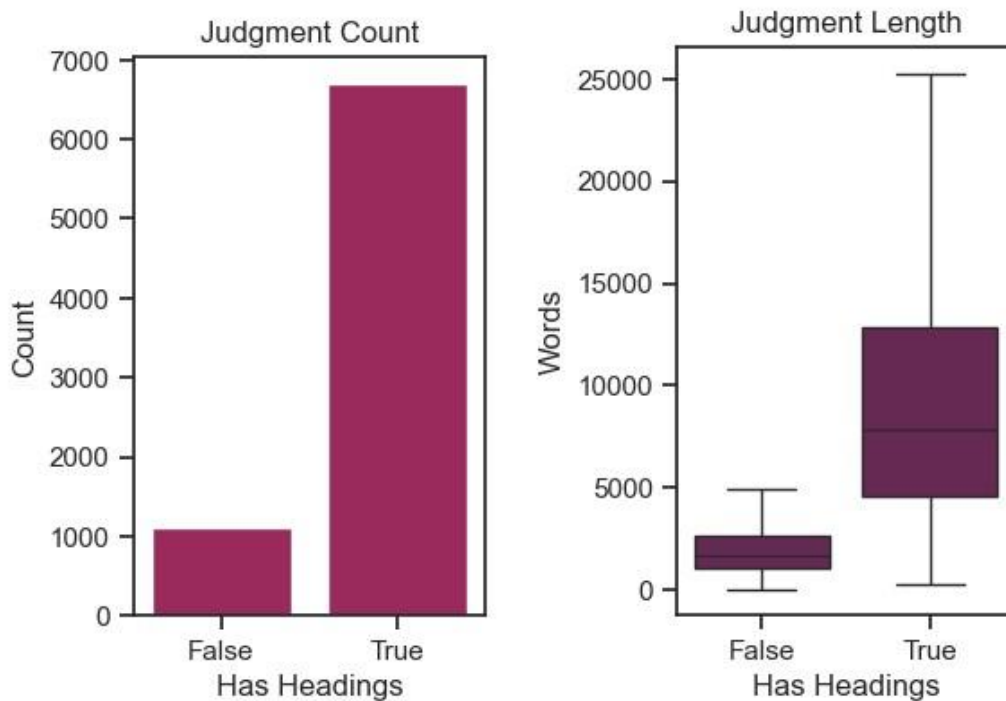


Figure 14: Judgments Without Headings and Judgments With Headings

(or half a page). There was no discernible difference in the use of first person pronouns in judgments over time. I also did not observe any pattern among individual judges.

Second person pronouns I found only two judgments in which the judge used second person pronouns that were not a quotation. Both were criminal cases.

The first case is *Public Prosecutor v Pramanik Liton* [2017] SGHC 110 in which Justice Choo Han Teck explained directly to the accused his reasons for finding him guilty of rape. Perhaps this was done because the accused was not legally represented and Justice Choo felt it was appropriate to address the accused directly instead of referring to him in the third person in his judgment.

The other case is *Public Prosecutor v DAM* [2023] SGHC 265. Save for a brief foreword, Justice Aedit Abdullah directed his sentencing remarks

entirely at the accused person, with some force, for physically abusing his own young daughter to the point of death. And in a moment of compassion and regret, Justice Aedit seemed to also address everyone else reading his words. This is what he said (at [4]):




Figure 15: Ratio of paragraphs and sentences beginning with connectives.

“It is important to my mind that we all remember her by her name, and not by cold impersonal nouns, such as the “deceased” or “the victim”. I fear Umaisyah will only be remembered as the child who was killed by her father; whose body was burnt by her parents; and whose charred remains were kept in a pot by them. Umaisyah deserved so much more.”

The troubling facts of the case may have occasioned this departure from the norm. Apart from the terrible abuse she suffered, Umaisyah’s undignified death had not been discovered for five years. Here, the words in the judgment were intended to condemn the accused person’s actions, restore some dignity to Umaisyah as a person, and to make a plea to humanity.

4.6 Comparison with Straits Times articles

Having discussed and compared judgments internally, how does the average Supreme Court judgment compare with other forms of written texts?

To compare against the downloaded Straits Times articles, I decided against using the entire judgment corpus. Instead, I filtered for judgments from the following subject areas that overlap with the areas under “community law”.³⁴ The point of this comparison is to use Singapore’s most-read English daily newspaper as benchmark for the judgments covering subject areas that have more direct impact on the average Singaporean.

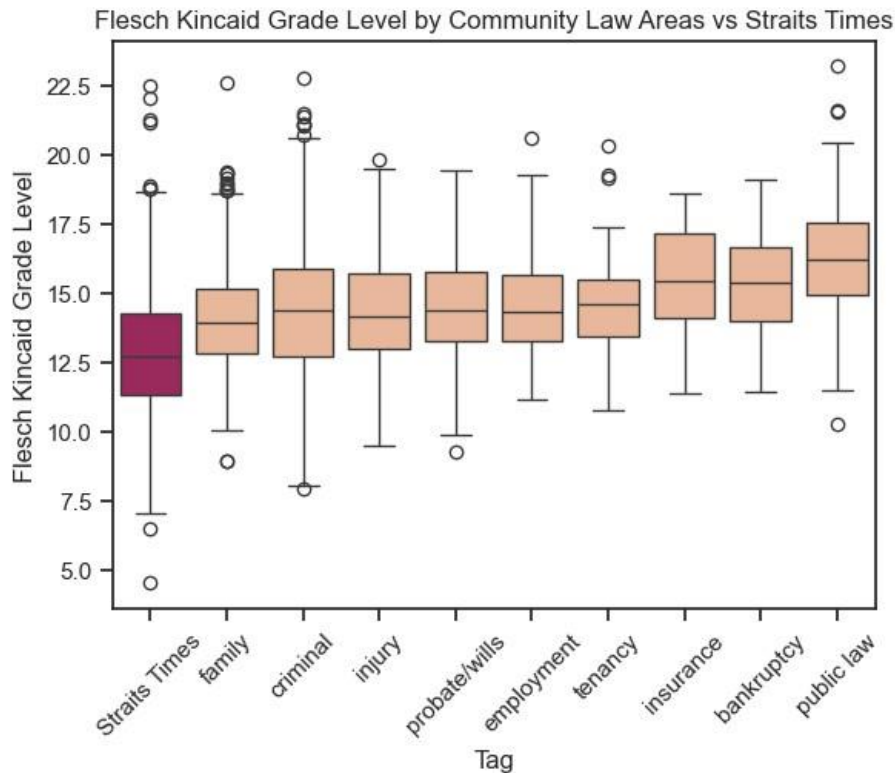


Figure 16: Flesch-Kincaid Grade Levels.

Here we see in Figures 16 and 17 that the judgments do not score as well as the Straits Times articles on the Flesch-Kincaid Grade Level score and coverage of familiar words, except in Family Law cases, which have better text coverage, on average.

³⁴ <https://www.lawsociety.org.sg/find-lawyer/community-law/>

4.7 Interesting Openings

Interesting openings can take many forms and there are more than a few gems that have been written. There are too many to include here but below are a few notable examples.

One common device is to quote well-known literature such as Shakespeare's works or classics like *Alice in Wonderland* and to use the familiarity of literary characters to engage the reader's attention. Take for example, the opening paragraph of *NCL Housing Pte Ltd v Sea-Shore Transportation Pte Ltd* [2021] SGHC 29 that dealt with the issue of a guarantor's liability:

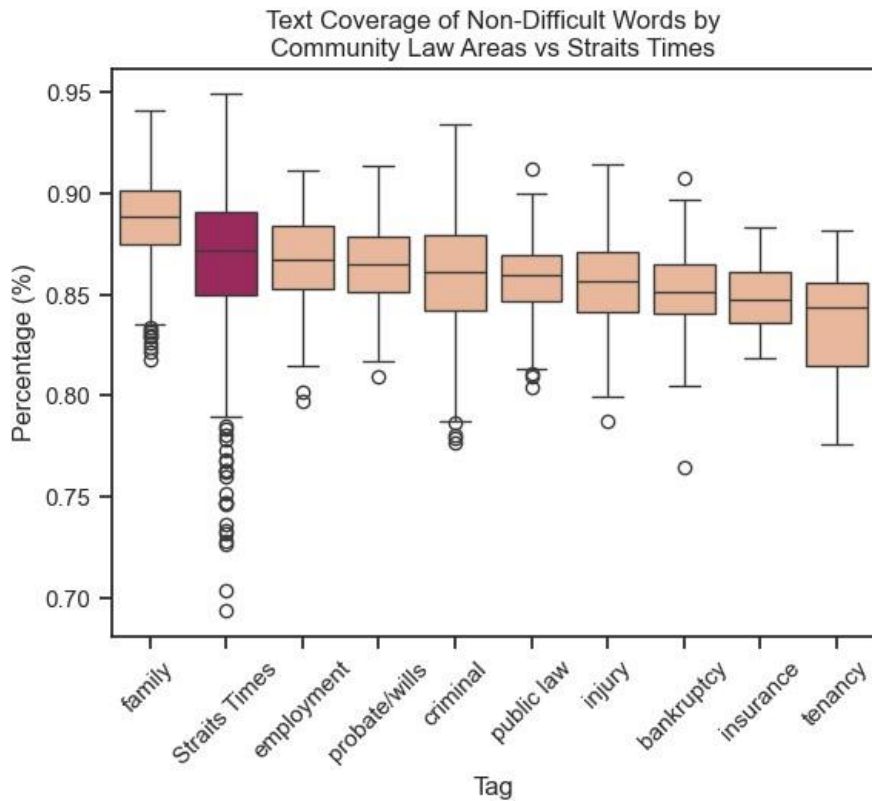


Figure 17: Text coverage ratios.

The Merchant of Venice warns of the perils of standing guarantor for the debt of another. As guarantor, Antonio boldly promises a pound of his flesh. When that is demanded of him, he does not shy away from the promise he made, though it might cost him his life.

What saves him, instead, is the interpretation placed on the guarantee.

The opening from *First Asia Capital Investments Ltd v Soci'et'e G'en'erale Bank & Trust* [2017] SGHC 78 sets the context with a well-known event and uses the informal and vernacular, injecting dark-humour and raising the stakes of the case for the reader:

Following the Global Financial Crisis in 2008, many investors incurred substantial losses to their equity portfolios. One particular financial product, known as a share "accumulator", was the source of many investors' financial woes during that period, so much so that it earned itself the infamous epithet, the "I kill you later" contract.

References to food, a pet topic of Singaporeans, also lends a relatable quality to this opening in *Consorzio di Tutela della Denominazione di Origine Controllata Prosecco v Australian Grape and Wine Incorporated* [2023] SGCA 37:

Much of the food that is consumed in Singapore is produced abroad. Generally, the consumer has little interest in the provenance of the produce that ends up on the dinner table. This, however, is not always the case. Many would say that the limestone-permeated water of Ipoh lends a special quality to the kway teow produced there. Also famous is the pork that comes from the Black Iberian pig which has been raised on the Iberian Peninsula and fed a diet of acorns grown there. In both these examples the special qualities of the landscape (or terroir to use the French phrase) have given the food product a unique flavour or quality which consumers actively seek out.

Other openings are written as a story-like narrative of the facts. The one in *Huatong Contractor Pte Ltd v Choon Lai Kuen (trading as Yishun Trading Towing Service)* [2020] SGHC 129 is replete with increasing tension and even contains a clear statement of the issue before the court:

Most of the journey from Fort Road to Benoi Crescent is conducted on the Ayer Rajah Expressway ("the AYE"). The route is relatively flat, gradient-wise, and has few traffic lights or bends

to negotiate. Most drivers can expect an uneventful journey, especially on a late Friday afternoon. This was exactly Mr. Ong Sen Lian's ("Ong") experience as he operated a concrete pump truck ("the Truck") which was towed across the AYE from Fort Road to Benoi Crescent on 20 July 2012. As he neared Benoi Crescent, he found smoke emitting from the Truck. Before long, the Truck's chassis and tyres had caught fire. The present action centred on the nature and cause of this fire.

Another vivid example of how a dispute can be distilled to its core characters and the resulting dilemma can be found in the opening to *Pertamina Energy Trading Limited v Credit Suisse* [2006] SGCA 27:

A fraudulent employee causes his company substantial losses in an audacious scam. An eager bank officer fails to heed alarm bells in his zeal to facilitate and sustain a banking relationship with that company. Who should bear the losses?

Describing a party's struggle in mythical proportions is also a way to gain the reader's attention, as was done in *Ong Jane Rebecca v PricewaterhouseCoopers* [2011] SGHC 146:

Jane Rebecca Ong ("the plaintiff") is the modern day Odysseus. She first arrived in the courts here in 1991 when she sued her former husband and his mother for various claims including the fraudulent transfer and concealment of assets to which she had a share. The present proceedings, taking place about 20 years later, seem remote to the Originating Summons she filed in 1991 but is connected to and arose from that action.

Finally, some openings place the dispute within the wider legal or social context, effectively highlighting the broader implications. So in *Ong Kian Peng Julian v Singapore Medical Council* [2022] SGHC 302,

The medical profession is held in high esteem because its members have been called to the work of healing. Those who avail of the services of medical professionals must repose trust and confidence in them. This is essential, not least because it is necessary to enable a frank and open exchange of information. But this engenders an expectation on the part of patients that they will be treated with dignity and respect, and that any information they provide will be used for proper purposes. It is

true that doctors have a life outside their profession. In general, they are not to be punished for moral failings in their personal lives. But what happens when the line between a doctor's personal and professional life is obscured? That is the question presented in this case.

5 Conclusion

In conclusion, this study has shown that Singapore judgments, like most legal texts, are still not very accessible to the wider public for a number of reasons. Firstly, they tend to contain more unfamiliar or difficult words, and have longer sentences on average when compared to newspaper articles on similar topics. With the changes to the ROC 2021, it is likely that some legal jargon, especially Latin phrases, will decline in usage. At the document level, most judgments are organised well with clear headings and use explicit connectives to help readers transition between sentences and paragraphs. While judgments have generally increased in length over time, this should not be seen as an unwelcome development because this has arrived with an increase in legal reasoning. Coupled with the fact that judges tend to write more judgments on topics that are more directly relevant to the average Singaporean, the result may be increased clarity and certainty in the application and understanding of the law in these areas.