

RESEARCH ARTICLE

The impact of the #MeToo movement on language at court A text-based causal inference approach

Henrika Langen *

Department of Economics, University of Fribourg, Fribourg, Switzerland

* rikalangen@t-online.de

Abstract

This study assesses the effect of the #MeToo movement on the language used in judicial opinions on sexual violence related cases from 51 U.S. state and federal appellate courts. The study introduces various indicators to quantify the extent to which actors in courtrooms employ language that implicitly shifts responsibility away from the perpetrator and onto the victim. One indicator measures how frequently the victim is mentioned as the grammatical subject, as research in the field of psychology suggests that victims are assigned more blame the more often they are referred to as the grammatical subject. The other two indices designed to gauge the level of victim-blaming capture the sentiment of and the context in sentences referencing the perpetrator. Additionally, judicial opinions are transformed into bag-of-words and tf-idf vectors to facilitate the examination of the evolution of language over time. The causal effect of the #MeToo movement is estimated by means of a Difference-in-Differences approach comparing the development of the language in opinions on sexual offenses and other crimes against persons as well as a Panel Event Study approach. The results do not clearly identify a #MeToo-movement-induced change in the language in court but suggest that the movement may have accelerated the evolution of court language slightly, causing the effect to materialize with a significant time lag. Additionally, the study considers potential effect heterogeneity with respect to the judge's gender and political affiliation. The study combines causal inference with text quantification methods that are commonly used for classification as well as with indicators that rely on sentiment analysis, word embedding models and grammatical tagging.

OPEN ACCESS

Citation: Langen H (2024) The impact of the #MeToo movement on language at court A text-based causal inference approach. PLoS ONE 19(5): e0302827. <https://doi.org/10.1371/journal.pone.0302827>

Editor: Satabdi Mitra, KPC Medical College and Hospital, INDIA

Received: September 15, 2023

Accepted: April 10, 2024

Published: May 15, 2024

Copyright: © 2024 Henrika Langen. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data is available from the CourtListener database. The processed data and the code for the analysis are available under the following link: <https://doi.org/10.5281/zenodo.10428646>.

Funding: The author received no specific funding for this work.

Competing interests: The author has declared that no competing interests exist.

Introduction

The present study evaluates how the #MeToo movement affected the way sexual offenses are handled within the U.S. justice system by analyzing the language used in judicial opinions published in U.S. state and federal appellate courts between 2015 and 2020. It examines the movement's impact by means of a Difference-in-Differences (DiD) and a panel data-based event study approach. In the DiD approach, the development of language in opinions on sexual offenses is compared to that in opinions on other crimes against persons. Meanwhile, the

panel event study approach assesses how the development of language in judicial opinions published by different judges changed as a result of the #MeToo movement.

This study develops novel text indicators aimed at measuring the extent to which actors in courtrooms employ language that implicitly shifts responsibility away from the perpetrator and onto the victim. The primary indicator measures the frequency with which the victim is mentioned as the grammatical subject, drawing upon psychology research which suggests that victims tend to be assigned greater blame when they are often mentioned as the grammatical subject. (see e.g. [1, 2]) Additionally, the study develops two other indicators to capture the development of victim-blaming language, which build on the word2vec algorithm and the SentiWordNet 3.0 lexicon [3] respectively, as well as an approach to use text vectorization methods typically used for text categorization in order to quantify evolution of language over time. The study then demonstrates how the different text quantifiers can be integrated into a DiD and an event study approach in order to estimate the causal effect of the #MeToo movement on language in judicial opinions. In doing so, this study contributes to the recently growing literature on text-based causal inference.

The study proceeds as follows. Section reviews the current state of literature on text analysis, particularly the emerging field of text-based causal inference, and provides background information on the #MeToo movement and the U.S. justice system. Section describes the corpus of judicial opinions examined in this study. The following section, Section, illustrates how the judicial opinions are quantified by means of various indicators and text vectorization approaches. Section outlines the identification strategy underlying this paper and discusses the assumptions necessary to identify the impact of the #MeToo movement using a DiD model and a panel data-based event study approach. The subsequent section () summarizes the results and Section concludes.

Background

Text-based causal inference

While there are several socio-economic and legal studies on text classification based on Natural Language Processing (NLP) and also some NLP-based analyses on language development, literature on text-based causal inference is scarce. The social science literature on text classification ranges from studies on differences in the linguistic style of posts in different online communities [4] and of comments on #MeToo articles in different news outlets [5] to studies that develop classifiers for political speeches in order to predict the speaker's ideology [6] or identify their sentiment towards the topic discussed in the speech [7]. In the legal domain, [8] developed a document classifier for judicial opinions from U.S. circuit courts that classifies the opinions according to the predicted ideological direction (conservative vs. liberal) of the decision. In addition, there are several studies on classifying legal documents by topic (see, e.g., [9–11]).

The evolution of language over time has been analyzed using quantifiers for the context in which words are used (see, e.g., [12–14]) as well as based on indicators for the sentiment of that context (see, e.g., [3, 15]). [16] analyze how new members in a medical forum gradually adapt their language to the forum's linguistic standards during their first year of forum participation.

Literature on text-based causal inference has emerged only in recent years. Some studies integrate NLP elements into causal inference for confounding adjustment (see [17] for a review). [18], e.g., develop a framework for estimating treatment effects which combines text-based matching and confounding adjustment based on text. They estimate how a scholar's gender affects the number of citations of their publications, while controlling for and matching

based on the content of publications. Other studies that rely on text-based confounding adjustment include those by [19, 20]. [21, 22] propose a text matching approach based on distance metrics rather than text classification. [23] integrate text classifiers into causal inference in order to tackle problems with missing data and measurement error.

Other studies use text as treatment or outcome. [24] analyze how a judge's score on an indicator of gender-stereotyped language affects their decisions on women's rights' issues. [25] assess how wording in tweets affects the number of re-tweets. To do so, they apply different text vectorization approaches and quantify wording by means of sentiment, lexical distinctiveness and readability indicators, in order to identify words that in-/decrease re-tweet propensity. Similarly, [26–30] evaluate how to estimate causal effects of wording, semantics and lexical choices in texts.

[31] develop a sample splitting framework for estimating treatment effects with text as outcome, building on the classification of outcome texts based on a model trained in the training sample. [32] assesses how troll activity that promotes a pro-government agenda on Russian social media affects the evolution of online discussions using a regression discontinuity approach. To do so, he models the development of conversations on social media as changes in the mixture of topics with topics identified through NLP-based classification. Other examples of studies with texts as outcome include a study by [33] on how Reddit's 2015 anti-harassment policy affected the usage of hate speech, a paper by [34] examining the influence of social norms on the rise of hate speech following terrorist attacks, as well as an analysis by [35] on the effect of tagging articles as not written in a "neutral point of view" on the development of lexical patterns in the labeled articles.

The present study contributes to this expanding field of text-based causal inference by integrating different text quantification approaches with causal inference methods. It develops a novel indicator drawing from psychology literature on the relationship between grammatical structure and perceptions of responsibility, while also employing established text quantifiers such as sentiment analysis and word2vec. Additionally, it proposes an approach to use text vectorization for quantifying language development within a text corpus for later integration into causal inference. In essence, it provides a diverse set of approaches to quantify text that can then be incorporated as outcomes in causal analyses.

The #MeToo movement and its societal, cultural and political impact

After starting as an online campaign against sexual harassment, #MeToo soon evolved into a movement that led to extensive and sustained media coverage of the prevalence of sexual violence in society. It prompted discussions about abuse of power, rape myths and the importance of supporting victims of sexual assault.

The phrase "Me Too" was initially coined by social justice activist Tarana Burke who began using this phrase in 2006 to campaign for the empowerment of sexual violence victims, particularly among women of color. It gained widespread attention in late 2017 after sexual misconduct allegations against Harvey Weinstein were exposed in a New York Times exposé by [36] published on October 5th, 2017. On October 15th, the actress Alyssa Milano tweeted a post encouraging victims of sexual harassment and assault to come forward by using the hashtag #MeToo. On Twitter, the hashtag was tweeted about 300,000 times on the day after Milano's post, reached a peak of 750,000 tweets within 24h and was used on average more than 55,000 times per day during the year following the initial tweet [37]. On Facebook, the #MeToo conversation peaked at 4.7 million participating users within 24 hours, who engaged with over 12 million posts, comments, and reactions [38].

The discussion on social media grew into a movement that led to protest marches in the U. S. and around the globe, resulted in extensive and sustained media coverage of the issues of sexual violence and abuse of power, and shaped the public discourse in the months following October 2017. [39] estimate that in the first 8 months after the movement's emergence, the number of google searches on sexual harassment and assault exceeded the expected amount by some 86%. According to the Women's Media Center, the number of articles on sexual assault and harassment in a sample of 14 leading U.S. newspapers was more than double the pre-#MeToo average in November 2017 and still exceeded the pre-#MeToo average by 30% some 10 months after the onset of the movement [40]. [41] named the "Silence Breakers"—victims of sexual harassment or assault who came forward and thereby started the global dialogue on sexual violence—as its 2017 "Person of the Year".

About 65% of social media users report having regularly encountered at least some content related to sexual harassment or assault on social media platforms in the months following the start of the #MeToo movement, with little difference across demographic groups [37]. The observation that a large share of society has been confronted with the issue of sexual violence is also reflected in Google search data. Following [39, 42], I look at how often the terms "sexual assault" and "sexual harassment" were searched for on Google in the U.S. during the study period. Fig 1 shows that public interest in these topics has never been higher than at the start of the #MeToo movement, with a similarly pointed peak of searches on Google News. Following [42], Fig 1 further shows that the #MeToo movement, rather than the Weinstein scandal,

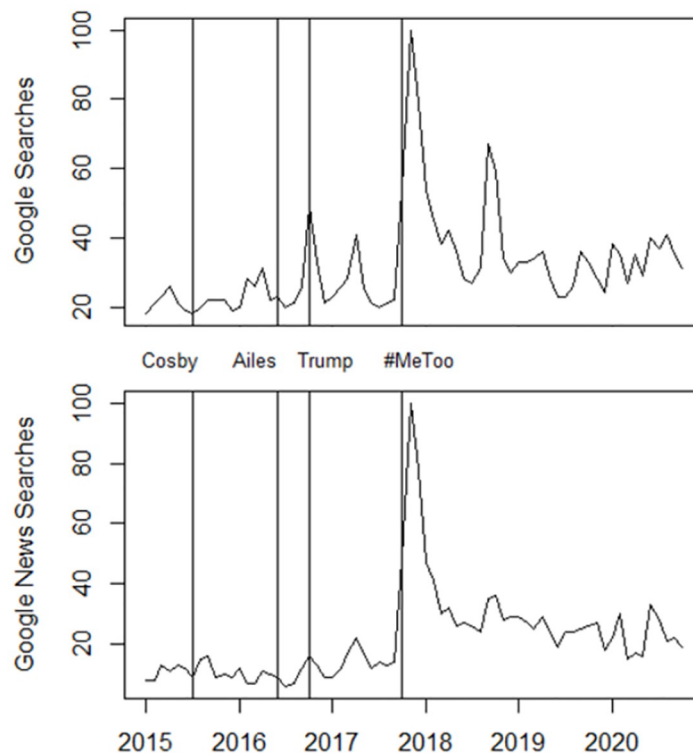


Fig 1. Google searches for “sexual assault” & “sexual harassment”. Searches on Google and Google News in relation to the highest point in the period January 2015 to November 2020. The vertical lines indicate when the sexual harassment/assault allegations against stand-up comedian Bill Cosby, former U.S. President Donald Trump and Fox News CEO Roger Ailes became public.

<https://doi.org/10.1371/journal.pone.0302827.g001>

generated the rise in searches, as scandals involving men with similar or even higher profile than Harvey Weinstein did not attract nearly as much public attention as the #MeToo movement.

The #MeToo movement has led to changes in attitudes and responses to sexual violence, but also provoked some disillusionment and backlash. While several surveys indicate that both men and women have observed others to avoid sexually harassing behavior in response to #MeToo in public or private settings [43–46], others suggest that employers anticipate or observe increased caution among men in their interactions with female colleagues, potentially hindering women's professional advancement (see, e.g., [47–51]). Furthermore, several articles argue that #MeToo (almost) exclusively benefited affluent white women, while women from lower socioeconomic backgrounds and women of color struggle to gain public attention and support (see, e.g., [52–54]), a criticism also supported by several studies (see e.g. [55, 56]). However, different studies [42, 57, 58] suggest that the under-representation of victims from certain demographic groups does not necessarily translate into a lesser impact of the movement on those groups.

The societal, cultural and political impact of the #MeToo movement has been analyzed extensively with mixed findings about the movement's influence. There are studies showing that the #MeToo movement did not significantly increase self-reported interest in political participation [59], that the portrayal of male entrepreneurs in Swedish media changed only marginally after the onset of the movement [60], that the #MeToo movement caused a significant increase in the propensity of female dating platform users in South Korea to decline dating requests [61] and that it induced a change in gender norms in Swedish-language tweets [62]. A study by [63] reveals that support for the #MeToo movement is associated with a stronger belief in the sexual misconduct allegations against Donald Trump among Democrats but not among Republicans. [64] examined how the movement affected women's perception of safety by comparing the development of men's and women's perceptions of safety in subway stations with a DiD approach, finding a significant decline in women's perceptions of safety after #MeToo.

In the study most relevant to the present one, [42] assess the effect of the #MeToo movement on the propensity to report a sexual offense to the police. They apply a triple-difference approach over time, across 31 OECD countries, and between sexual and non-sexual offenses. For countries where the #MeToo movement attracted a great deal of attention, they find that it caused a significant 10% increase in the number of reported sexual offenses during the first six months after the movement began. A DiD analysis comparing the development of reported sexual and non-sexual crimes in the U.S. reveals a similar long-term effect for the 15 months following Alyssa Milano's tweet. The study further shows that the movement did not affect the reporting of all types of sexual assault equally: reports of rape and fondling increased, while reports of statutory rape and sodomy remained unaffected. Finally, the study assesses the impact of the movement on sexual offense arrests, finding that the increase in reported sex offenses in the U.S did not bring about a similar surge in the number of arrests. Rather, the authors estimate that the movement raised the number of arrests by only about 6%, which they attribute to the fact that the movement had a stronger effect on cases with a low likelihood of arrest due to delayed reporting, less severe offenses, or insufficient evidence. The study does not provide any conclusions about the effect of the movement on convictions.

Victim blaming

Besides raising awareness about the prevalence of sexual harassment and assault in society, the #MeToo movement has also fueled discussions about the acceptance of rape myths. For this

reason, this study assesses the judicial opinions specifically with respect to the reinforcement of rape myths and victim blaming.

The term “rape myths” refers to stereotypical and inaccurate beliefs about sexual assaults that are prevalent in society. These myths are frequently used to redirect blame from the perpetrator to the victim and to downplay the seriousness of the incident. Rape myths may not always be explicitly stated; they can also be reinforced through rhetorical techniques or by highlighting particular details when discussing sexual assault. For example, studies show that sexual violence reports that emphasize external circumstances and the victim’s behavior can increase the likelihood of readers to accept rape myths and view the victim as partly responsible for the assault (see, e.g., [65–67]).

Victim blaming and reinforcement of rape myths can be observed in a variety of contexts, albeit with varying degrees of explicitness. While often more bluntly expressed in social media and other informal contexts (see, e.g., [68, 69] for studies on identification of victim-blaming language in informal contexts), victim blaming and reinforcement of rape myths are also present in traditional media (see, e.g., [66, 70, 71]) and within the justice system. In court, there are general victim rights and sexual assault-specific regulations, such as rape shield laws, that to some extent prohibit explicit victim blaming, stigmatization, and stereotyping. Yet, several studies, based on testimonies from various actors, such as victims, barristers, sex crime investigators and independent observers, suggest that victim blaming and rape myths are still prevalent in courts (see, e.g., [72–75]). Although it is often the defense that introduces rape myths for strategic purposes, some judges do not always intervene and may even endorse such arguments [72, 76]. In addition, many judges tend to employ terminology of affection and consensual sex in cases where the perpetrator is familiar to the victim [76].

The U.S. justice system

The legal documents examined in this paper are judicial opinions from 51 U.S. state and federal appellate courts. Appellate courts review legal cases that have already been heard in a lower court (trial court) after a party appeals the trial court’s decision. In civil cases, both parties have the right to appeal the trial court’s decision, while in most states, only the defendant can appeal in criminal cases. To appeal a decision, the appealing party (the appellant) must file a brief, i.e., a written argument setting forth the facts and arguing why the trial court’s decision was erroneous, to which the other party (the appellee) must respond with an appellee’s brief.

The appellate court does not usually admit new evidence or witnesses; it may rule solely on the basis of the written briefs or after hearing oral arguments. The appellate court often issues what is called a judicial opinion, i.e., a written decision outlining the court’s reasoning; it is usually written by a single judge and reviewed by the other judges on the panel. If a judge disagrees, they can issue a dissenting opinion. Judges who disagree with the reasoning but agree with the result can issue a concurring opinion. Sometimes, judges issue an unsigned opinion called a *per curiam* opinion [77].

Text data

The judicial opinions used in this study were obtained from CourtListener, an archive of court data operated by [78]. The CourtListener database collects judicial opinions issued by state and federal courts from various sources. The body of opinions for this study is restricted to precedential opinions from state and federal appellate courts. The reason for this choice is that the corpus of precedential opinions available in the CourtListener database is complete for the available courts, unlike that of non-precedential opinions, i.e., by restricting the sample to precedential opinions, no selection problems arise. In addition, precedents reflect developments

in courts and case law. As integral components of the body of law, they contribute to the continuous development of the legal system.

The body of judicial opinions

The body of judicial opinions examined in this paper includes opinions from 51 courts. Most judicial opinions of state and federal appellate courts have a similar structure. They usually begin with a summary of the trial court's ruling, followed by the arguments of defense and prosecution that were presented and admitted before the trial court, as well as the appellate brief filed by the defense. They typically conclude with reasoning and the decision of the appellate judges. Thus, the opinions reflect the atmosphere during the trial court hearing and the appeal proceedings, as well as the attitude of the defense and the judges towards the victim.

To single out the opinions on sexual offenses and those on other crimes against persons from the full body of precedential opinions, I take advantage of the fact that the opinions have a similar structure and usually begin with an introduction specifying the offense(s) for which the offender was convicted in trial court. The opinions can therefore be reliably classified based on term search in the introduction. To distinguish opinions on sexual-violence related cases and those on other cases of interpersonal crimes, I rely on regular-expression-based (`regex`-based) search of the legal terms for such crimes in the introduction of the respective opinion (for more information on the legal terms used, see [S1 Appendix](#)).

As the headers of the opinions differ between courts, their divisions and even the judges who author the opinions, the introductions cannot be reliably identified based on `regex` rules. Therefore, the first third of each opinion but no more than 5,000 characters are defined as the introduction. The introductions of all opinions available at CourtListener are searched for matches to legal terms related to sexual violence; the remaining opinions are then searched for matches to other interpersonal-violence related legal terms. While certainly not adequate to accurately classifying opinions into crime categories, the identification procedure described above ensures that opinions are selected into the sample based on the same rules throughout the observation period. Focusing on the introduction of opinions prevents erroneously classifying opinions to one of the two groups when an opinion mentions crimes from a party's past.

Using the `regex` procedure described above, I can identify 43,088 opinions on crimes against persons published between January 2015 and November 2020, including 15,307 on sexual offenses and 27,781 on non-sexual offenses against persons (performance assessment in [S1 Appendix](#)). The number of opinions per court is provided in [S1 Table](#).

Text quantification

In order to assess the impact of the #MeToo movement on language in court, the judicial opinions need to be quantified. The quantifiers described in Section aim at quantifying the amount of victim blaming in each opinion, while the text vectorization methods outlined in Section are later used to capture the general evolution of language in judicial opinions.

Victim-Blaming indicators

To quantify the extent of victim blaming in judicial opinions, three indicators are constructed based solely on sentences in which the victim or perpetrator is named. These indicators aim at capturing the extent to which opinions contain wording that implicitly shifts some blame from the perpetrator onto the victim, where such wording may come from the defense, the prosecution, or the judges involved in the case. The first indicator represents a novel approach based on findings from the field of psychology and will be the primary focus here. The remaining two indices assess the context in which the victim and perpetrator are discussed, where the

first one measures the sentiment of the context, while the second one more comprehensively assesses the linguistic environment in which they are referenced.

In the appellate court opinions, the victim is often referred to by their name or by the term “victim”. The appealing party, i.e., the person found guilty in trial court, is usually called the “appellant” or by their name, but may also be referred to as the “petitioner”. As the victims’ and the appellants’ names are not clearly identifiable, only sentences containing the words “victim”, “appellant” and “petitioner” as well as inflected forms of these words are considered in the construction of the indicators (As the gender of victim and offender cannot be unmistakably identified, personal pronouns cannot be used to construct the indicators). If the names of victim and offender were identifiable in all texts, replacing the occurrences of the names with the corresponding roles of the individuals (victim/offender) and then applying the approaches below would certainly add substantial value to this study. Similarly, the approaches outlined below would be particularly interesting for text corpora where the names of the victim and offender are known.

The semantic role of victim: Subject vs. Object. The first victim-blaming indicator captures the semantic structure of the sentences in which the victim is mentioned. This approach is grounded in psychological research, which suggests that grammatical structure can convey an author’s perception of a fact and influence how recipients perceive it. [2] asked participants in an experiment to read fictional reports of rape in which either the victim or the perpetrator was the grammatical subject in some 75% of the sentences. Participants who read reports in which the victim was primarily the grammatical subject were more likely to shift some responsibility for the assault to the victim. These findings are in line with a study by [1] in which study participants were asked to judge the intentionality of the grammatical object and subject in a set of sentences that were ambiguous in terms of intentionality. Study participants attributed significantly more intentionality to the grammatical subject than to the object.

The above findings also apply to sentences and texts written in passive voice. In an experiment by [65], study participants were asked to describe an uncommented video showing a rape scene and to complete a questionnaire measuring rape myth acceptance. The study reveals that describing the scene primarily in the passive voice (with the victim as the subject) is positively correlated with attributing responsibility to the victim. [79] confronted study participants with fabricated news reports on violence against women written in either active or passive voice. They found that males but not females rated the perpetrator’s responsibility higher after reading reports in the active voice.

These findings suggest developing a victim-blaming indicator that captures the semantic role of the victim in judicial opinions. I construct an indicator that measures the relative frequency of sentences in which the victim functions as grammatical subject, where the semantic roles of all words are identified using the part-of-speech tagger from the Python package `spacy` which is trained based on both active- and passive-voice sentences.

Sentiment orientation. The second indicator aims to capture the sentiment orientation of sentences mentioning the perpetrator. [3] develop such an indicator based on SentiWordNet, a database with information on word sentiments (see [S1 Appendix](#) for more information on the SentiWordNet 3.0). [3] propose to calculate sentiment scores for a word of interest as the average of the positivity or negativity scores of all context words (which they define as the 5 words surrounding the word of interest). As the meaning of the context words is not identifiable without deeper content analyses, they take the weighted average of the scores for all meanings of each context word, with the weights calculated based on the meaning ranks from the SentiWordNet dataset.

To assess the sentiment orientation of contexts in which the perpetrator is named, I determine the negativity score as suggested by [3]. The negativity score is zero for neutral and

positively connotated words and takes on positive values for negatively connotated words. The problem with calculating the average sentiment of context words is that negations cannot be adequately taken into account. Addressing this issue would necessitate employing more advanced methodologies. Nevertheless, as long as the frequency of negation usage does not change over time, the imprecision in the sentiment indicator used in this study will not introduce bias to the causal effect estimator.

The reason for focusing on the context words surrounding the mention of the perpetrator is that positively connotated context words of the victim may not point at the absence of victim-blaming language, as, e.g., comments on the victim's clothing or attitude towards the perpetrator are common examples of rape myth reinforcement.

Word embedding. In a third approach to identifying the effect of the #MeToo movement on the use of victim-blaming language, the development of the context in which the words “victim” and “appellant”/“petitioner” appear is assessed by means of Word2Vec, a neural network-based word embedding method (for more information on Word2Vec see [S1 Appendix](#) and [80]). The Word2Vec algorithm is trained for each calendar quarter and for sexual and personal crimes separately. In order to make the Word2Vec models comparable, they are aligned using Compass Aligned Distributional Embeddings (CADE). As suggested by [3], I calculate the cosine similarity between the vectors of the first quarter of 2015 (for the words “victim” and “appellant”) and those for each other quarter, both for the group of opinions on sexual offenses and the group of non-sexual offenses, in order to be able to assess the development of the word embedding vectors.

Since the word embedding approach only identifies one vector per time period and group, this approach only allows for estimating the effect of the #MeToo movement, but not the uncertainty of the estimate, i.e., the standard errors.

Text vectorization

For analyzing the linguistic development in the judicial opinions, the opinions are vectorized by means of the Bag of Words (BoW) and the term frequency—inverse document frequency (tf-idf) algorithms, both of which are frequently used for text classification.

Both algorithms convert text documents into fixed-length vectors, where the vector length is determined by the number of unique words in the entire corpus of text documents. The vector elements in the BoW representation of a document correspond to the relative frequency of each word. For a tf-idf representation, these vector elements are multiplied by the logarithmically scaled inverse proportion of documents in the corpus that contain the respective word. This way, words that are characteristic of one or a set of document(s) are given a higher weight, while vector elements are comparably small if the corresponding word occurs in (almost) every document.

The BoW approach can be refined by excluding common words like “court”, “trial”, or “judge” that appear in almost all opinions and offer little insight into the evolution of court language. Additionally to the BoW representation of opinions based on the entire vocabulary, I also construct a second set of BoW vectors including only words that occur in less than 95% of all documents. This alternate BoW approach will in the following be referred to as reduced-corpus BoW.

Concerning tf-idf, applying this approach to the corpus of judicial opinions without additional text preprocessing would pose a problem: names of persons, places and organizations that only occur in one or a few opinion(s) receive disproportionately high weights, despite their lack of relevance to the document's content. To circumvent this, I apply the Stanford PoS

Tagger to identify and exclude all proper names before applying the tf-idf algorithm (see [S1 Appendix](#) for more information on the Stanford PoS Tagger).

Assessing the development of text vectors over time. The high-dimensional vectors determined with the three approaches outlined above are then projected using principal component analysis (PCA) in order to reduce the dimensionality and thereby the computational complexity, while preserving as much of the variance of the original text vectors as possible. Since the dimensionality reduction is applied to all text vectors together, the resulting vectors are suitable for computing similarity between the opinions they represent. [3] take a similar approach to evaluate the evolution of word context vectors over time.

The data is reduced to as many principal components as needed to capture 90% of the variation in the original opinion vector dataset. The resulting reduced vectors have 64 (BoW), 65 (reduced corpus BoW) and 36 (tf-idf) dimensions respectively. Studies on text classification usually reduce the data to only 1–5 dimensions to avoid capturing noise. The goal of the present study, however, is not to classify texts, but rather to detect even minor changes in the language used in judicial opinions. These changes may indicate, e.g., shifts in the treatment of victims in court, while in the context of the thematic classification, they might be considered noise.

Following this data reduction step, I average the vectors of all opinions from the first half of 2015 for each court separately and calculate the distance between these 2015 average vectors and all other opinion vectors of the corresponding court. To do this, I choose the L_1 distance metric as proposed by [81]. They show that for data of 20 dimensions and more the L_1 metric is the best distance measure in terms of contrasting two points.

When employing the text vectorization approach detailed above, it is important to acknowledge that the method, by its nature, captures general language developments in court opinions. While the approach is capable of detecting shifts in language use, it does not distinguish between changes related to victim blaming and more general linguistic developments, such as changes in the formal textual requirements for judicial opinions. This approach is therefore primarily suitable for applications that aim to capture a broad spectrum of linguistic variations.

Identification

The effect of the #MeToo movement on the language at court is assessed by means of a DiD and an event study approach. To account for the fact that the process between the initial trial court hearing and the publication of a judicial opinion from an appellate court often takes several months, sometimes years, I particularly look at the effect of #MeToo on court opinions published at least one year after the movement began; that is, while the estimates for the year following the movement's onset are also reported, a stronger emphasis is put on subsequent years, particularly November 2018 through April 2020, when the Covid-19 pandemic hit the U.S. While appeals must be filed promptly after publication of the trial court's decision (usually within 30 days), the process in appellate courts often takes considerably longer: the U.S. Courts of Appeals report the median disposition time, i.e., the time between the filing of an appeal and the appellate court's decision, to be 8.6 months for 2015 (see <https://www.uscourts.gov/news/2016/12/20/just-facts-us-courts-appeals>). Thus, opinions published at least one year after the movement's onset are likely to contain closing arguments, trial court judgements, appellee and appellant briefs, as well as appellate court reasoning and decisions written after the #MeToo movement began. Opinions on trial court proceedings lasting more than one year, however, may reference pre-#MeToo hearings, which is essential to consider when

interpreting the estimated effect in order to avoid underestimating the movement's impact on court language.

Difference-in-Differences approach

I apply a DiD approach both with two and multiple time periods in order to identify the Average Treatment Effect on the Treated (ATET) (see, e.g., [82–84]). The sample consists of all judicial opinions on cases that can be classified as “crimes against persons”. Sexual offense opinions form the treatment group, i.e., the set of opinions that are affected by the #MeToo movement, opinions on other crimes against persons serve as control group. One advantage of choosing the sample of “crimes against persons” is that all crimes against persons involve a victim and a perpetrator, which is important for detecting victim-blaming language. Further, the comparison of sexual offenses and other crimes against persons is also common in the literature on victim blaming and rape-myth acceptance (see, e.g., [42, 85, 86]).

In order to identify the ATET of the #MeToo movement using a DiD approach, certain assumptions must hold. For simplicity, I will discuss these assumptions for the simple case of two time periods, but the discussion is easily transferable to the multiple time period approach.

Common support assumption. The common support assumption states that for each post-treatment observation from the treatment group, there must be comparable observations in the other three groups, i.e., the pre-treatment observations from the treatment and the control group, as well as the post-treatment observations from the control group. Translated to the present study, there must be opinions in these three groups that are comparable to post-treatment opinions on sexual offenses in terms of court and judge characteristics. This assumption is likely to hold because I only consider judicial opinions from courts that handle both sexual offenses and other crimes against persons. Furthermore, judges are usually assigned their cases randomly, which is why the characteristics of judges in all 4 groups should be similar.

No anticipation assumption. The no anticipation assumption states that the treatment may not have any effect on the outcome in pre-treatment periods, which would have been the case if judges and other actors in court had anticipated the movement and changed their behavior accordingly before the movement went viral. This assumption is also likely to hold as the #MeToo movement was launched immediately after the sexual assault allegations against Harvey Weinstein came to light, and to an extent that no one anticipated.

Common trend assumption. The third assumption is the common trend assumption. It states that in absence of the treatment, outcomes in the treatment and control group would follow a parallel trend, or in other words, that the gap between the outcome in the two groups would be constant over time.

There are various circumstances that may challenge the validity of this assumption. For one, legislative reforms are likely to affect language in court, as they, e.g., bring about changes in how certain crimes are sanctioned, what evidence is admitted in court, or what role parties and witnesses may take in the court process. To the best of my knowledge, and as noted by [42], there have been no major legislative changes on crimes against persons (neither on sexual nor on non-sexual offenses) in the U.S. during the years under study.

A second potential violation of the Common Trend Assumption could stem from external factors influencing the judges' attitudes on sexual offense cases, such as other scandals, movements, or media releases addressing sexual violence. All of these events could potentially draw public attention to the issue of sexual violence and lead to a change in attitudes toward sexual violence cases and their victims. These concerns can be mitigated by looking at search

history, traditional media and social media data from the U.S. (see Section). The data show that at no other time during the period studied there was as much attention drawn to sexual harassment and assault as in the weeks following the onset of the #MeToo movement.

Finally, the language in trials on sexual offenses might generally, even in absence of the #MeToo movement, evolve faster than that in trials on other crimes against persons. One can argue that crimes against persons, being relatively “old” crimes, are less likely to change in nature, thus not requiring frequent reinterpretation of corresponding laws (unlike, e.g., cyber-crimes). However, recent decades have witnessed a trend toward more gender-sensitive language in parts of society, which may—if also perceptible in courts—result in a violation of the parallel trend assumption. Furthermore, various feminist developments over recent decades might have encouraged judges to more frequently reinterpret the law on sexual offenses, even if it was not for the #MeToo movement. The results from the DiD approach should therefore be interpreted with some reservation, and placebo tests will be conducted to rule out serious violations of the common trend assumption.

Stable Unit Treatment Value Assumption (SUTVA). The SUTVA assumption, finally, rules out spill-over effects between observations. In the present study this implies that the #MeToo movement must solely influence the language used in opinions related to sexual offenses, without affecting language in opinions concerning other crimes against persons. Furthermore, the SUTVA rules out compositional changes in the treatment or control group over time.

Spill-over effects might arise insofar as the #MeToo movement might have not only altered how courtroom actors perceive and treat victims of sexual offenses but also potentially those of other violent crimes, leading to a potential underestimation of the movement’s impact on victim-blaming language. However, besides the fact that #MeToo discourse did not tackle victim blaming in crimes other than sexual offenses, studies have shown that victim blaming is more prevalent in the context of sexual offenses than in that of other crimes [86], indicating a greater potential for developing a language of compassion toward the victim in sexual offense cases.

Arguing for the validity of the second implication of SUTVA presents a greater challenge. The composition of sexual offense cases heard in court may have changed as the #MeToo movement led to an increase in the reporting of such crimes. [42] find a rise in reports and subsequent arrests related to sexual offenses following the #MeToo movement. Some of these additional arrests may have resulted in convictions, some of which in turn may have been appealed and thus become part of the sample. Similarly, there may be cases in my sample that went to trial only because of the prohibition on non-disclosure agreements that some states enacted in response to #MeToo.

Although the increase in cases in itself is not problematic for identifying the effect of #MeToo on language in court, it may have led to changes in the composition of sex offenses addressed within the judicial opinions in my sample. Compositional changes, in turn, would likely result in language shifts that cannot be attributed to a change in how a given case is treated in court. Concerns about #MeToo-induced compositional changes in the present sample can be debunked to at least some extent: For one, the findings by [42] suggest that the set of additional sexual offense reports includes a disproportionate number of comparatively lighter offenses and cases with less pressing evidence, where the likelihood of a criminal trial is low. This, in turn, reduces the chances of these cases reaching an appeals court and ending up in my sample.

Furthermore, S2 Table shows a consistent share of sexual offense opinions in the sample, ranging from 35% to 36% annually. The table also reports the share of sexual offense opinions that address different sexual offenses (Note: The opinions are categorized based on which sexual offense-related terms could be identified in the opinion’s introduction, i.e., one opinion

may be counted in more than one category). The opinions are grouped into six categories of sex offenses similar to those defined by the FBI National Incident-Based Reporting System, which [42] use in their study. Additionally, I include a subcategory for sexual assault of children and minors. The table reveals no substantial shifts in the composition of sexual offense cases over time. This holds true when focusing on the years 2017 to 2019, despite the increase in reports found by [42]. The p-values obtained when controlling for court fixed effects indicate that there are no statistically significant differences between the years under study. When not controlling for court fixed effects, a few of these differences are moderately statistically significant, both before and after the #MeToo movement, suggesting that the slight inter-temporal differences in the composition are attributable to general variations in the number of opinions per court and differences across courts in the composition of opinions.

To deal with this, the DiD analysis is complemented by an Inverse Probability Weighting (IPW) DiD approach [87], in which I weigh all control observations and the pre-movement sexual offense opinions to have the same distribution of courts as the group of post-treatment sexual offense opinions (IPW based on offense categories would be problematic in that the offense categories of the control and treatment groups do not overlap by design). Further, I conduct a second robustness check, building on the results of [42]. The authors find that the #MeToo movement did not affect the reporting of certain types of criminal sexual offenses, namely sodomy and rape. I therefore re-run my analyses, considering only opinions of these types.

Nevertheless, there may still be #MeToo-induced shifts in the composition of cases that cannot be controlled for or accounted for in robustness checks: the composition of offenses of a given type could still change in terms of the strength of the evidence and/or the severity of the offense. This would be the case if the #MeToo movement had led to an increase in trial court cases involving sexual offenses and, at the same time, a decrease in the proportion of convicted offenders who appeal their convictions. In this case, the proportion of appeal cases with inconclusive evidence may have increased, which could have accelerated the development of language in the sex offense sample and thus biased upward the effect estimates from the text vectorization approaches. In contrast, the estimates for the impact on the use of victim-blaming language would in this scenario constitute lower bounds of the actual decline, since the more reasons there are to doubt the credibility of the victim or the seriousness of the incident, the more likely it is that victim-blaming language will be used.

In the DiD approach, I control for court fixed effects, since the court- or state-specific laws and rules, as well as the terminology therein, are likely to differ. In addition, I control for the word count to account for the fact that there are some very short and formal opinions in the sample that have little or no flexibility in how they are written. However, I also report the estimates for when no controls are included. The reported standard errors are heterogeneity-robust and clustered at the court level [88, 89].

Event study approach

I complement the DiD analysis with an event study approach, in which I assess the development of the text quantifiers before and after the #MeToo movement in a panel setting. In doing so, I attempt to address the problem that the set of opinions on non-sexual offenses may represent an imperfect control group. In the event study approach, I examine the judge-specific developments in the text quantifiers and victim blaming language indicators before and after the onset of the #MeToo movement. The purpose of this event study application is not to identify the causal effect of the movement, but rather to observe whether there was a shift in

the overall development of language in sex offense cases and the use of victim-blaming language from before to after the movement.

I identify 1,382 judges, who, on average, publish roughly 11 sexual offense case opinions during the observation period, with the median number of opinions being 4, i.e., there are a few judges who authored a large amount of opinions (up to 184 opinions) while many others only published a handful during the study period (for more detail on the identification of the authoring judges, see [S1 Appendix](#)). For each judge in the sample, I calculate the six-month average of each text quantifier and victim blaming indicator to obtain a panel data structure with one or no observation per individual and time period. To avoid losing too many observations, I keep all judges who published at least one opinion before the onset of the #MeToo movement and at least one a year or more after the start of the movement. Then, I apply a Fixed Effects (FE) approach while weighting the observations by the number of opinions they were calculated from. Through weighting the observations, I account for the fact that the judges differ greatly in how many sex offense opinions they publish per six-month period, which makes some judges much more important for the development of language in court than others.

Although some information is lost by averaging the observations per judge and 6-month period, the panel approach might eliminate some of the noise typical of text analyses by increasing the amount of text per observation. On the other hand, however, it requires me to exclude several observations from the sample (3,843 opinions authored by 785 judges) because I do not have observations from either before or after #MeToo for these judges. By excluding judges who, for whatever reason, do not frequently publish precedential opinions on sex offenses, important information may be lost.

While the DiD approach captures language changes resulting from shifts in linguistic conduct of court actors as well as compositional shifts in the judiciary, this panel approach does not capture the latter. To some extent, comparing the estimates from these two approaches can help discern whether the observed evolution in language is primarily attributable to changes in the linguistic conduct of individual judges or rather to compositional shifts in the judiciary. There is a growing literature on the drivers of attitudinal, cultural and literary change in society. Most studies conclude that such change is driven to a greater extent by generational turnover than by shifts in the behavior of individuals (see, e.g., [90–92]), and that intra-personal changes are more pronounced among younger people (see, e.g., [91]). The strong influence of cohort change on cultural shifts would suggest higher DiD than panel event study effect estimates.

Assessing effect heterogeneity

Finally, I also assess whether there is evidence of effect heterogeneity with respect to a judge's gender or political affiliation, as well as with regard to the political orientation of the state in which a court is located. This is because different studies on victim blaming and rape myths acceptance show that females are less likely to accept rape myths and shift the blame for an assault upon the victim (e.g., [93–96]). [97] finds that this is particularly true for female Democrats.

Most studies in which participants were confronted with a sexual assault scenario find that men were more likely than women to blame the victim and show signs of rape myth acceptance, while other studies find no significant effect of gender on the likelihood of victim blaming (see [98–100] for reviews). Other research shows that study participants with politically conservative views are more likely to (partially) blame the victim for a sexual assault [101, 102]. For the judicial context, [97] finds that female Democratic judges are less likely to use

rape myths than male judges (regardless of political affiliation), while her results show no significant difference between female Republican judges and male judges.

In addition to the differences in victim blaming and rape myth acceptance noted above, there is also evidence that the perception of and reaction to the #MeToo movement differ across genders and political camps. [59] find in a poll that Democrats are more likely than Republicans to say they were aware of the movement and mobilized by it. An analysis of the members of Congress' communications on their public Facebook pages reveals that in the wake of the #MeToo movement, far more female than male members addressed the issue of sexual violence in their posts, with this pattern evident in both political parties [37]. In light of the findings by [91] that intra-personal changes are more pronounced among younger people, it might also be interesting to assess effect heterogeneity with regard to the age of judges.

For the scope of this analysis, I will however concentrate on whether the judges' language in court is affected differently by the #MeToo movement depending on their gender and political affiliation. The research cited above suggests that female and/or politically liberal judges were more receptive to the #MeToo movement and more willing to change their behavior. Then again, these judges may have been more cautious in their choice of words prior to the movement and may have already avoided language that implied victim blaming, leaving them little room for change toward language that attributed less blame to the victim.

The analysis of treatment effect heterogeneity by judge characteristics is restricted to courts for which the names of the authoring judge is available on CourtListener. Information on the judges' gender and political affiliation is obtained from ballotpedia, an online encyclopedia on American politics and elections. Ballotpedia only provides the political affiliation of some judges. To determine the political affiliation of the other judges, I draw on the party for which the judge ran in the judicial election or the political affiliation of the politician who appointed the judge, depending on the process used to select judges. For about 21% of the judges no political affiliation can be determined. To assess effect heterogeneity with respect to a state's political orientation, I categorize those states as predominantly Democratic (Republican) that were won by the Democratic (Republican) party in at least three of the four 2008–20 presidential elections. Opinions from swing states and courts at the supra-state level are excluded for this analysis.

Results

Victim blaming indicators

Fig 2 shows the development of the victim blaming semantics and sentiment indicators in opinions on sexual offenses and on other crimes against persons. Neither plot suggests any substantial differences between the development of these two indicators in the treatment and the control group. Further, the DiD estimates provided in Table 1 do not indicate a significant impact of the #MeToo movement on either of the two indicators. The DiD estimates for the semantics indicator suggest a slight, but not statistically significant, #MeToo-induced decline in the number of victim mentions as a grammatical subject and hence a decline in victim blaming. The estimates for the sentiment indicator are also not statistically significant and even indicate, contrary to the research hypothesis, a decrease in the use of words with negative connotations in the context of mentions of the perpetrator. Likewise, the event study approach (see Fig 3) shows no substantial changes in either of the two victim blaming indicators.

While assessing the impact of the #MeToo movement on the victim blaming indicators does not yield significant results, the plot on the development of the semantics estimator is nevertheless interesting. It shows that the use of the victim as grammatical subject is generally substantially higher in opinions on sexual assault cases than in cases on other crimes against

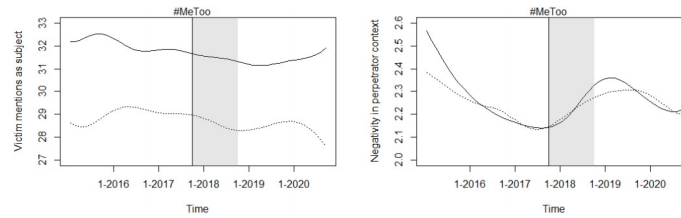


Fig 2. Development of victim blaming indicators in treatment & control group. Kernel-smoothed plot of victim blaming indicators, with the solid line representing sexual offense opinions and the dotted line representing the control group. Curves were smoothed using the default settings of the `sm.regression` function in R. The semantics indicator measures mentions of the victim as a subject as a percentage of total mentions, and the sentiment indicator measures the negativity of the context words of offender mentions, ranging from 0 to 100, with the words being more negatively connotated the higher the score.

<https://doi.org/10.1371/journal.pone.0302827.g002>

persons. Given that several studies indicate higher prevalence of victim blaming in sexual assault cases than in cases on other crimes against persons (see Section) and that the indicator is constructed based on scientific findings (see Section), it seems reasonable to explore whether this indicator may be useful for measuring the extent of victim blaming in judicial opinions—as long as the purpose is to classify opinions or quantify the status quo rather than measure changes over time. The sentiment indicator, on the other hand, seems to vary little both over time and crime types. It generally takes very low values, which may be because judges deliberately avoid sentiment-charged language. It therefore does not seem appropriate for use in the context of court opinions and may be better suited for contexts with more colloquial language such as social media.

Table 1. DiD estimates of effect heterogeneity for victim blaming indicators.

	Share of Victim as Subject			Negativity of Offender Context		
sex cr. x post	-0.687			-0.079		
	(1.087)			(0.056)		
sex cr. x '19		-0.446			0.002	
		(1.088)			(0.064)	
sex cr. x '20		-0.872			-0.087	
		(1.340)			(0.057)	
sex cr. x H1-'19			-1.214			0.000
			(1.528)			(0.062)
sex cr. x H2-'19			0.273			0.003
			(1.144)			(0.117)
sex cr. x H1-'20			-1.155			-0.107
			(1.407)			(0.098)
sex cr. x H2-'20			-0.621			-0.069
			(1.587)			(0.098)
post	X			X		
year FE		X			X	
half-year FE			X			X
court FE	X	X	X	X	X	X
# words	X	X	X	X	X	X

DiD estimates of effect heterogeneity for victim blaming indicators. Significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

<https://doi.org/10.1371/journal.pone.0302827.t001>

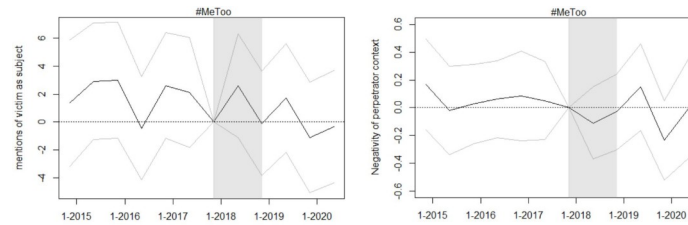


Fig 3. Event study estimates for evolution of victim blaming indicators. Event study estimates for the evolution of the victim blaming indicators in sexual offense opinions. The gray lines indicate the 90 percent confidence interval around the estimates.

<https://doi.org/10.1371/journal.pone.0302827.g003>

The results of the robustness checks for the DiD with victim blaming indicators are provided in [S3 Table](#). Given that no significant effects on either indicator can be identified using the DiD approach, the fact that the placebo test does not indicate a violation of the parallel trend assumption is not particularly informative. For the sentiment indicator, both the IPW estimates as well as the DiD estimates based on sodomy and rape cases only indicate a significant decrease in negatively connoted words in the context of perpetrator mentions. In the case of the semantics indicator, both estimates have different signs and are not statistically significant.

The effect heterogeneity estimates obtained using the DiD and event study approach, respectively, can be found in [S5 Table](#), [S1](#) and [S2 Figs](#). For the semantics indicator, the estimates suggest a larger decline in mentions of the victim as subject among female and Democratic judges as well as in predominately Democratic states, although again the differences are not statistically significant. The DiD estimates for the sentiment indicator imply a greater decline in the use of words with negative connotations among females and no difference in the effect of #MeToo on this indicator between Democrats and Republicans. The event study approach, on the other hand, suggests that Democrats have increased their use of negatively connoted words when compared to the development of this indicator among Republicans.

Finally, [Fig 4](#) displays the development of the Word2Vec representation of the words “victim” and “appellant” in both the control and treatment group. Contrary to the research hypothesis, the context in which victim and offender are mentioned does not evolve faster for opinions on sexual offenses than for those on other crimes against persons. Again, this approach may be better suited to contexts with more flexible and rapidly evolving language.

Text vectorization

[Fig 5](#) illustrates the evolution of the BoW (the reduced sample BoW quantifier evolves similarly to the BoW quantifier in both groups, which is why it is not shown here). and the tf-idf text

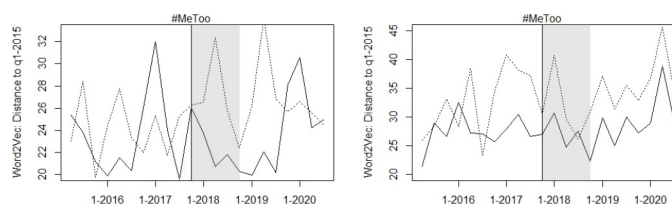


Fig 4. Development of the word2vec representation of the word “victim” and “appellant”. Development of the word2vec representation of the word “victim” (left) and “appellant” (right), with the solid line representing sexual offense opinions and the dotted line representing the control group.

<https://doi.org/10.1371/journal.pone.0302827.g004>

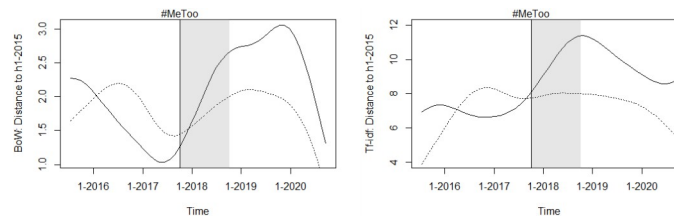


Fig 5. Distance of text vectors from their H1 2015 average over time. Kernel-smoothed plot of the distance of text vectors from their H1 2015 average, with the solid line representing sexual offense opinions and the dotted line representing the control group. The distance of opinions to the H1 2015 average is expressed relative to the median distance of all opinions to the H1 2015 average, in percent.

<https://doi.org/10.1371/journal.pone.0302827.g005>

quantifiers in opinions on sexual offenses and on other crimes against persons. Both charts suggest that language in opinions on sexual offenses evolves more rapidly than that in opinions on other crimes against persons between the onset of the movement and 2019. However, the graphs also indicate that there may be problems with the parallel trend assumption. Moreover, the BoW graph shows a narrowing of the distance between the opinion vectors and their H1–2015 average in 2016 and 2017, as well as in 2020, which could be due to changes in the composition of courts, but also to other unobservable factors, which in turn would be critical for identifying the causal effect (other smoothing methods and smaller bandwidths yield similar curves. Thus, it does not seem to be an (over-)smoothing issue).

The DiD estimates for the text vectorization-based opinion quantifiers can be found in Table 2. The results point at a slight #MeToo-induced change in courtroom language, which however materializes with a substantial time lag. The estimates suggest that the language in sexual offense opinions deviates more quickly from the 2015 average than in opinions on other cases of crimes against persons. The DiD estimate for the second half of 2020 is statistically insignificant for all three text quantifiers and numerically small for the two BoW quantifiers, indicating that the language change in sexual offense opinions is not likely due to the COVID-19 crisis.

The event study estimates in Fig 6 show a decrease in the distance to the H1–2015 average in 2016 that is similar to, though less pronounced than, that observed in Fig 5. Further, the event study estimates do not indicate a stronger deviation of language from the H1–2015 average in the years after #MeToo than in the years before #MeToo, suggesting that the effect estimated with the DiD approach may be attributable to changes in case composition or personnel rather than changes in judges' attitudes.

The results of the robustness checks are presented in S4 Table. While the placebo tests do not reveal a significant violation of the common trend assumption, they do not allow me to rule out such a violation, especially because the estimated effects of the placebo treatment on the (reduced sample) BoW quantifiers are positive, just like the observed effect in the main DiD analysis. I therefore also estimate the effect of placebo treatments at other points in time during the pre-#MeToo period, all of which turned out to be statistically insignificant, with some of them having a positive and others a negative sign. The fact that the IPW-based DiD and the reduced sample DiD estimate positive effects for all quantifiers, some of which are statistically significant, support the finding that the #MeToo movement has caused a slight increase in language development in sexual assault opinions, with the IPW results ruling out that the observed effect in the main analysis is due to changes in the composition of courts, while the latter rules out that it is due to an increase of reports of sexual offenses. When considering these results in conjunction with the event study results, one possible explanation for the

Table 2. DiD estimates for text vectorization-based opinion quantifier.

	BoW			Reduced Sample BoW			tf-idf		
sex cr. x post	0.815*			0.700			1.665		
	(0.469)			(0.511)			(1.562)		
sex cr. x '19		0.629			0.269			1.338	
		(0.535)			(0.507)			(1.732)	
sex cr. x '20		0.919**			1.165**			1.106*	
		(0.454)			(0.496)			(0.646)	
sex cr. x H1-'19			0.658			0.474			0.940
			(0.715)			(0.718)			(1.835)
sex cr. x H2-'19			0.603			0.079			1.716
			(0.563)			(0.564)			(1.813)
sex cr. x H1-'20			1.461			1.806*			0.281
			(0.905)			(0.955)			(0.966)
sex cr. x H2-'20			0.317			0.459			1.958
			(0.692)			(0.670)			(1.312)
post	X			X			X		
year FE		X			X			X	
half-year FE			X			X			X
court FE	X	X	X	X	X	X	X	X	X
# words	X	X	X	X	X	X	X	X	X

DiD estimates for text vectorization-based opinion quantifiers. The distance of opinions to the H1 2015 average is expressed relative to the median distance of all opinions to the H1 2015 average, in percent. Significance levels:

- * p<0.1,
- ** p<0.05,
- *** p<0.01.

<https://doi.org/10.1371/journal.pone.0302827.t002>

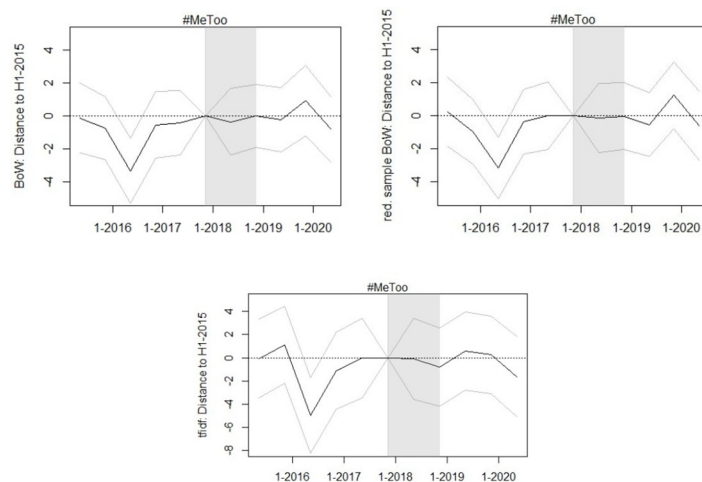


Fig 6. Event study estimates for evolution of text vectorization-based opinion quantifiers. Event study estimates for the evolution of the text vectorization-based opinion quantifiers. The gray lines indicate the 90 percent confidence interval around the estimates. The distance of opinions to the H1 2015 average is expressed relative to the median distance of all opinions to the H1 2015 average, in percent.

<https://doi.org/10.1371/journal.pone.0302827.g006>

observed slight linguistic change in sexual offense opinions is the change in judicial appointments toward judges with more progressive views on sex offenses. However, given the numerically small effect estimates, the noise in these quantifiers, and the fact that the effect does not appear until two years after the movement began, it is difficult to attribute the observed effect to the #MeToo movement.

A look at the effect heterogeneity estimates in [S6 Table](#), [S1](#) and [S2 Figs](#) does not give a clear picture. While the DiD estimates for the (reduced sample) BoW suggest a faster evolution of language in opinions written by females and Democrats, the event study plots point to a similar evolution of language in opinions written by female and male or Democratic and Republican judges, respectively. For the tf-idf approach, the DiD estimates suggest that the language of female and Democratic judges evolves less rapidly than that of their counterparts.

Conclusion

In this study, I quantified judicial opinions by means of different indicators and text vectorization methods to assess how the #MeToo movement has affected the evolution of language in sexual assault opinions and whether it has led to a decrease in victim blaming. Although I did not obtain statistically significant estimates for the impact of the movement on most quantifiers, the point estimates suggest a faster evolution of language in sexual assault opinions as well as a decline in victim blaming. The reasons for not obtaining statistically significant results may be manifold. For one, the language in judicial opinions is generally not very flexible, i.e., there is regulation on what different parties are allowed to say in court and the structure of court opinions, particularly in certain paragraphs, is highly formalized. Therefore, any treatment should be expected to have a smaller effect on language in court opinions than on language in any other more flexible text corpus. Moreover, the text vectorization methods in particular, but also the victim blaming indicators, capture a lot of noise, leading to large standard errors in the estimators.

The study's real merit lies in its potential to make a valuable contribution to the expanding body of literature on text-based causal inference: for one, I have developed indicators that can be useful proxies for victim blaming and may be used as treatment, outcome or control, if they are applied to a text corpus with more flexible language. In text corpora in which the identity of victim and perpetrator is known, the proposed indicators could also be applied to their names and personal pronouns. Furthermore, the victim-blaming indicators might also be used in descriptive contexts to gauge the level of victim blaming in a text body or to compare the extent of victim blaming across different text corpora.

I presented a text-vectorization-based approach for evaluating language development, enabling the integration of text vectorization methods within DiD analyses or panel data methods. This approach may prove beneficial when evaluating bodies of text with more adaptive language, analyzing larger text bodies, or evaluating longer-term effects of a treatment in a panel setting. When examining the impact of the #MeToo movement on more flexible and faster-responsive text bodies, one could also take advantage of temporal and regional variations in hashtag usage following the onset of the movement.

Supporting information

S1 Appendix. Additional information on text processing.
(PDF)

S1 Table. Descriptives. Number of opinions per court.
(PDF)

S2 Table. Descriptives. Composition of sexual-violence related opinions, by crime type & year.

(PDF)

S3 Table. DiD: Robustness checks. Victim blaming indicators.

(PDF)

S4 Table. DiD: Robustness checks. Text vectorization.

(PDF)

S5 Table. DiD: Effect heterogeneity. Victim blaming indicators.

(PDF)

S6 Table. DiD: Effect heterogeneity. Text vectorization.

(PDF)

S1 Fig. Event study approach: Effect heterogeneity. Event study estimates of difference in development of female vs. male judges.

(TIF)

S2 Fig. Event study approach: Effect heterogeneity. Event study estimates of difference in development of Democrats vs. Republicans.

(TIF)

Author Contributions

Conceptualization: Henrika Langen.

Data curation: Henrika Langen.

Formal analysis: Henrika Langen.

Methodology: Henrika Langen.

Writing – original draft: Henrika Langen.

Writing – review & editing: Henrika Langen.

References

1. Strickland B, Fisher M, Keil F, Knobe J. Syntax and intentionality: An automatic link between language and theory-of-mind. *Cognition*. 2014; 133(1):249–261. <https://doi.org/10.1016/j.cognition.2014.05.021> PMID: 25058414
2. Niemi L, Young L. When and why we see victims as responsible: The impact of ideology on attitudes toward victims. *Personality and social psychology bulletin*. 2016; 42(9):1227–1242. <https://doi.org/10.1177/0146167216653933> PMID: 27340155
3. Jatowt A, Duh K. A framework for analyzing semantic change of words across time. In: *IEEE/ACM Joint Conference on Digital Libraries*. IEEE; 2014. p. 229–238.
4. Khalid O, Srinivasan P. Style matters! Investigating linguistic style in online communities. In: *Proceedings of the International AAAI Conference on Web and Social Media*. vol. 14; 2020. p. 360–369.
5. Rho EHR, Mark G, Mazmanian M. Fostering civil discourse online: Linguistic behavior in comments of# metoo articles across political perspectives. *Proceedings of the ACM on Human-Computer Interaction*. 2018; 2(CSCW):1–28. <https://doi.org/10.1145/3274416>
6. Yu B, Kaufmann S, Diermeier D. Classifying party affiliation from political speech. *Journal of Information Technology & Politics*. 2008; 5(1):33–48. <https://doi.org/10.1080/19331680802149608>
7. Abercrombie G, Batista-Navarro RT. 'Aye' or 'no'? Speech-level sentiment analysis of Hansard UK parliamentary debate transcripts. In: *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*; 2018.

8. Hausladen CI, Schubert MH, Ash E. Text classification of ideological direction in judicial opinions. *International Review of Law and Economics*. 2020; 62:105903. <https://doi.org/10.1016/j.irl.2020.105903>
9. Undavia S, Meyers A, Ortega JE. A comparative study of classifying legal documents with neural networks. In: 2018 Federated Conference on Computer Science and Information Systems (FedCSIS). IEEE; 2018. p. 515–522.
10. Filtz E, Kirrane S, Polleres A, Wohlgenannt G. Exploiting eurovoc’s hierarchical structure for classifying legal documents. In: OTM Confederated International Conferences “On the Move to Meaningful Internet Systems”. Springer; 2019. p. 164–181.
11. Alekseev A, Katasev A, Kirillov A, Khassianov A, Zuev D. Prototype of Classifier for the Decision Support System of Legal Documents. In: SSI; 2019. p. 328–335.
12. Kulkarni V, Al-Rfou R, Perozzi B, Skiena S. Statistically significant detection of linguistic change. In: Proceedings of the 24th International Conference on World Wide Web; 2015. p. 625–635.
13. Hamilton WL, Leskovec J, Jurafsky D. Cultural shift or linguistic drift? comparing two computational measures of semantic change. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing. vol. 2016. NIH Public Access; 2016. p. 2116.
14. Frermann L, Lapata M. A bayesian model of diachronic meaning change. *Transactions of the Association for Computational Linguistics*. 2016; 4:31–45. https://doi.org/10.1162/tacl_a_00081
15. Hellrich J, Buechel S, Hahn U. Modeling word emotion in historical language: Quantity beats supposed stability in seed word selection. *arXiv preprint arXiv:180608115*. 2018;.
16. Nguyen D, Rose C. Language use as a reflection of socialization in online communities. In: Proceedings of the Workshop on Language in Social Media (LSM 2011); 2011. p. 76–85.
17. Keith KA, Jensen D, O’Connor B. Text and causal inference: A review of using text to remove confounding from causal estimates. *arXiv preprint arXiv:200500649*. 2020;.
18. Roberts ME, Stewart BM, Nielsen RA. Adjusting for confounding with text matching. *American Journal of Political Science*. 2020; 64(4):887–903. <https://doi.org/10.1111/ajps.12526>
19. Sallin A. Estimating returns to special education: combining machine learning and text analysis to address confounding. *arXiv preprint arXiv:211008807*. 2021;.
20. Veitch V, Sridhar D, Blei D. Adapting text embeddings for causal inference. In: Conference on Uncertainty in Artificial Intelligence. PMLR; 2020. p. 919–928.
21. Mozer R, Miratrix L, Kaufman AR, Anastasopoulos LJ. Matching with text data: An experimental evaluation of methods for matching documents and of measuring match quality. *Political Analysis*. 2020; 28(4):445–468. <https://doi.org/10.1017/pan.2020.1>
22. Field A, Park CY, Tsvetkov Y. Controlled Analyses of Social Biases in Wikipedia Bios. *arXiv preprint arXiv:210100078*. 2020;.
23. Wood-Doughty Z, Shpitser I, Dredze M. Challenges of using text classifiers for causal inference. In: Proceedings of the Conference on Empirical Methods in Natural Language Processing. Conference on Empirical Methods in Natural Language Processing. vol. 2018. NIH Public Access; 2018. p. 4586.
24. Ornaghi A, Ash E, Chen DL. Stereotypes in High-Stakes Decisions: Evidence from US Circuit Courts. *Center for Law & Economics Working Paper Series*. 2019;2.
25. Tan C, Lee L, Pang B. The effect of wording on message propagation: Topic-and author-controlled natural experiments on Twitter. *arXiv preprint arXiv:14051438*. 2014;.
26. Deshpande S, Li Z, Kuleshov V. Multi-Modal Causal Inference with Deep Structural Equation Models. *arXiv preprint arXiv:220309672*. 2022;.
27. Pryzant R, Card D, Jurafsky D, Veitch V, Sridhar D. Causal effects of linguistic properties. *arXiv preprint arXiv:201012919*. 2020;.
28. Wang Z, Culotta A. When Do Words Matter? Understanding the Impact of Lexical Choice on Audience Perception Using Individual Treatment Effect Estimation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 33; 2019. p. 7233–7240.
29. Fong C, Grimmer J. Discovery of treatments from text corpora. In: Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers); 2016. p. 1600–1609.
30. Feuerriegel S, Heitzmann SF, Neumann D. Do investors read too much into news? How news sentiment causes price formation. In: 48th Hawaii International Conference on System Sciences. IEEE; 2015. p. 4803–4812.
31. Egami N, Fong CJ, Grimmer J, Roberts ME, Stewart BM. How to make causal inferences using texts. *arXiv preprint arXiv:180202163*. 2018;.

32. Sobolev A. How pro-government “trolls” influence online conversations in Russia. University of California Los Angeles[7]. 2018;.
33. Chandrasekharan E, Pavalanathan U, Srinivasan A, Glynn A, Eisenstein J, Gilbert E. You can't stay here: The efficacy of reddit's 2015 ban examined through hate speech. *Proceedings of the ACM on Human-Computer Interaction*. 2017; 1(CSCW):1–22. <https://doi.org/10.1145/3134666>
34. Álvarez-Benjumea A, Winter F. The breakdown of antiracist norms: A natural experiment on hate speech after terrorist attacks. *Proceedings of the National Academy of Sciences*. 2020; 117(37):22800–22804. <https://doi.org/10.1073/pnas.2007977117> PMID: 32873640
35. Pavalanathan U, Han X, Eisenstein J. Mind Your POV: Convergence of Articles and Editors Towards Wikipedia's Neutrality Norm. *Proceedings of the ACM on Human-Computer Interaction*. 2018; 2(CSCW):1–23. <https://doi.org/10.1145/3274406>
36. Kantor J, Twohey M. Harvey Weinstein Paid Off Sexual Harassment Accusers for Decades. *The New York Times*, October 5, 2017. 2017;.
37. Anderson M, Toor S. How social media users have discussed sexual harassment since #MeToo went viral. Pew Research Center url: <https://www.pewresearch.org/fact-tank/2018/10/11/how-social-media-users-have-discussed-sexual-harassment-since-metoo-went-viral/>. 2018;.
38. Santiago C, Criss D. An activist, a little girl and the heartbreaking origin of 'Me too'. *CNN*, October 17, 2017. 2017;.
39. Caputi TL, Nobles AL, Ayers JW. Internet Searches for Sexual Harassment and Assault, Reporting, and Training Since the #MeToo Movement. *JAMA Internal Medicine*. 2019; 179(2):258–259. <https://doi.org/10.1001/jamainternmed.2018.5094> PMID: 30575847
40. Ennis E, Wolfe L. Media and #MeToo: How a movement affected press coverage of sexual assault. *Women's Media Center Report*. 2018;.
41. Time Magazine. Time Person of the Year 2017—The Silence Breakers. <https://time.com/time-person-of-the-year-2017-silence-breakers/>. 2017;.
42. Levy R, Mattsson M. The effects of social movements: Evidence from# MeToo. Available at SSRN 3496903. 2021;.
43. Keplinger Ksenia K JF Johnson Stefanie K, Barnes LY. Women at work: Changes in sexual harassment between September 2016 and September 2018. *PloS one*. 2019; 14(7). <https://doi.org/10.1371/journal.pone.0218313> PMID: 31314792
44. Jackson C, Newall M. The #MeToo Movement: One Year Later. Ipsos. 2018;.
45. Careerarc. Survey: 76% of Employed Americans Say #MeToo Positively Impacted How Sexual Harassment Is Addressed In The Workplace, While More Than 2 In 5 Say It Has Damaged Trust Between HR and Employees. Press Release, October 3, 2020. 2020;.
46. Greenfield R. Powerful Men Have Changed Their Behavior at Work Since #MeToo. *Blomberg*, April 10, 2018. 2018;.
47. Atwater LE, Tringale AM, Sturm RE, Taylor SN, Braddy PW. Looking Ahead: How What We Know About Sexual Harassment Now Informs Us of the Future. *Organizational Dynamics*. 2019; 48(4):100677. <https://doi.org/10.1016/j.orgdyn.2018.08.008>
48. French MT, Mortensen K, Timming AR. A multivariate analysis of workplace mentoring and socializing in the wake of #MeToo. *Applied Economics*. 2021; p. 1–19.
49. McGregor J. #MeToo backlash: More male managers avoid mentoring women or meeting alone with them. *Washington Post*, May 17, 2019. 2019;.
50. Bertotti C, Maxfield D. Most People Are Supportive of #MeToo. But Will Workplaces Actually Change? *Harvard Business Review*, July 10, 2018. 2018;.
51. NBC News and Wall Street Journal. Study #17409, Question 22d. 2017;.
52. Fileborn B, Loney-Howes R. Introduction: Mapping the emergence of #MeToo. In: *#MeToo and the Politics of Social Change*. Springer; 2019. p. 1–18.
53. Kagal N, Cowan L, Jawad H. Beyond the Bright Lights: Are Minoritized Women Outside the Spotlight Able to Say #MeToo? In: *# MeToo and the Politics of Social Change*. Springer; 2019. p. 133–149.
54. Taub A. #MeToo Paradox: Movement Topples the Powerful, Not the Ordinary. *The New York Times*, February 11, 2019. 2019;.
55. Mueller A, Wood-Doughty Z, Amir S, Dredze M, Nobles AL. Demographic representation and collective storytelling in the me too Twitter hashtag activism movement. *Proceedings of the ACM on Human-Computer Interaction*. 2021; 5(CSCW1):1–28. <https://doi.org/10.1145/3449181> PMID: 35295189
56. Evans A. # MeToo: A study on sexual assault as reported in the New York Times. *Occam's Razor*. 2018; 8(1):3.

57. Palmer JE, Fissel ER, Hoxmeier J, Williams E. # MeToo for whom? Sexual assault disclosures before and after# MeToo. *American journal of criminal justice*. 2021; 46(1):68–106. <https://doi.org/10.1007/s12103-020-09588-4>
58. Pew Research Center. American Trends Panel—Wave 35, May 29—July 11, 2018. <https://www.pewresearch.org/internet/dataset/american-trends-panel-wave-35/>. 2018;.
59. Castle JJ, Jenkins S, Orbals CD, Poloni-Staudinger L, Strachan JC. The Effect of the #MeToo Movement on Political Engagement and Ambition in 2018. *Political Research Quarterly*. 2020; 73(4). <https://doi.org/10.1177/1065912920924824>
60. Jernberg F, Lindbäck A, Roos A. A new male entrepreneur? Media representation of male entrepreneurs before and after #metoo. *Gender in Management: An International Journal*. 2020;. <https://doi.org/10.1108/GM-01-2019-0004>
61. Yoon S, Choe JS, Han YM, Kim SH. #MeToo Hits Online Dating, too: An Empirical Analysis of the Effect of the Me-Too Movement on Online Dating Users. 2020;.
62. Moricz S. Using Artificial Intelligence to Recapture Norms: Did# metoo Change Gender Norms in Sweden? arXiv preprint arXiv:190300690. 2019;.
63. Klar S, McCoy A. The #MeToo movement and attitudes toward President Trump in the wake of a sexual misconduct allegation. *Politics, Groups, and Identities*. 2021; 0(0):1–10.
64. Ait Bihi Ouali L, Graham DJ. The impact of the MeToo scandal on women's perceptions of security. *Transportation research part A: policy and practice*. 2021; 147:269–283.
65. Bohner G. Writing about rape: Use of the passive voice and other distancing text features as an expression of perceived responsibility of the victim. *British Journal of Social Psychology*. 2001; 40(4):515–529. <https://doi.org/10.1348/014466601164957> PMID: 11795065
66. Franiuk R, Seefeldt JL, Cepress SL, Vandello JA. Prevalence and effects of rape myths in print journalism: The Kobe Bryant case. *Violence against women*. 2008; 14(3):287–309. <https://doi.org/10.1177/1077801207313971> PMID: 18292371
67. McCoy VL. The effect of language used in newspaper report of rape: Measuring readers' judgments of victim blame. The University of Tulsa; 2004.
68. Suvarna A, Bhalla G, Kumar S, Bhardwaj A. Identifying Victim Blaming Language in Discussions about Sexual Assaults on Twitter. In: *International Conference on Social Media and Society*; 2020. p. 156–163.
69. Suvarna A, Bhalla G. # NotAWhore! A Computational Linguistic Perspective of Rape Culture and Victimization on Social Media. In: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*; 2020. p. 328–335.
70. Sacks M, Ackerman AR, Shlosberg A. Rape myths in the media: A content analysis of local newspaper reporting in the United States. *Deviant Behavior*. 2018; 39(9):1237–1246. <https://doi.org/10.1080/01639625.2017.1410608>
71. Northcutt Bohmert M, Allison K, Ducate C. "A rape was reported": construction of crime in a university newspaper. *Feminist Media Studies*. 2019; 19(6):873–889. <https://doi.org/10.1080/14680777.2018.1509104>
72. Temkin J, Gray JM, Barrett J. Different functions of rape myth use in court: Findings from a trial observation study. *Feminist Criminology*. 2018; 13(2):205–226. <https://doi.org/10.1177/1557085116661627>
73. Spencer D, Dodge A, Ricciardelli R, Ballucci D. "I think it's re-victimizing victims almost every time": police perceptions of criminal justice responses to sexual violence. *Critical criminology*. 2018; 26(2):189–209. <https://doi.org/10.1007/s10612-018-9390-2>
74. Smith O, Skinner T. Observing court responses to victims of rape and sexual assault. *Feminist Criminology*. 2012; 7(4):298–326. <https://doi.org/10.1177/1557085112437875>
75. Temkin J. Prosecuting and defending rape: Perspectives from the bar. *Journal of Law and Society*. 2000; 27(2):219–248. <https://doi.org/10.1111/1467-6478.00152>
76. Ehrlich S. Perpetuating—and resisting—rape myths in trial discourse. *Sexual Assault in Canada*. 2012; p. 389–408.
77. American Bar Association. How Courts Works. https://www.americanbar.org/groups/public_education/resources/law_related_education_network/how_courts_work/appeals/ Retrieved August 15, 2021. 2019;.
78. The Free Law Project. RECAP Archive Accessed November, 2020, <https://www.courtlistener.com/recap/>. 2020;.
79. Henley NM, Miller M, Beazley JA. Syntax, semantics, and sexual violence: Agency and the passive voice. *Journal of Language and Social Psychology*. 1995; 14(1-2):60–84. <https://doi.org/10.1177/0261927X95141004>

80. Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. Distributed representations of words and phrases and their compositionality. In: *Advances in neural information processing systems*; 2013. p. 3111–3119.
81. Aggarwal CC, Hinneburg A, Keim DA. On the surprising behavior of distance metrics in high dimensional space. In: *International conference on database theory*. Springer; 2001. p. 420–434.
82. Snow J. On the mode of communication of cholera. *Edinburgh medical journal*. 1856; 1(7):668. PMID: [29647347](https://pubmed.ncbi.nlm.nih.gov/29647347/)
83. Card D, Krueger AB. Minimum wages and employment: A case study of the fast-food industry in New Jersey and Pennsylvania. *The American Economic Review*. 1994; 84(4):772–793.
84. Acemoglu D, Angrist JD. Consequences of employment protection? The case of the Americans with Disabilities Act. *Journal of Political Economy*. 2001; 109(5):915–957. <https://doi.org/10.1086/322836>
85. Reich CM, Pegel GA, Johnson AB. Are survivors of sexual assault blamed more than victims of other crimes? *Journal of interpersonal violence*. 2021; p. 08862605211037423. PMID: [34376082](https://pubmed.ncbi.nlm.nih.gov/34376082/)
86. Bieneck S, Krahe B. Blaming the victim and exonerating the perpetrator in cases of rape and robbery: Is there a double standard? *Journal of interpersonal violence*. 2011; 26(9):1785–1797. <https://doi.org/10.1177/0886260510372945> PMID: [20587449](https://pubmed.ncbi.nlm.nih.gov/20587449/)
87. Abadie A. Semiparametric Difference-in-Differences Estimators. *Review of Economic Studies*. 2005; 72:1–19. <https://doi.org/10.1111/0034-6527.00321>
88. Zeileis A. Econometric computing with HC and HAC covariance matrix estimators. *Institut für Statistik und Mathematik, WU Vienna University of Economics and ...* 2004;.
89. Bertrand M, Duflo E, Mullainathan S. How much should we trust differences-in-differences estimates? *The Quarterly journal of economics*. 2004; 119(1):249–275. <https://doi.org/10.1162/003355304772839588>
90. Edelmann A, Kiley K, Kesintürk T, Bouklas I, Vaisey S. Quantifying the Importance of Change for Understanding Differences in Personal Culture. 2023;.
91. Underwood T, Kiley K, Shang W, Vaisey S. Cohort succession explains most change in literary culture. *Sociological Science*. 2022; 9:184–205. <https://doi.org/10.15195/v9.a8>
92. Kiley K, Vaisey S. Measuring stability and change in personal culture using panel data. *American Sociological Review*. 2020; 85(3):477–506. <https://doi.org/10.1177/0003122420921538>
93. Pinciotti CM, Orcutt HK. Understanding gender differences in rape victim blaming: The power of social influence and just world beliefs. *Journal of interpersonal violence*. 2021; 36(1-2):255–275. <https://doi.org/10.1177/0886260517725736> PMID: [29294886](https://pubmed.ncbi.nlm.nih.gov/29294886/)
94. Russell KJ, Hand CJ. Rape myth acceptance, victim blame attribution and Just World Beliefs: A rapid evidence assessment. *Aggression and Violent Behavior*. 2017; 37:153–160. <https://doi.org/10.1016/j.avb.2017.10.008>
95. Davies M, Rogers P, Whitelegg L. Effects of victim gender, victim sexual orientation, victim response and respondent gender on judgements of blame in a hypothetical adolescent rape. *Legal and Criminological Psychology*. 2009; 14(2):331–338. <https://doi.org/10.1348/978185408X386030>
96. Schneider LJ, Mori LT, Lambert PL, Wong AO. The role of gender and ethnicity in perceptions of rape and its aftereffects. *Sex Roles*. 2009; 60(5):410–421. <https://doi.org/10.1007/s11199-008-9545-9>
97. Boux HJ. Sexual assault jurisprudence: Rape myth usage in state appellate courts. *Georgetown University*; 2016.
98. Gravelin CR, Biernat M, Bucher CE. Blaming the victim of acquaintance rape: Individual, situational, and sociocultural factors. *Frontiers in psychology*. 2019; 9:2422. <https://doi.org/10.3389/fpsyg.2018.02422> PMID: [30719014](https://pubmed.ncbi.nlm.nih.gov/30719014/)
99. Grubb A, Turner E. Attribution of blame in rape cases: A review of the impact of rape myth acceptance, gender role conformity and substance use on victim blaming. *Aggression and violent behavior*. 2012; 17(5):443–452. <https://doi.org/10.1016/j.avb.2012.06.002>
100. Suarez E, Gadalla TM. Stop blaming the victim: A meta-analysis on rape myths. *Journal of interpersonal violence*. 2010; 25(11):2010–2035. <https://doi.org/10.1177/0886260509354503> PMID: [20065313](https://pubmed.ncbi.nlm.nih.gov/20065313/)
101. Anderson KB, Cooper H, Okamura L. Individual differences and attitudes toward rape: A meta-analytic review. *Personality and Social Psychology Bulletin*. 1997; 23(3):295–315. <https://doi.org/10.1177/0146167297233008>
102. Lambert AJ, Raichle K. The role of political ideology in mediating judgments of blame in rape victims and their assailants: A test of the just world, personal responsibility, and legitimization hypotheses. *Personality and Social Psychology Bulletin*. 2000; 26(7):853–863. <https://doi.org/10.1177/0146167200269010>