American Economic Journal: Applied Economics 2024, 16(1): 314–350
https://doi.org/10.1257/app.20210435

Gender Attitudes in the Judiciary: Evidence from US Circuit Courts†

By Elliott Ash, Daniel L. Chen, and Arianna Ornaghi*

Do gender attitudes influence interactions with female judges in US circuit courts? In this paper, we propose a judge-specific measure of gender attitudes based on use of gender-stereotyped language in the judge's authored opinions. Exploiting quasi-random assignment of judges to cases and conditioning on judges' characteristics, we validate the measure showing that higher-slant judges vote more conservatively in gender-related cases. Higher-slant judges interact differently with female colleagues: they are more likely to reverse lower court decisions if the lower court judge is a woman than a man, are less likely to assign opinions to female judges, and cite fewer female-authored opinions. (*JEL* D91, J16, K41)

Women are underrepresented at the top of the legal profession. Although since the 1990s close to 45 percent of law school graduates have been female (Croft 2016), women still compose only 20 percent of equity partners in large law firms (National Association of Women Lawyers 2019) and 30 percent of state and federal judgeships (George and Yoon 2019; Root 2019). In this paper, we focus on US Circuit Courts of Appeals and explore a possible explanation for this disparity: differential treatment of female judges on the part of their colleagues due to variation in gender attitudes.

Attitudes toward social groups—most notably women and racial minorities—are highly predictive of judgments and choices (Bertrand, Chugh, and Mullainathan 2005). Attitudes matter even in high-stakes settings such as physician treatments (Green et al. 2007), voting (Friese, Bluemke, and Wänke 2007), hiring decisions (Rooth 2010), employer-employee interactions (Glover, Pallais, and Pariente 2017), and teacher effectiveness (Carlana 2019). If gender attitudes imply differential treatment of female judges as well, they might play a role in explaining the representation gap of women in the judiciary.

*Ash: ETH Zurich (email: ashe@ethz.ch); Chen: Toulouse School of Economics (email: daniel.chen@iast.fr); Ornaghi: Hertie School (email: ornaghi@hertie-school.org). Ilyana Kuziemko was coeditor for this article. We thank Jacopo Bregolin, David Cai, Christoph Goessmann, Ornelie Manzambi, and Seraina Zuber for helpful research assistance. We thank Ryan Copus, Arthur Dyevre, Mirko Draca, Patrick Kline, Emily Leslie, David Mimno, Amanda Pallais, and Evan Soltas as well as conference and seminar participants for helpful comments and suggestions. Funding from the Russell Sage Foundation is gratefully acknowledged (grant number G-1905-16078).
† Go to https://doi.org/10.1257/app.20210435 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

The major challenge to investigating these issues is that the measures of gender attitudes traditionally used in the social sciences, such as Implicit Association Tests, are not available for circuit court judges.[1] We address this challenge by proposing a novel measure of gender attitudes that exploits a unique feature of our setting—the large corpus of written text that is available for appellate judges—and the idea that text can provide important insights into human social psychology (Jakiela and Ozier 2021). In particular, we draw on recent developments in natural language processing (NLP) and proxy judges' attitudes toward gender by measuring use of gender-stereotyped language in their writing (Pennington, Socher, and Manning 2014; Caliskan, Bryson, and Narayanan 2017; Antoniak and Mimno 2018; Kozlowski, Taddy, and Evans 2019). That is, we develop a measure of *gender slant* based on how strongly judges associate men with careers and women with families in the opinions that they write.

The key NLP technology powering our approach, word embeddings, is an algorithm that distributes words in a vector space based on their co-occurrence in a corpus (Mikolov et al. 2013; Pennington, Socher, and Manning 2014). We use word embeddings because they are a language representation that preserves semantic relationships. First, words with similar meaning have similar representations: their vectors are close together in the space. But in addition to this, the relative position of word vectors also conveys meaning. For example, male and female words (e.g., **man**, **woman** or **king**, **queen**) tend to hold similar relative positions to each other, and as a result we can identify a gender dimension in the space (equivalent to taking a step in the "male" direction) by taking the vector difference between male and female words. More generally, dimensions induced by word vector differences can be used to identify cultural concepts (Kozlowski, Taddy, and Evans 2019).

We operationalize these ideas to identify gender attitudes in language. If men are associated with career and women are associated with family, the relative position of male-to-female words should be similar to the relative position of career-to-family words. That is, the gender dimension (**male** − **female**) should be similar to the dimension representing stereotypical gender roles (**career** − **family**). More precisely, we measure the intensity of gender attitudes by looking at the cosine similarity between the gender dimension (**male** − **female**) and the stereotypical dimension (**career** − **family**). If the cosine similarity between the two dimensions is high, the corpus uses stereotyped language: men are associated with career, while women are associated with family. If the correlation is around zero, stereotyped language is not present in the text.

In the empirical analysis, we exploit the relative strength of the association to proxy for differential gender attitudes of judges. To construct a judge-specific gender slant measure, we consider the majority opinions authored by a given judge as a separate corpus. We then train embeddings for each judge, which allows us to calculate

[1] To the best of our knowledge, there exist only two papers that have collected IAT scores for judges, and neither of them links the scores to the real-world behavior of these judges. Rachlinski et al. (2009) measure implicit bias against African Americans for 133 state and local trial judges, finding some evidence that more biased judges prefer harsher sentences when primed with hypothetical cases about African Americans in a vignette experiment. Levinson, Bennett, and Hioki (2017) collect implicit bias scores against Asians and Jews for 239 federal and state judges, and find that judges biased against Jews display different hypothetical sentencing behavior.

a judge-specific gender slant measure. To ensure that we obtain a high-quality representation notwithstanding the potentially small size of a judge's corpus, we train embeddings on different bootstrap samples following Antoniak and Mimno (2018) and define gender slant as the median value of the measure across the samples. Using this method, we calculate the gender slant of 139 circuit judges.

As a measure of stereotyped gender attitudes, gender slant behaves convincingly in a number of validation checks. Descriptively, we find that female and younger judges display lower gender slant and that having a daughter reduces gender slant. Getting into the language, we find that lower gender slant is associated with more frequent use of gender-neutral pronoun constructions, such as "he or she." Finally, in a human annotation exercise performed by a reader with legal training, we find that judges with higher slant tend to express less empathy toward women in their writing.

This approach to measuring gender attitudes is related to a growing literature using word embeddings to analyze bias in text. Bolukbasi et al. (2016) and Caliskan, Bryson and Narayanan (2017) demonstrate biases in general web corpora in terms of gender and a number of other dimensions. Garg et al. (2018) train separate embeddings by decade using the Google Books corpus and show that gender associations track demographic and occupational shifts. Two recent papers apply word embeddings to the judicial setting: Rice, Rhodes and Nteta (2019) detect racial slant in a corpus that includes US circuit court opinions, and Ash, Chen and Galletta (2022) document sentiment toward social groups in a similar setting.[2] We build on this literature to construct author-specific measures of gender slant that we link to real-world behaviors.

There exist of course a number of alternative measures we could have used to proxy for these judges' gender attitudes. First, we could have asked human evaluators to qualitatively score the writing of judges. In addition to potential concerns regarding the subjectivity of the coding, the main problem with this approach is its prohibitive cost, considering that our corpus includes over 14 million sentences. Second, we could have proxied for gender attitudes by looking at judges' rulings or votes. While in principle this is indeed a reasonable approach, the small number of gender-related votes for which this information is available means that voting rates in gender decisions are unlikely to produce an empirically useful measure of gender attitudes for circuit judges.

The central research question of the paper is whether judges with different gender slant interact with female judges differently. Our empirical strategy relies on the fact that, within each circuit and year, judges are quasi-randomly assigned to cases. This

[2] Ash, Chen, and Galletta (2022) measure text-based sentiment (positive/warm versus negative/cold) toward 19 social groups (e.g., Black, White, women, Catholics, businesses, etc.) in US circuit court opinions. The sentiment measure is based on embedding models broadly related to the ones used in this paper, although the implementation is different. Specifically, the authors use Doc2Vec (Le and Mikolov 2014), an algorithm to represent documents in an embedding space, and apply it to each sentence appearing in the opinions. For each sentence, they compute the proximity of the vector representing the sentence to a vector representing positive/negative sentiment and the proximity to vectors representing the different social groups. They obtain each opinion's sentiment toward each social group by taking the average of the sentence-level sentiment weighted by the sentence's proximity to the group. This methodology differs from ours in the use of different embedding models (a single embedding model that allows for paragraph representation versus different embedding models for each judge as we do here) and in the objective of creating a case-level sentiment measure rather than understanding judge-level attitudes.

ensures that higher-slant judges do not self-select into cases systematically based on the expected case outcome: cases assigned to higher- and lower-slant judges are comparable. In addition to this, we condition on detailed judges' characteristics to provide supportive evidence that our strategy identifies the effect of judges having different levels of slant, and not judges differing along some other characteristic such as gender or conservative ideology.

We study whether higher-slant judges treat female judges differently focusing on three sets of interactions: reversals of district court decisions, opinion assignment, and citations. We find evidence supporting this hypothesis. First, we show using a differences-in-differences design that higher-slant judges are more likely to vote to reverse lower court decisions authored by female district judges with respect to those authored by male district judges. The magnitude of the effect is sizable: being assigned a judge with a one standard deviation higher gender slant increases the probability that a female relative to a male district judge is reversed by 1 percentage point, which corresponds to around 5 percent of the outcome mean. Second, we consider the decision of a senior judge tasked with assigning the writing of the majority opinion to a panel member. Assigning judges with a one standard deviation higher gender slant are less likely to assign opinion authorship to a female judge by 1.7 percentage points, or 4.5 percent of the outcome mean. Finally, higher-slant judges are also less likely to cite opinions of female judges, although identification is weaker and the result is less robust. To the extent that these outcomes are relevant for future opportunities, interacting with judges with different gender attitudes as measured by gender slant has the potential to hinder the career progression of female relative to male judges.

To strengthen the interpretation that gender slant is a proxy for gender attitudes, we validate the measure by studying whether higher-slant judges take different decisions in gender-related cases. Consistent with the proposed interpretation, we find that judges with higher slant tend to vote more conservatively in gender-related issues (that is, against expanding women's rights). The magnitude of the effect is large: being assigned a judge with a one standard deviation higher gender slant increases the probability of voting against expanding women's rights by 4.1 percentage points, which corresponds to a 7 percent increase over the outcome mean. In addition to providing a validation of the measure, this is an interesting result per se, as it shows that by affecting how precedent is set, gender attitudes have the potential to impact real-word outcomes even outside the judiciary (Chen and Yeh 2014, 2020a; Chen and Sethi 2018).

Finally, we investigate whether higher-slant judges also respond differently to nongender characteristics. While we do find that higher-slant judges tend to vote more conservatively in nongender-related but ideologically divisive cases, the effect of being assigned a higher-slant judge is smaller than in gender-related cases. In addition, we find that gender slant has little to no effect on how judges interact with colleagues who were appointed by a Democratic president, who are minorities, or who are in a different age group. Overall, these results support the view that gender slant captures attitudes that are specific to gender.

If we think of courts as a workplace, this paper speaks to the literature on how gender shapes the labor market outcomes of women and why (see among others

Bohren, Imas, and Rosenberg 2019; Bordalo et al. 2019; Card et al. 2019; Sarsons 2019; Hengel 2022), in particular for women employed at the top end of the earnings distribution (Bertrand, Goldin, and Katz 2010; Bertrand 2013; Bursztyn, Fujiwara, and Pallais 2017). Despite the richness of the data, the setting we study is quite novel: there is a scarcity of existing work that takes this approach toward the courts in particular to study the potential for gender discrimination toward female judges.[3] In addition, we contribute to this literature by providing evidence that gender attitudes might play a direct role in determining differential labor market outcomes for men and women, even in high-stakes environments such as appellate courts.

More generally, this paper contributes to the growing literature demonstrating the importance of attitudes in decision making. For example, Glover, Pallais, and Pariente (2017) show that attitudes regarding minorities influence manager-employee interactions in a way that impacts performance, and Carlana (2019) shows that teachers with stereotypical views of gender negatively impact the test scores and future scholastic careers of female students. Our novel text-based measure of gender attitudes allows us to study the role these attitudes play in determining the behavior of high-skilled professionals, who might otherwise be hard to reach through surveys or tests traditionally used in the literature.

In addition, by studying the effect of being assigned judges with different levels of slant on voting, this paper builds on the literature on the determinants of judicial decisions, which broadly shows that demographic and ideological characteristics of judges matter (Sunstein et al. 2006; Boyd and Spriggs II 2009; Kastellec 2013; Cohen and Yang 2019). Recent work has shown that in criminal cases judges systematically demonstrate racial (Rehavi and Starr 2014; Arnold, Dobbie, and Yang 2018) and gender (Starr 2014) disparities in their sentencing decisions. Instead of inferring bias from the decisions themselves, we take a different approach by considering how gender attitudes impact voting. The existing evidence on gender attitudes in the judiciary is limited to studies that look at the correlation between attitudes measured by Implicit Association Tests and decisions in hypothetical scenarios (Rachlinski et al. 2009; Levinson, Bennett, and Hioki 2017). Instead, our measure allows us to link gender attitudes to real world outcomes.

The remainder of the paper is organized as follows. Section I describes our measure of gender slant, and Section II discusses the empirical strategy and presents additional data sources. Section III provides descriptive statistics. Section IV shows the relationship between higher-slant judges and decisions, while Section V describes the relationship between higher-slant judges and interactions with female judges. Finally, Section VI concludes.

I. Gender Slant: Measuring Gender Attitudes in Text

We begin by constructing a measure of how gender is characterized in the language used by judges. Specifically, our measure aims to capture the strength of

[3] The one exception is contemporaneous work by Battaglini, Harris, and Patacchini (2023), who show that random exposure to female judges on panels increases the probability of hiring a female clerk.

the association between gender identifiers and stereotypical gender roles, with men being more career-oriented and women being more strongly associated to family.

A. *Word Embeddings*

In order to construct our measure, we use word embeddings, a language modeling technique from NLP that relies on word co-occurrence to create a representation in a low-dimensional Euclidean space in such a way as to preserve semantic meaning (Mikolov et al. 2013; Pennington, Socher, and Manning 2014).

Consider the simplest way of representing language. For a given vocabulary V, one possibility is to represent words as one-hot-encoded vectors, which are vectors with all values equal to 0 except the one entry corresponding to the word itself. This approach presents two issues. First, the dimensionality of the vector space grows linearly in the size of the vocabulary, as the one-hot-encoded vectors are V-dimensional by construction. Second, it is impossible to infer anything about the relationship between words in the resulting space: all word vectors are orthogonal to each other. Word embeddings offer a solution to both issues. The word representations are low-dimensional—in our case, 300 dimensional—dense vectors that can accommodate large vocabularies and corpora without increasing dimensionality. In addition, the positions of word vectors in the space encode relations between words.

Word embeddings encode semantic meaning in two principal ways. First, distance in the word embedding space conveys semantic similarity between words. This is because the position of a word's representation in the vector space is assigned based on the context the word appears in: words that appear frequently in the similar context have representations close to each other in the space, while words that rarely do have representations that are far apart. Importantly, given that a vector's position is defined based on appearance in given contexts, word embedding distance can identify that two words are similar even if they do not necessarily often appear together, as long as the neighboring words tend to be similar.

Second, the relative position of word vectors in the space also conveys meanings. For example, as Figure 1 illustrates, male and female words tend to be in similar relative positions to each other. This means that taking the difference between word vectors representing male and female words can identify a gender dimension, equivalent to taking a step in the male direction.

B. *GloVe Embeddings Implementation*

The specific model we use is Global Vectors for Word Representation (Pennington, Socher, and Manning 2014). GloVe is a weighted least squares model that trains word vectors on global co-occurrence counts. GloVe first computes a global co-occurrence matrix, which reports the number of times two words have occurred within a given context window. It then obtains word vectors $w_i \in \mathbf{w}$ to minimize the following objective function:

$$J(\mathbf{w}) = \sum_{i,j} f(X_{ij}) \left[w_i^T w_j - \log(X_{ij}) \right]^2, \tag{1}$$

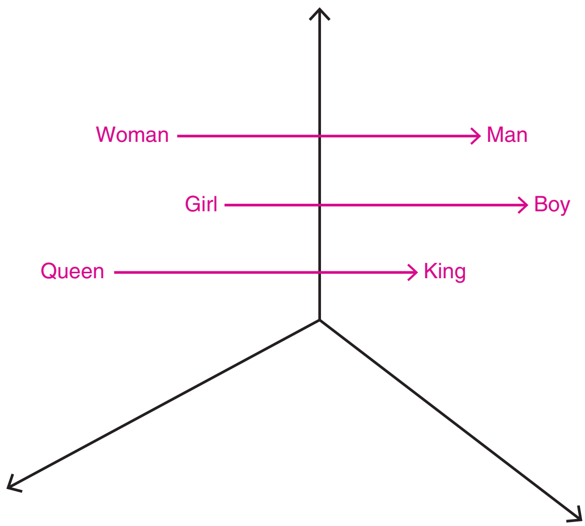

Figure 1. Identifying a Gender Dimension in a Vector Space

Note: The figure exemplifies how the gender dimension can be identified in the vector space based on the difference between vectors representing male and female words.

where X_{ij} is the co-occurrence count between words i and j and $f(\cdot)$ is a weighting function that serves to down-weight particularly frequent words.[4] The objective function $J(\cdot)$ effectively trains the word vectors to minimize the squared difference between the dot product of the vectors representing two words and their empirical co-occurrence in the corpus. Our GloVe implementation minimizes $J(\cdot)$ by stochastic gradient descent.

The two key hyperparameters for GloVe are the dimensionality of the vectors and the window size for computing co-occurence statistics. Previous experiments by NLP researchers suggest that increasing dimensionality beyond 300 has negligible improvements for downstream tasks (Pennington, Socher, and Manning 2014; Rodriguez and Spirling 2022), so we follow that literature and train 300-dimensional vectors. In turn, we choose a standard window size of 10, a middle ground between shorter windows, which would tend to capture syntactic/functional relations between words, and longer windows, which tend to capture topical relations between words.[5]

A practical feature of GloVe is that the algorithm goes through the full corpus only once in order to build the initial co-occurrence matrix. This feature accounts

[4] Following the original GloVe paper (Pennington, Socher, and Manning 2014), we use the following weighting function:

$$f(x) = \begin{cases} (x/x_{max})^a, & \text{if } x < x_{max}; \\ 1, & \text{otherwise}; \end{cases}$$

where $a = 3/4$.

[5] Online Appendix Figure 1, panels A and B show that embeddings trained using 5, 10, or 15 word windows produce highly correlated embeddings. The same is true for 100- and 300-dimensional embeddings (panel C).

for the considerable improvements in training time compared to other popular word embeddings algorithms such as word2vec (Mikolov et al. 2013), while obtaining embeddings of comparable quality (Pennington, Socher, and Manning 2014). Given that our approach requires the training of a large number of separate embeddings, this is a particularly attractive feature for our application.

C. *Word Vectors and Gender Slant*

We use word embeddings to identify cultural dimensions in language (Kozlowski, Taddy, and Evans 2019). As mentioned in the previous subsection, a key feature of word embeddings is that the direction of the difference between word vectors in the space conveys meaning. Consider the vector representing the word **man** and the vector representing the word **woman**. The vector difference between the two, i.e., the vector identified by **man** $-$ **woman**, identifies a dimension in the space that corresponds to a step in the male direction.

In practice, this is true for **boy** $-$ **girl**, **he** $-$ **she**, and so on: we can identify a gender dimension in the space by taking the difference between the average normalized vector across a set of male words and the average normalized vector across a set of female words:

$$(2) \qquad \mathbf{male} - \mathbf{female} = \frac{\sum_n \mathbf{male\,word}_n}{|N_{male}|} - \frac{\sum_n \mathbf{female\,word}_n}{|N_{female}|},$$

where $|N_{male}|$ is the number of words used to identify the male dimension and $|N_{female}|$ is similarly defined.

A desirable feature of the Euclidean geometry of the vector space and the gender dimension is that other words meaningfully project onto it. It is then possible to understand the connotation of other words along the gender dimension by looking at the cosine of the angle between the vector representing the word and the dimension itself. Formally, we use the cosine similarity, defined as:

$$(3) \qquad \mathrm{sim}(\mathbf{x},\mathbf{y}) = \cos(\theta) = \frac{\mathbf{x}\cdot\mathbf{y}}{\|\mathbf{x}\|\,\|\mathbf{y}\|} = \frac{\sum x_i y_i}{\sqrt{\sum x_i^2}\sqrt{\sum y_i^2}},$$

where $\mathbf{x}$ and $\mathbf{y}$ are nonzero vectors, θ is the associated angle, and $\|\cdot\|$ is the 2-norm. Note that $\mathrm{sim}(\mathbf{x},\mathbf{y})$ varies between -1 and $+1$. Continuing with the example, words with male (female) connotations—e.g., male (female) first names—are going to be positively (negatively) correlated with the gender dimension defined by **male** $-$ **female**.

We put these dimensions in service of constructing a gender slant measure. The goal is to capture the strength of the association between gender and stereotypical attitudes identifying men more closely with career and women with family. Word embeddings provide an intuitive metric. If men are associated with career and women are associated with family in the corpus, the relative positions of male-to-female words should be similar to the relative positions of career-to-family words. Specifically, we use the cosine similarity between the vector representing

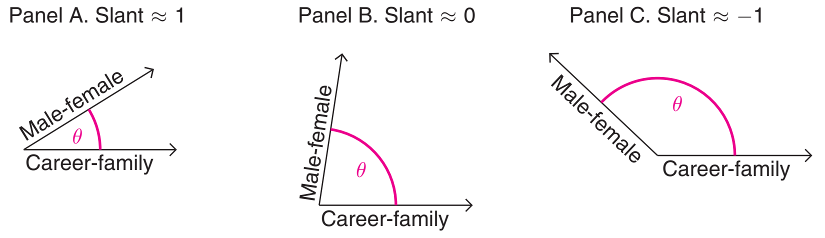

Figure 2. Measuring Gender Attitudes Using Cosine Similarity

Note: The figure exemplifies how gender slant varies depending on the relative position of the gender and the career-family dimensions in the vector space.

the gender dimension, defined by **male** – **female**, and the vector representing the career-family dimension, defined by **career** – **family**. The statistical similarity summarizes how closely related the gender and stereotypical dimension are in the space, as illustrated in Figure 2. When the two concepts are strongly associated in a corpus, the two vectors are close together ($\theta \approx 0$), and the slant measure is close to 1 (Figure 2, panel A). If there is no association between the two, then $\theta \approx 90°$, and the slant measure is 0 (Figure 2, panel B). Finally, if the concepts are negatively associated in a corpus (e.g., male is associated to family and female is associated to career), the two vectors are far apart ($\theta \approx 180°$), and the slant measure will tend to -1 (Figure 2, panel C).

For this task, there are many potential combinations of male, female, career, and family words that could be used to identify the gender and career-family dimension. We select the word sets to identify the dimensions using the following procedure. First, we identify potential word sets using Linguistic Inquiry and Word Count Dictionaries, which provide a human-validated list of words and word stems that correspond to certain concepts. We use the word sets for male, female, work and family. From these word sets, we eliminate words that could be ambiguous or have specific legal meanings in our setting (e.g., tribe, tribes for family; line, situation, trade for work). From each list, we then select the ten most frequent words in the full judicial corpus.[6] Table 1 reports the resulting word sets. In addition, we show how the results change when selecting the top 5 to top 15 most frequent words and when dropping one word at a time in the career versus family dimension.

We focus on the stereotypical association between men and career versus women and family as opposed to other associations typically studied in the literature—for example, the association between men and science versus women and arts, and between men and positive attributes versus women and negative attributes—for the following reasons. First, words related to sciences and arts do not frequently appear in the corpus, which makes it difficult to identify the science/art dimension

[6] We select words using these procedures as opposed to the word sets usually used in the gender-career IAT because we want to ensure that we are using words that meaningfully define gender in judicial language. For example, first names are rarely used in legal language as opposed to gender pronouns, which would introduce substantial noise in the slant measure.

Table 1—Word Sets

Male	his, he, him, mr, himself, man, men, king, male, fellow
Female	her, she, ms, women, woman, female, herself, girl, girls, queen
Career	company, inc, work, business, service, pay, corp, employee, employment, benefits
Family	family, wife, husband, mother, father, parents, son, brother, parent, brothers

in the space (see online Appendix Table 1). Second, we only find limited evidence that the full judicial corpus presents a stereotypical association between men and positive attributes versus women and negative attributes, which suggests limited scope for this measure to capture judge-specific variation in gender attitudes. In fact, online Appendix Figure 2 shows that in the full judicial corpus, the cosine similarity between the gender and the positive-negative dimension is generally smaller and closer to zero with respect to the cosine similarity between the gender and the career-family dimension. When we perform a permutation test following Caliskan, Bryson, and Narayanan (2017), we indeed find p-values lower than 0.10 in only 3 of the 24 bootstrapped embeddings trained on the full judicial corpus.

D. *Measuring Judge Gender Slant*

Our goal is to produce measures of gender attitudes in the writing of circuit court judges. Our starting corpus is the universe of published opinions in circuit courts for the years 1890–2013, which consists of 380,000 opinions in 13 courts from Bloomberg Law (Bloomberg Law 2022).

As a preprocessing step, we clean that text and exclude punctuation and numbers, although we retain hyphenated words. To avoid the word vectors being case sensitive, we transform all words to be lowercase. We then retain only the most common 50,000 words in all judicial opinions. Opinions are separated into sentences using punctuation, and each sentence is further tokenized into words. These tokenized sentences are the starting point of the model.

To obtain judge-specific, gender slant measures, we take the set of majority opinions authored by each judge as a separate corpus. We train separate GloVe embeddings on each judge's corpus, and we used the resulting vectors to compute the gender slant measure as described in subsection C. To ensure convergence, we train vectors for 20 iterations with a learning rate of 0.05.

Creating judge-specific embeddings implies the use of relatively small corpora, which is potentially a problem given that word embeddings perform best when they are trained on large collections. To address this issue, we follow the approach suggested by Antoniak and Mimno (2018) and train embedding models on 25 bootstrap samples of each judge corpus. Specifically, we consider each sentence written by a judge as a document, and then create a corpus by sampling with replacement from all sentences. The number of sentences contained in the bootstrapped sample is the same as the total number of sentences in the original judge corpus. We then calculate our measure for all bootstrap samples and assign to each judge the median value of the measure across the samples. Given that embeddings trained on small corpora tend to be sensitive to the inclusion of specific documents, the bootstrap procedure produces more stable results. In addition, bootstrapping ensures stability with respect to the initialization of the word vectors, a potential concern given that GloVe presents a non-convex objective function (Rodriguez and Spirling 2022). Importantly, this constraint on corpus size is why we are not able to construct time-varying measures of gender slant.

Even following the bootstrapping procedure, we might still worry about the quality of the judge-specific embeddings, and in particular, whether they are able to capture meaningful information about gender. To ensure that the judge-specific embeddings are able to do so, we compute the cosine similarity between the gender dimension and each of the vectors representing the most common 25 male and female names according to the 1990 census for each judge and bootstrap sample.[7] We then regress a dummy for whether the name is male on the median cosine similarity between the vector representing the name and the gender dimension across bootstrap samples, separately for each judge.

Figure 3 shows the cumulative distribution of the t-statistics resulting from these regressions for sets of judges with a different number of tokens. The figure shows that, for judges with a small number of tokens, the t-statistics are rarely above conventional significance level and are at times even negative. As the number of tokens increases, the t-statistics start to become significant and are never lower than zero, showing that the gender dimension identified in the embeddings does indeed contain meaningful gender information. Based on these experiments, all of our main results focus on the 139 judges whose corpus includes at least 1,500,000 tokens.[8] Out of the 139 judges with 1.5 million tokens, 12 percent are women (17 judges) and 42 percent were appointed by a Democratic president (58 judges).

The relatively small sample of judges for which we are able to define the measure might raise concerns related to the external validity of the findings. However, the judges in our sample do not appear to be highly selected as compared to other circuit judges: they are not significantly different in terms of party of appointing president, gender, and race, although they are more likely to be born after 1920 (see online Appendix Table 4).

E. *Variation in Judge Gender Slant*

Figure 4 shows the distribution of gender slant across judges in the sample. Gender slant is almost always positive, reflecting that most judges display a stereotypical association between men and career and between women and family. However, there is significant variation in how strong this association is. Our

[7] We get the most common 25 male and most common 25 female first names from the "Frequently Occurring Surnames from US Census (1990).

[8] For comparison, around 1,000 judges served in circuit courts from 1890 to 2013. The online Appendix discusses in detail the robustness of our results to decreasing or increasing the tokens threshold. Consistent with the idea that our measure does not do a good job at capturing gender additudes for corpora that are too small, our results are not consistently robust to decreasing the token threshold below 1,500,000 tokens.

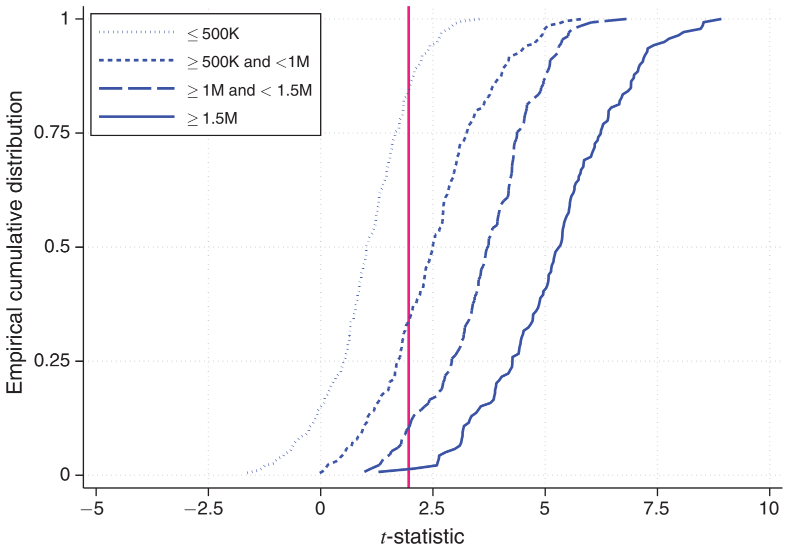

Figure 3. Judge-Specific Word Embeddings Capture Gender Information

Notes: The graph shows that the gender dimension in judge-specific word embeddings captures gender information when the corpus of the judge is sufficiently large. In particular, we test whether male first names have a higher cosine similarity with the gender dimension than female names in the judge-specific embeddings. The graph reports the cumulative distribution of the t-statistics resulting from a series of regressions of an indicator variable equal to 1 if the name is male on the median cosine similarity between the vector representing the name and the gender dimension across bootstrap samples for each judge, by number of tokens included in the judge's corpus.

empirical analysis leverages this variation to relate judges with higher or lower gender slant to judicial outcomes.[9]

To put the measure in perspective, the vertical lines on the same graph show the degree of gender slant for a sample of recent US Supreme Court justices for which we were able to measure gender slant. Interestingly, the conservative Justice Antonin Scalia has the highest slant of the group, while more liberal judges such as Justice Anthony Kennedy have relatively lower slant.[10]

Where is the variation in gender slant across judges coming from? If judges are writing a legal opinion to support the decision in a case, do they still have latitude in their writing? Judges and legal scholars who have studied this topic have documented qualitative and quantitative differences in judge writing (see Posner 1995, 2008). In the US court system especially, judges tend to express their individual

[9] Because we cannot ground our gender slant measure in an external measure of gender attitudes, we are unable to identify an objectively "correct" gender slant measure in this context. Nonetheless, we believe that our results can still provide important insights on the role of gender attitudes on judicial behavior exploiting the variation across judges in the relative strength of the stereotypical association. Because the units of the measure are not easily interpretable, we standardize it in the empirical analysis.

[10] The gender slant of Supreme Court justices is measured by training embeddings on the US Supreme Court majority opinions that were authored by these judges. Finding judges with corpora of sufficient size is even more difficult in this setting: the figure shows the gender slant of justices with at least 1,000,000 tokens in their corpus. Unfortunately, none of the female justices (e.g., Ginsburg and O'Connor) reached the minimum token threshold. The text of Supreme Court opinions is from LexisNexis (LexisNexis 2022b).

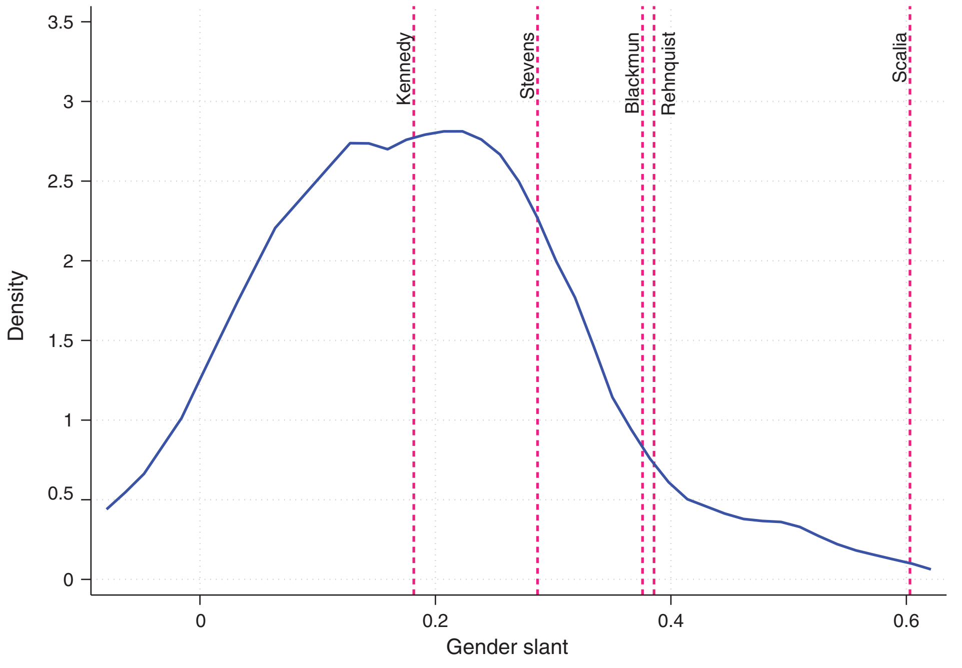

Figure 4. Distribution of Gender Slant

Notes: The graph shows the distribution of the gender slant measure for the 139 judges in the sample. The vertical lines indicate the gender slant of five Supreme Court justices who were appointed after 1970 and have a corpus of at least 1,000,000 tokens. The gender slant of Supreme Court justices is measured by training embeddings on the Supreme Court majority opinions that were authored by them. Gender slant is the cosine similarity between the gender and the career-family dimensions.

personalities and attitudes in their opinions. These attitudes come out not just in the ruling and the legal arguments in support, but also, importantly, in how the facts and moral and policy issues are framed.

In particular, conditional on the facts and legal boundaries, a judge's choice of words can be expressive of attitudes toward gender. A familar example would be the choice of pronouns when discussing hypotheticals. A judge could either default to the masculine pronoun ("he"), as is the historical convention, or could use a feminine pronoun ("she"), or they could use a gender-neutral expression (e.g., "he or she"). The latter, especially, tends to express commitment to gender-equality norms. Consistent with this intuition, online Appendix Figure 3 shows that higher-slant judges are less likely to use gender-neutral pronoun constructions.

To add texture to these quantitative metrics, we asked a law student to read a selection of judicial opinions and annotate examples of passages demonstrating empathy for women. Online Appendix Table 2 provides a selection of these annotated snippets. They are interesting to read and illustrate how the personalities of judges are expressed in their opinions. The passages are notably expressive and not just bland technical writing, offering some qualitative evidence in our setting for judicial discretion in writing style.

F. *Discussion of Alternative Measures*

There are (perhaps many) other potential approaches to measure gender attitudes in judicial language. First, we could have had human evaluators qualitatively score the writing of judges. These coders could have done a deep reading of the text or a subjective coding of important themes (Glaser and Strauss 2017). This qualitative approach is somewhat subjective and therefore lacks a rigorous method of replication (Ricoeur 1981; DiMaggio 1997). The binding constraint to the annotation approach, however, is the cost: it would be prohibitively expensive to implement on a large scale such as we have in our corpus, which contains over 14 million sentences. In this sense, our automated approach can be seen as a way of efficiently extracting information regarding how gender is talked about in large corpora.

To compare our approach to human coding, we tasked a law student with the following annotation exercise (Ash, Chen, and Ornaghi 2022). The annotator read two randomly paired judicial opinions and assessed which opinion demonstrated more empathy toward women. To ensure that gender was a relevant element of the case, and to hold the direction of the ruling constant, the pairs were selected among gender-related cases that were decided in favor of expanding women's rights.[11] The pairs contained a higher-slant judge's opinion and a lower-slant judge's opinion (the top/bottom 10 percent of the slant distribution, to maximize power). Consistent with the idea that gender slant captures gender attitudes that are expressed in text, online Appendix Figure 4 shows that opinions authored by higher-slant judges are less likely to be identified as having more empathic language. In addition, online Appendix Figure 5 shows that they are less likely to include at least one empathic statement. Even among cases that are on a close topic (i.e., gender-related cases) and that are decided in the same direction (i.e., in favor of expanding women's rights), there is significant variation in the language that judges with higher and lower slant use.

While this annotation exercise provides some reassurance about our embedding-based measure, it also illustrates the infeasibility of hand coding a comparable dataset. The law student needed about 30 minutes to annotate each case. At $30 per hour, it would be prohibitively expensive to annotate a sufficiently large corpus to do a comparable empirical analysis.

A second alternative to our embedding-based measure of gender attitudes would be to construct a count-based measure based on word or phrase frequencies. A count-based measure has intuitive appeal given that the language representation provided by word embeddings is itself built on proximate co-occurrence of words. Thus, co-occurrence statistics of male/female words with career/family words can be thought of as the building blocks of the gender slant measure. Consistent with this intuition, the relative co-occurrence of male words with career words relative to the co-occurrence of female words with family words is positively associated with the gender slant measure (online Appendix Figure 6). In online Appendix Table 3, we report ten randomly selected sentences that present these co-occurrences for reference.

[11] Section II and IV provide more information on gender-related cases.

Looking at the raw counts, however, we see that these explicit word co-occurrences are rare. For example, the median judge writes only 38 sentences where female and family words co-occur. This sparsity in the count-based measure is why a word embedding approach is preferrable, as embeddings do not require words to appear directly next to each other to register an association, and they are not strictly bound to the specified lexicon. As a result, embeddings are able to take into account more information contained in judges' writing. Instead, the count-based approach misses implicit and nuanced gender-stereotyped language and as a result provides a less precise measure.

On a different tack, it is also worth asking whether a linguistic analysis is needed at all. Why not just measure gender attitudes from the judge's rulings or votes? After all, in Section IV, we ask whether higher-slant judges vote differently in gender-related cases as a way of validating our measure. In principle, one could measure gender attitudes by directly computing the average decision direction (progressive or conservative) in gender-related cases. However this approach does not work well in our setting in practice. As we will discussed further, we are constrained by our judges showing up in datasets hand-coded for vote valence. The average judge has only 27 gender-decision votes, and many judges have fewer than 10. This sparsity of data points means that voting rates in gender decisions are unlikely to produce an empirically useful measure of gender attitudes.

II. Empirical Strategy and Additional Data Sources

A. *Empirical Strategy*

The main objective of the paper is to study whether judges displaying higher or lower levels of gender slant interact differently with female judges. The empirical strategy relies primarily on the quasi-random assignment of judges to cases, which means that higher- and lower-slant judges do not self-select into cases systematically based on the expected outcome. In addition, we condition on detailed judges' biographic characteristics to ensure that we are not confounding the effect of the judge having higher slant with other judge characteristics such as gender or conservative ideology.

Quasi-Random Assignment.—A major concern for identification is the endogeneous selection of judges to cases, as we might worry that higher-slant judges might decide which panels to sit on based on the expected outcome of the case. We are helped in this respect by the fact that in circuit courts, cases are quasi-randomly assigned to panels formed by three judges, which ensures comparability of cases seen by judges with higher and lower gender slant.

Qualitative research based on court documents and interviews with court officials describe the process as follows (Bowie, Songer, and Szmer 2014; Chen and Sethi 2018). Judges are chosen from a pool of 8–40 judges, depending on the specific circuit. Before oral arguments, available judges (including visiting judges) are assigned to cases following a random process, in recent years using computer programs. In some cases, randomization might be constrained by organizational considerations, such as accommodating judges' schedules. Importantly, in case of conflict of interest, judges might recuse themselves, but they would generally do so before random assignment.

Evidence supporting the quasi-random assignment of judges to cases has been provided by Chen and Sethi (2018), who show for a sample of gender discrimination cases that pretrial characteristics are uncorrelated with the demographic composition of the panel, and by Chen, Levonyan, and Yeh (2020), who show that the panel composition does not appear to be autocorrelated over time. Chilton and Levy (2015) test whether the ideological composition of panels is indeed random by comparing a simulated panel formation process with actual panels sitting together in different sessions of the court, and they find evidence of nonrandom assignment for four circuits (namely, the second, eighth, ninth, and DC Circuit). Because running their test requires detailed information on court calendars that are not available for the full period we study, we cannot replicate their test looking at the slant composition of panels. Nonetheless, we show that our results are robust to dropping the circuits where nonrandom assignment is suspected. Most importantly, based on extensive qualitative research, Chilton and Levy (2015) see the lack of random assignment as driven by organizational constraints, such as the need to ensure spacing of judicial assignments or to accommodate vacation schedules, and not by judges specifically selecting into panels, which would be more problematic in our setting.

Here, we provide two pieces of additional evidence supporting quasi-random assignment of judges to cases. First, we test whether higher-slant judges are systematically assigned to cases with different topics (including whether the case is gender related) or that were originally decided by district judges with different characteristics. We do so by regressing case characteristics on gender slant, judge demographic controls, and circuit-year fixed effects. Online Appendix Table 5 shows that gender slant does not appear to be systematically correlated with the case characteristics we consider. Although slant is negatively correlated with the district judge being female, this is in line with statistical chance. Consistent with this interpretation, a joint test that the effect of slanted judges fails to reject the null hypothesis (p-value: 0.162).

Second, we test whether the variation in case characteristics across judges that we see in our data is consistent with quasi-random assignment. Following Abrams, Bertrand, and Mullainathan (2012), we do this by comparing the variation we see in the data with the one resulting from simulating random assignment. Online Appendix Figure 7 shows, for each characteristic we consider, the distribution of the interquartile range of the mean characteristic across judges for 1,000 assignment simulations.[12] The vertical line shows the true value we observe in the data. Overall, the graphs show that the cross-judge variation we observe in the data is in line with the one we should expect given quasi-random assignment.

[12] For each characteristic, we estimate the actual interquartile range by taking the average value of the characteristic for each judge across all the cases that the judge is assigned to. Then, we simulate the process of quasi-random assignment by randomly assigning case characteristics to each judge. Because case characteristics vary substantially across circuits and years, we simulate the data to match the mean characteristics in each circuit and year. For each simulated dataset, we compute the average value of the characteristic for each judge and estimate the interquartile range of this value across the judges. We perform this exercise focusing on the 139 judges that are part of our main sample.

We interpret the evidence as supporting the fact that judges are quasi-randomly assigned to cases.

Conditioning on Observable Characteristics.—Random assignment of judges to cases implies that the effect of being assigned a higher- or lower-slant judge is well-identified, but we might still worry of this effect being confounded by other judge characteristics, such as judge gender or conservative ideology. We address this issue by exploiting detailed information on the demographic characteristics of these judges, which allows us to directly control for other characteristics of the judge that might correlate both with gender slant and with behavior. In online Appendix Table 6, we assess this assumption using the method from Oster (2019).

B. *Additional Data Sources*

This paper combines four principal data sources, in addition to the text data used to construct the gender slant measure for each judge described in the previous section. Table 2 shows descriptive statistics for the data used in the analysis.

Judges' Demographic Characteristics.—The data on judge characteristics are from the Federal Judicial Center (Federal Judicial Center 2022) and the Appeals and District Court Attributes Data (Gryski and Zuk 2008a,b). The final dataset has information on gender, party of appointing president, region of birth, cohort of birth, religion, law school attended, and whether the judge had federal experience prior to being appointed to an appellate court. It also includes information on the full history of federal court appointments held by the judge.

Judicial Decisions.—To study judicial decisions, we use two existing legal datasets with hand-coded vote direction and topic from Epstein, Landes, and Posner (2013b) and Glynn and Sen (2015b). The datasets were originally constructed by searching for cases related to a given set of topics, and then having research assistants read through the opinions to code whether each judge's vote was liberal or conservative.[13,14] For gender-related issues—namely, reproductive rights, gender discrimination, and sexual harassment—a liberal vote corresponds to a vote in support of extending women's rights. The analysis pools the two datasets. In addition, we have information on circuit court cases from the US Court of Appeals datasets (Songer 2008; Kuersten and Haire 2011). These data include a 5 percent random

[13] The need to code the direction of the vote (in favor or against expanding women's rights) is why we are constrained to using these preexisting datasets, as we do not know the "directionality" of votes in the larger circuit court dataset that we use in the analysis regarding the treatment of female judges.

[14] In particular, the Epstein, Landes, and Posner (2013b) dataset contains all published opinions related to abortion, the Americans with Disabilities Act, affirmative action, campaign finance, capital punishment, the contracts clause, criminal appeals, environmental regulation, federalism, piercing the corporate veil, race discrimination, sex discrimination, sexual harassment, and takings. The dataset is updated until 2008, but the starting years of the dataset vary by issue, ranging from 1982 for abortion to 1995 for capital punishment. The original dataset was constructed by searching for cases related to each issue backward from the present, and stopping when a sufficient number of cases was reached for that issue. The Glynn and Sen (2015b) data contain all published and unpublished opinions from 1996 to 2002 that contain the words "gender," "pregnancy," or "sex" in the case headings. When the two datasets are pooled, we drop duplicate cases present in both datasets.

Table 2—Descriptive Statistics

	N (1)	Mean (2)	SD (3)	N (4)	Mean (5)	SD (6)
Panel A. Judges						
Gender slant	139	0.186	0.132			
Democrat	139	0.417	0.495			
Female	139	0.122	0.329			
Minority	139	0.079	0.271			
Born in 1920s	139	0.252	0.436			
Born in 1930s	139	0.295	0.458			
Born after 1940	139	0.230	0.422			
Panel B. Cases						
	Analysis sample			All observations		
Conservative vote, gender-related cases	3,086	0.605	0.489	5,185	0.591	0.492
Conservative vote, non gender-related cases	54,777	0.569	0.495	9,318	0.546	0.498
Voted to reverse	145,862	0.177	0.382	357,920	0.203	0.402
Female district judge	145,862	0.120	0.325	357,923	0.100	0.300
Author is female	32,052	0.383	0.486	51,183	0.389	0.488
Cites at least one female judge	107,923	0.383	0.486	257,923	0.322	0.467

Notes: This table reports descriptive statistics for judges' characteristics (panel A) and for the main variables considered in the analysis (panel B). In panel B, columns 1 to 3 restrict the sample to observations included in the main analysis, while columns 4 to 6 include all observations.

sample of circuit cases that were hand-coded for vote valence (liberal, conservative, or neutral/hard to code).

Circuit Court Cases.—To study interactions with female judges, we exploit detailed records from all 380,000 circuit court cases for the years 1890–2013, which we obtained from Bloomberg Law (Bloomberg Law 2022). For each case, the records include information on year, circuit, and topic, as well as the panel of judges assigned to decide on it. For each assigned judge, we have information on whether they voted to affirm or reverse the district court decision, whether they authored the majority opinion, and whether they dissented or concurred.

For a subsect of the cases, we are also able to match the case to the identity of the district judge who decided the corresponding lower court case. The judge's name is obtained by parsing the name from the circuit opinion's case history, using data on circuit court cases from LexisNexis (LexisNexis 2022a).[15] We then match the name of the district judge to their biographic information using information from the two sources described above. We are able to assign a unique district judge (and associated gender) to 121,944 circuit cases (32 percent of all cases).

[15] More precisely, the algorithm starts with all district judges, then checks for every case/case history for the corresponding court. Out of the judges within that specific district court, it then narrows it down to the judges that were active during the time of the case plus two years, to allow for appeal proceedings. Finally, based on a name similarity measure (Levenshtein distance), district judges are assigned if the score is above 70 (out of a maximum 100).

Citations.—We use the full text of the opinions to extract information on which cases are cited in the same opinion. We use this dataset to define a citation network that includes both backward and forward citations for each case.

Clerks.—Information on circuit court clerks are from Bonica et al. (2019).

III. Gender Slant and Judge Demographics

This section explores descriptively how the gender slant measure varies based on judge characteristics. We begin by correlating gender slant and judge characteristics using separate univariate regressions with judge-level data. Note that here and in the rest of the paper, the gender slant measure is standardized for ease of interpretation.

Table 3 reports estimates from these regressions. Column 1 shows that judges nominated by presidents of different political parties do not display different levels of gender slant. Column 2 shows that female judges display gender slant that is 0.5 standard deviations lower than male judges, on average. The difference is statistically significant at the 10 percent level. Column 3 shows that there is no difference depending on judge race, possibly an artifact of there being very few minority judges included in the sample. As one would expect, older judges tend to have significantly higher gender slant (column 4): judges that were born before 1920 display between 0.5 and 0.765 standard deviation higher slant than judges born between 1930 and 1939 and after 1940. While this is consistent with older judges holding more socially conservative views, this variation might reflect differences in the cases that were tried by the judges, as older judges served in court in periods with lower female labor force participation. Interestingly, judge cohort appears to have the strongest explanatory power across all different demographic characteristics.

Column 5 includes all characteristics in the same regression and additionally controls for judge religion, law school attended, whether the judge had federal experience prior to being appointed to circuit courts, and circuit fixed effects. The previously discussed correlations remain.[16] Overall, female judges and younger judges display the lowest gender slant.

What explains the variation in gender slant across judges? Gender slant might be partially representational and reflect variation in the facts of the cases tried by the judges. Here, we explore a different possibility: exposure to women. In particular, we ask whether judges that have daughters display different levels of slanted language.[17] Given that, conditional on total number of children, gender should be as good as randomly assigned, we can estimate the causal effect of having daughters on slant (Washington 2008; Glynn and Sen 2015a).

[16] There is no difference in the gender slant of judges born in different regions. Jewish judges appear to be slightly less slanted than Protestant judges. As far as law schools are concerned, judges who received their JDs from Yale display lower slant than judges who received their JDs from Harvard, while Stanford judges have higher slant. Finally, judges from the third (PA, DE, MD), sixth (MI, OH, KT, TN), and Federal Circuits have higher slant than judges from the First Circuit (ME, MA, RI, CT), while judges from the seventh Circuit (WI, IL, IN) display lower gender slant.

[17] The two explanations are not mutually exclusive. Exposure to women, and in particular having daughters, might matter for slant for potentially different reasons, including learning, empathy, or preference realignment (Glynn and Sen 2015a).

Table 3—Correlates of Gender Slant

Dependent variable	Gender slant					
	(1)	(2)	(3)	(4)	(5)	(6)
Democrat	−0.027				−0.003	0.083
	(0.172)				(0.178)	(0.269)
Female		−0.502			−0.592	−0.713
		(0.288)			(0.202)	(0.276)
Minority			−0.098		−0.164	0.453
			(0.329)		(0.194)	(0.283)
Born in 1920s				−0.069	0.080	0.152
				(0.191)	(0.208)	(0.299)
Born in 1930s				−0.765	−0.740	−0.606
				(0.203)	(0.234)	(0.336)
Born after 1940				−0.537	−0.558	−0.381
				(0.229)	(0.258)	(0.338)
Daughter						−0.490
						(0.275)
Observations	139	139	139	139	139	98
Outcome mean	0.000	0.000	0.000	0.000	0.000	−0.085
Adjusted R^2	−0.007	0.020	−0.007	0.087	0.440	0.529
Circuit FE					X	X
Additional controls					X	X
Number of children FE						X

Notes: The table shows the correlation between demographic characteristics and gender slant. We regress gender slant on demographic characteristics of the judge in separate regressions (columns 1 to 4) and in a multivariate regression that includes additional controls and circuit fixed effects (column 5). The additional controls are region of birth, religion, law school attended, and whether the judge had federal experience prior to being appointed to a circuit court. The omitted category for judge cohort are judges born before 1920. In column 6 we additionally include an indicator variable for having at least one daughter, and number of children fixed effects. Standard errors are robust. Gender slant is the standardized cosine similarity between the gender and the career-family dimensions.

To perform the analysis, we combine our measure of slant with information on judges' family composition from Glynn and Sen (2015b). We estimate the following specification:

$$(4) \qquad GenderSlant_j = \beta Daughter_j + \mathbf{X}_j' \gamma + \delta_{c(j)} + \delta_{n(j)} + \epsilon_j,$$

where $GenderSlant_j$ is the slant (standardized cosine similarity between the gender and career-family dimensions) of judge j, $Daughter_j$ is an indicator variable equal to 1 if judge j has at least one daughter, $\mathbf{X}_j$ are demographic characteristics of judge j (gender, party of nominating president, race, religion, region of birth, cohort of birth, law school attended, and whether the judge had federal experience prior to appointment), $\delta_{c(j)}$ are circuit fixed effects, and $\delta_{n(j)}$ are number of children fixed effects. Standard errors are robust.

Table 3 column 6 reports the estimates. We find that conditional on the number of children, having a daughter lowers gender slant by 0.490 standard deviations. In comparison, female judges tend to have about 0.713 standard deviations lower gender slant than male judges in this sample. The effect is only significant at the 10 percent level. While these estimates should be interpreted carefully, it is interesting to note that they are potentially consistent with the view that gender exposure may be important for gender attitudes, in line with the recent literature on the effect of direct contact on attitudes toward specific groups (Alesina et al. 2018; Lowe 2021; Corno, La Ferrara, and Burns 2022).

IV. Slanted Judges and Decisions in Gender-Related Cases

This section asks whether higher-slant judges take different decisions in gender-related cases, as a way of validating our interpretation that gender slant can be seen as a proxy for gender attitudes. We estimate the following specification:

$$(5) \qquad \mathit{Conservative\,Vote}_{ij} \;=\; \beta\,\mathit{Gender\,Slant}_j + \mathbf{X}_j'\gamma + \delta_{c(i)t(i)} + \epsilon_{ij},$$

where $\mathit{Conservative\,Vote}_{ij}$ is an indicator variable equal to 1 if judge j voted conservatively (against expanding women's rights) in case i, $\mathit{Gender\,Slant}_j$ is the gender slant (standardized cosine similarity between the gender and the career-family dimensions) of judge j, $\mathbf{X}_j$ are demographic characteristics of judge j (gender, party of nominating president, race, religion, region of birth, cohort of birth, law school attended, and whether the judge had federal experience prior to appointment), and $\delta_{c(i)t(i)}$ are circuit-year fixed effects. The circuit-year fixed effects ensure that we are using within circuit-year variation, for which cases are quasi-randomly assigned to panels. Instead, controlling for demographic characteristics ensures that the effect of being assigned higher- or lower-slant judges is not confounded by other features of the judge. The dataset is at the vote level, i.e., for each case it includes one observation for the vote of each judge. Standard errors are clustered at the judge level.

Table 4 shows that judges with higher slant are more likely to vote conservatively in gender-related cases. In particular, column 2, which reports estimates from the baseline specification, shows that judges with a one standard deviation higher slant are less likely to vote in favor of expanding women's rights in gender-related cases by 4.1 percentage points, which is significant at the 1 percent level.[18,19] The magnitude of the effect is sizable, corresponding to 7 percent of the outcome mean. To put this in perspective, the effect of being assigned a judge with a one standard deviation higher slant has around one-third of the effect of being assigned a judge that was nominated by a Democratic president. Given that gender slant is measured with error, meanwhile, the estimates are likely to be attenuated toward zero. In addition to providing an important validation for our interpretation of the measure, this result is interesting per se in that it shows that slanted judges have the potential to

[18] Online Appendix Table 6 displays the result of the test proposed in Oster (2019) for bounding selection of unobservables based on selection on observable characteristics. The test assesses the amount of selection on unobservables that is plausible in a setting by looking at the change in the coefficient of interest from an uncontrolled regression to a regression that includes all controls, scaled by the change in R^2. Overall, selection seems to be limited in this setting: selection on unobservable characteristics would need to be twice as large as selection on observables for the treatment effect to equal 0.

[19] Online Appendix Figure 8 shows a binned scatterplot of the main relationship between gender slant and decisions in gender-related cases, conditional on demographic controls and circuit-year fixed effects.

Table 4—Slanted Judges and Decisions in Gender-Related Cases

Dependent variable	Conservative vote				
	(1)	(2)	(3)	(4)	(5)
Gender slant	0.041	0.041	0.041	0.052	0.046
	(0.016)	(0.013)	(0.012)	(0.014)	(0.012)
Democrat		−0.144	−0.141	−0.134	−0.148
		(0.025)	(0.025)	(0.024)	(0.025)
Female		−0.031	−0.042	−0.016	−0.034
		(0.033)	(0.032)	(0.025)	(0.034)
Observations	3,086	3,086	3,086	3,086	3,086
Clusters	113	113	113	113	113
Outcome mean	0.606	0.606	0.606	0.606	0.606
Circuit-year FE	X	X	X	X	X
Additional controls		X	X	X	X
Year of appointment			X		
Exposure FE				X	
No gender-related cases					X

Notes: The table shows the effect of slanted judges on decisions in gender-related cases, i.e., cases related to reproductive rights, gender discrimination, and sexual harassment. We regress an indicator variable equal to 1 if the judge voted conservatively in a gender-related case on the judge's gender slant, demographic controls, and circuit-year fixed effects (equation (2)). Demographic controls are gender, party of appointing President, race (i.e., whether minority), region of birth, cohort of birth, religion, law school attended, and whether the judge had federal experience prior to being appointed to a circuit court. Column 1 does not include demographic controls. Column 2 estimates the baseline specification. Column 3 controls for year of first appointment of the judge to a circuit court. Column 4 includes exposure fixed effects, which are indicator variables equal to 1 if the judge sat on at least one panel in a given circuit over a given 25-year period. In column 5, gender slant is calculated using embeddings trained excluding gender-related cases. Gender slant is the standardized cosine similarity between the gender and the career-family dimensions. The dataset is at the vote level. Standard errors are clustered at the judge level.

influence real-world outcomes (Chen and Yeh 2014; Chen and Sethi 2018; Chen and Yeh 2020b). Gender slant is policy-relevant.

Even if we control for detailed demographic characteristics of the judge, we might still worry that the results are picking up exposure to different types of cases. In fact, while random assignment of judges to cases ensures that judges are exposed to similar cases in a given circuit and year, it is still possible that the cases tried across circuits and years might differ quite significantly. We address this concern in two ways. First, in column 3, we show that controlling for year of first appointment of the judge to an appellate court does not impact the results. Second, in column 4, we include exposure fixed effects, i.e., indicator variables equal to 1 if the judge sat on at least one case in a given circuit over a 25-year period, which ensures that we are only comparing judges who served in similar circuits and years. Again, the result is not affected.[20,21]

[20] Exposure fixed effects are different from circuit-year fixed effects: they are a biographic characteristic that refers to the career of the judge, and not a characteristic of the specific case. We are constrained in how detailed the exposure fixed effects can be by the relatively small number of judges in our sample.

[21] Online Appendix Table 7 shows the results are the same if we separately estimate the regression for the Epstein, Landes, and Posner (2013b) and Glynn and Sen (2015b) datasets.

Robustness Checks.—A potential concern with the gender slant measure is that the corpus of authored opinions might offer a limited reflection of the judges' preferences, if judicial clerks are the ones responsible for drafting them. However, online Appendix Table 8, column 1 shows that controlling for the share of the clerks that are women does not impact the results, suggesting that gender slant is not just proxying for clerk characteristics. In addition, column 2 shows that the result is robust to dropping cases decided in circuits where quasi-random assignment of judges to cases was contested by Chilton and Levy (2015) (namely, the second, eighth, ninth, and DC).

Finally, the result is robust to two procedures that take into account the precision with which the slant of each judge is estimated. First, online Appendix Figure 9 shows that the point estimate is virtually unchanged if we shrink gender slant using Empirical Bayes techniques.[22] Second, online Appendix Table 8, column 3 shows that the point estimate is also not affected if we weight the regression by the inverse of the variance of gender slant across bootstrap samples, thus giving higher weight to judges whose slant is more precisely estimated. However, the fact that the result is robust to controlling for log number of tokens (column 4) shows that the effect we estimate is not driven by differences in corpus size.

Robustness to Word Set Choice.—The result does not depend on the specific choice of words used to construct the gender and career-family dimensions. We experiment with expanding or restricting the word sets, or dropping single words at a time, and present the results in online Appendix Figure 10. In particular, the graph to the left shows the coefficients and 95 percent confidence interval from separate regressions where slant is identified using the top 5 to top 15 most frequent male, female, career, and family words from LIWC. The graph to the right shows coefficients and 95 percent confidence intervals from separate regressions where slant is measured by dropping one attribute word at the time. Smaller word sets give larger confidence intervals and weaker explanatory power of decisions in gender-related cases, but the result is otherwise robust. At the same time, no single word is driving the result: the coefficient is remarkably stable across all the regressions displayed in the graph to the right.

Gender Slant Excluding Gender-Related Cases.—A potential concern with these results is that the gender slant measure could itself be determined by the text of opinions from gender-related cases. Under this argument, cases involving gender-normative situations (women at home and men at work) could be systematically correlated with more conservative decisions, and judges with higher slant might be more exposed to such cases in their careers. We believe this to be unlikely for the following reasons. First, as highlighted before, quasi-random assignment of judges to cases implies that cases are comparable within a given circuit and year. Still, while judges serving across circuits or time might be exposed to different types of cases, our results are robust to including judge cohort fixed effects, year of first appointment

[22] See online Appendix C for a detailed discussion of the procedure we use to implement the EB-adjustment, and why it does not seem to help us lower the tokens threshold.

to a circuit court, and exposure fixed effects, which work to control for the types of cases judges have been exposed to. Second, gender-related cases constitute a small proportion of the texts in the full corpus: only 17,523 cases out of the 114,702 circuit court cases on which we train the judge-specific word embeddings are gender related.[23] Third, the bootstrap procedure we employ to train the word embeddings ensures that the measure is not driven by any specific opinions. Nonetheless, we also show directly that the results are not driven by writing in gender-related cases. As shown in Table 4, column 5, the result is unchanged if the gender slant measure is computed on embeddings trained excluding gender-related cases.[24]

Nongender-Related Cases.—Finally, if gender slant proxies for gender preferences, we should expect it to have larger effects on gender-related cases as opposed to nongender-related cases. Two separate datasets allow us to explore this question. First, we use the Epstein, Landes, and Posner (2013a) dataset, which also includes decisions in ideologically divisive but nongender-related issues such as age discrimination or campaign finance, to study the effect of being assigned a slanted judge on decisions in these types of cases. Online Appendix Table 9 shows that higher-slant judges are more likely to vote conservatively in nongender-related cases, but the effect is around two thirds as big as the effect in gender-related cases. Two additional pieces of evidence are consistent with a gender-focused effect. If we estimate the baseline specification controlling for the share of conservative votes of the judge in nongender-related cases, the effect of being assigned a slanted judge on conservative votes in gender-related cases is unchanged (online Appendix Table 8, column 5). In addition, if we estimate a differences-in-differences specification in which we compare gender- and nongender-related cases that are assigned to judges with different levels of gender slant, we find that judges with higher slant are more likely to vote conservatively in gender-related as opposed to nongender-related cases (online Appendix Table 10).

Second, we use the US Court of Appeals datasets (Songer 2008; Kuersten and Haire 2011), which includes a 5 percent random sample of circuit cases that were hand-coded for vote valence (liberal, conservative, or neutral/hard to code). Online Appendix Table 10 shows that higher-slant judges are not differentially likely to cast a conservative vote across all specifications. Taken together, these results show that while gender slant may be correlated with holding liberal views, the measure does indeed capture attitudes that are specific to gender.

[23] We identify gender-related cases using a simple pattern-based, sequence-classification method. The method identifies a case to be gender-related if it contains a word that is much more likely to appear in gender-related cases as identified by the Epstein et al. (2013) and Glynn and Senn (2015) datasets, then in an equally sized random sample of cases not identified to be gender-related according to these two datasets. First, we match the cases identified to the related opinions string matching on citations and party names. We are able to match 86 percent of the cases. Second, we define a word to be gender-related if it is 25 times more likely to appear in gender-related cases versus non gender-related cases, and if it appears at least 500 times in gender-related cases.

[24] It is still possible, although unintuitive, for gender slant to be a proxy of nongender-related cases for which facts imply a closer association between men and career and women and family. In this case, we should interpret the effect of slanted judges as the effect of being exposed to such cases in such a way that affects not only decision in other gender-related cases but also interaction with female judges.

V. Slanted Judges and Interactions with Female Judges

In this section we ask whether gender attitudes, as proxied by our measure of slant, manifest themselves in differential treatment of female judges on the part of their colleagues. We focus on dimensions that are relevant to a judge's career. In particular, we explore the following outcomes: whether lower court decisions by female judges are more likely to be reversed with respect to lower court decisions by male judges, whether female judges are assigned the writing of majority opinions, and forward citations by future judges.

A. *Slanted Judges and Disparities in Reversals*

District court trials are presided by a single judge, and cases are assigned to district judges quasi-randomly within each district-year. Up to 40 percent of district cases are appealed and are therefore considered by circuit courts (Eisenberg 2004). Importantly, reversals matter for career outcomes: as shown in online Appendix Figure 11, district judges that have a higher share of their decisions reversed on appeal are less likely to be promoted to circuit courts.

Here, we ask whether judges with higher gender slant are differentially likely to reverse decisions authored by female relative to male district court judges. The empirical strategy we employ in this section is slightly different than the rest of the paper, although it also builds on quasi-random assignment of circuit judges to cases and on conditioning on observable characteristics. In particular, we identify the effect of being assigned a higher-slant judge on reversals using a differences-in-differences design that compares appealed cases decided by female and male district judges that are assigned to circuit judges with different levels of slant. Identification requires that cases originally decided by a male district judge and assigned to circuit judges with different level of slant provide a good control group for cases that were originally decided by a female district judge. The quasi-random assignment of cases to panels at the circuit level helps us in this respect: cases assigned to judges with higher or lower slant are comparable. Importantly, the identification strategy allows for cases decided by female and male judges to be different, for example, because they are appealed at different rates, as long as there is no systematic assignment of cases to higher and lower slant judges based on the likely reversal outcome.

We estimate the following baseline specification:

$$\begin{aligned} \mathit{Voted\,to\,Reverse}_{ij} \;=\; & \pi\,\mathit{Female\,District\,Judge}_{k(i)} \times \mathit{Gender\,Slant}_j \\ & + \mathit{Female\,District\,Judge}_{k(i)} \times \mathbf{X}_j'\gamma + \delta_{c(i)t(i)} + \delta_{k(i)} + \delta_j + \epsilon_{ij}, \end{aligned}$$

where $\mathit{Voted\,to\,Reverse}_{ij}$ is an indicator variable equal to 1 if circuit judge j voted to reverse the district court decision in case i, $\mathit{Gender\,Slant}_j$ is the slant (standardized cosine similarity between the gender and career-family dimensions) of circuit judge j, $\mathit{Female\,District\,Judge}_{k(i)}$ is a dummy equal to 1 if the district judge who originally decided the case (judge $k(i)$) is female, $\mathbf{X}_j$ are demographic controls for the circuit court judge (gender, party of nominating president, race, religion, region of birth, cohort of birth, law school attended, and whether the judge had federal experience prior to appointment), $\delta_{c(i)t(i)}$ are circuit-year fixed effects, $\delta_{k(i)}$ are district judge fixed effects, and δ_j are circuit judge fixed effects. The dataset is at the vote level. Standard errors are clustered at the circuit judge level.

The inclusion of the circuit-year fixed effects $\delta_{c(i)t(i)}$ ensures comparability of cases assigned to higher and lower slant judges because of quasi-random assignment of cases to judges. As before, controlling for the demographic characteristics $\mathbf{X}_j$ shows that the effect of higher-slant judges is not confounded by other judge characteristics. District judge fixed effects $\delta_{k(i)}$ ensure that we are comparing what happens when cases decided at the district court level by the same judge are assigned to circuit judges with higher and lower slant. Finally, the circuit judge fixed effects δ_j allow for circuit judges to differ in their baseline probability of reversing a district court decision.

Table 5, column 1 reports estimates from a specification that only includes the fixed effects, while column 2 estimates the baseline specification. Circuit judges with a one standard deviation higher slant are 1 percentage point (5.6 percent of the baseline mean) more likely to vote to reverse a district court decision if the district judge is female, relative to when the district judge is male. Other characteristics of the circuit judge do not make a difference: being appointed by a Democratic president or being female is unrelated to disparately voting to reverse female district judges. Female district judges are 2.3 percentage points less likely to be reversed, and going from a judge with the mean slant to a judge with a one standard deviation higher slant decreases this gap almost by half.[25]

The result is not explained by judges being exposed to different cases over the course of their careers: the coefficient on slant is unchanged if we control for year of first appointment to a circuit court (column 3), or if we include exposure fixed effects that ensure we are comparing judges who served in similar periods and areas (column 4). Finally, using a gender slant measure calculated based on embeddings trained excluding gender-related cases barely affects the estimate (column 5).

It is also interesting to explore whether the effect of assignment to a higher-slant judge changes over time. In particular, we might think that those who, in more recent years, still express stereotyped views of gender in their writings might display especially discriminatory behavior when dealing with female colleagues. Consistently with this hypothesis, in online Appendix Figure 13 we show that the effect we estimate seems to be almost entirely driven by cases from after 2000.

Are higher-slant judges reacting specifically to the gender of the district judge? Not necessarily: it could be that higher-slant judges are reacting to a characteristic

[25] Online Appendix Figure 12 shows a binned scatterplot of the main relationship between gender slant and reversals separately by gender of the district judge, conditional on demographic controls and circuit-year fixed effects. The figure shows that the reversal gap between male and female district judges is almost entirely driven by lower slanted circuit judges. Instead, the gap almost disappears for higher slanted circuit judges, who reverse male and female district judges at similar rates. This pattern is consistent with two interpretations. If the reversal gap between male and female district judges for low-slant circuit judges reflects what should be the "true" reversal rate based on decision quality, then higher-slant judges over-reverse female judges. Alternatively, it is possible that male and female district judges should be reversed at the same rate, and that lower-slant judges reverse decisions of male district judges more. Unfortunately, without a measure of decision quality that is independent of reversals, we cannot tease apart these two interpretations.

Table 5—Slanted Judges and Reversals

Dependent variable	Voted to reverse district decision				
	(1)	(2)	(3)	(4)	(5)
Gender slant × female district judge	0.006	0.010	0.010	0.010	0.009
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
Democrat × female district judge		−0.010	−0.010	−0.010	−0.011
		(0.006)	(0.006)	(0.006)	(0.006)
Female × female district judge		−0.003	−0.002	−0.003	−0.005
		(0.010)	(0.010)	(0.010)	(0.010)
Observations	145,862	145,862	145,862	145,862	145,862
Clusters	133	133	133	133	133
Outcome mean, male district judge	0.180	0.180	0.180	0.180	0.180
Outcome mean, female district judge	0.157	0.157	0.157	0.157	0.157
Circuit-year FE	X	X	X	X	X
Circuit judge FE	X	X	X	X	X
District judge FE	X	X	X	X	X
Additional controls		X	X	X	X
Year of appointment			X		
Exposure FE				X	
No gender-related cases					X

Notes: The table shows the differential effect of slanted judges on the reversal probability of cases originally decided by male and female district judges using a differences-in-differences design. We regress an indicator variable equal to 1 if the judge voted to reverse the district decision on the gender slant of the judge interacted with an indicator variable for whether the district judge is female, demographic controls interacted with an indicator variable for whether the district judge is female, circuit judge fixed effects, district judge fixed effects, and circuit-year fixed effects (equation (3)). Demographic controls are gender, party of appointing President, race (i.e., whether minority), region of birth, cohort of birth, religion, law school attended, and whether the judge had federal experience prior to being appointed to a circuit court and refer to the circuit judge. Column 1 does not include demographic controls. Column 2 estimates the baseline specification. Column 3 controls for year of first appointment of the judge to a circuit court interacted with an indicator variable for whether the district judge is female. Column 4 includes exposure fixed effects, which are indicator variables equal to 1 if the judge sat on at least one panel in a given circuit over a given 25-year period, interacted with an indicator variable for whether the district judge is female. In column 5, gender slant is calculated using embeddings trained excluding gender-related cases. Gender slant is the standardized cosine similarity between the gender and the career-family dimensions. The dataset is at the vote level. The sample is restricted to cases for which we were able to determine the identity of the district judge. Standard errors are clustered at the circuit judge level.

of appealed cases that correlates with the gender of the district judge. For example, female district judges might be systematically more lenient in criminal cases, and higher-slant judges might be responding along that margin. Because we don't observe such granular features of the district judge's decisions, we cannot fully rule it out. However, we note that with the differences-in-differences design, the pattern is not driven by level differences across cases. For example, if female district judges systematically make higher (or lower) quality decisions, or are appealed in different types of cases, that should be addressed by the fixed effects. These characteristics matter only to the extent that they interact with gender slant. Thus, the analysis is still informative as to how gender preferences affect gendered incidence of reversals, even if it is driven by female judges making decisions that higher-slant judges tend to find fault with.

Robustness Checks.— Online Appendix Table 12 shows that the main effect on reversals is robust to a number of additional robustness checks. Including district-year fixed effects, which are not needed for identification but might increase precision, does not impact the result (column 1). Controlling for the share of female clerks (column 2) and controlling for the circuit judge's vote record in ideologically divisive cases (column 3) do not make a difference. In addition, the main effect is stable to restricting the sample to the post-1980 period when there were likely to be more female district judges (column 4) or to circuits in which we are confident cases were quasi-randomly assigned to judges (column 5). The effect on reversals is also not explained by some slant measures being estimated more precisely than others. The point estimate is, if anything, slightly larger when we use an EB-adjusted version of slant (online Appendix Figure 14) or when we weight the regression by the inverse of the variance of slant across bootstrap samples (column 6). Controlling for the size of the corpus also does not make a difference (column 7). Finally, online Appendix Figure 15 shows that the results are robust to using different word sets to identify the gender and career-family dimensions.

Back-of-the-Envelope Calculation.—Given that reversals have a negative effect on district judges' promotion, higher-slant judges have the potential to hinder the career progression of female district judges with respect to male district judges. Using estimates of the effect of reversals on the probability that district judges get elevated to circuit courts, we can try to estimate the magnitude of this effect in a back-of-the-envelope calculation. In particular, we find that a female judge whose appealed decisions were assigned to circuit judges with on average a one standard deviation higher slant would be 6.3 percent less likely to be elevated than a male judge faced with similarly slanted circuit judges.[26]

B. *Slanted Judges and Disparities in Opinion Authorship*

In circuit courts, decisions are generally taken in conference by the three judges on the panel after oral arguments. The decision with the most votes becomes the majority position, and the judges in the majority then have to decide who is going to author the associated opinion. By custom, the most senior acting judge in the majority assigns the responsibility of writing, taking into consideration expertise, work load, and other factors (Bowie, Songer, and Szmer 2014). The majority opinion articulates the principles behind the decision, which are binding law for lower courts (Rohde and Spaeth 1976): the authoring of the opinion itself is an important task. Given the policy stakes of opinion assignment, a relevant question is whether the preferences of the senior judges affect this procedure. In particular, we investigate whether senior judges with higher gender slant are differentially likely to assign majority opinions to female judges.

We estimate the following specification:

$$(7) \qquad FemaleAuthor_i \ = \ \beta\, GenderSlant_{j(i)}^{SENIOR} + \mathbf{X}_{j(i)}^{SENIOR'}\gamma + \delta_{c(i)t(i)} + \epsilon_i,$$

[26] In fact, online Appendix Table 13 shows that an increase from 0 to 1 in the share of votes to reverse the district court decision on appeal implies a 38 percentage point decrease in the probability of being elevated. The calculation follows from the fact that female district judges have around a 6.8 percent baseline probability of being elevated. Interestingly, the relationship between reversals and promotions is not differential by gender of the district judge.

where $FemaleAuthor_i$ is a dummy equal to 1 if the author of the majority opinion of case i, $GenderSlant_{j(i)}^{SENIOR}$ is the gender slant (standardized cosine similarity between the gender and the career-family dimensions) of the most senior judge on the panel of case i, $\mathbf{X}_{j(i)}^{SENIOR}$ are demographic characteristics of the most senior judge (gender, party of nominating president, race, religion, region of birth, cohort of birth, law school attended, and whether the judge had federal experience prior to appointment), and $\delta_{c(i)t(i)}$ are circuit-year fixed effects. The dataset is at the case level. Standard errors are clustered at the senior judge level.[27] For opinion assignment to a female or a male judge to be a meaningful decision, we drop per curiam (unsigned) opinions and restrict the sample to cases that have at least one female judge on the panel. Since the decision to dissent or concur is possibly endogeneous, we also exclude cases that contain either a dissent or a concurrence.[28]

Table 6 reports the estimated effect of the most senior judge having higher or lower gender slant on whether the authoring judge is female. Column 1 reports the estimates from a specification that only controls for the circuit-year fixed effects, while column 2 estimates the baseline specification. When the most senior judge on the panel has a one standard deviation higher slant, the majority opinion is less likely to be authored by one of the female judges by 1.7 percentage points, corresponding to a 4.5 percent decrease over the outcome mean.[29,30] If the most senior judge is a woman instead, there is a 13.4 percentage points (35 percent) higher probability that the author of the opinion is a woman. While slant has a nontrivial effect on the gender of the authoring judge, the magnitude of the effect is substantially smaller than that of being female.[31]

Columns 3 and 4 show that the result is unchanged when we control for exposure to different types of cases. In addition, if we use a gender slant measure calculated based on embeddings trained excluding gender-related cases, the result is the same (column 5). As was the case for reversals, online Appendix Figure 17 shows that the effect of the most senior judge on the panel having higher slant is larger for more recent cases (i.e., those after 2000). Overall, these results show that judges with more conservative attitudes toward gender are less likely to assign an important career-relevant task to female judges.

[27] We determine the most senior judge on the panel using information on the career of appellate judges, and exclude from the analysis cases for which we could not precisely determine who the identity of the most senior judge on the case, mainly because of there being multiple judges appointed in the same year.

[28] In online Appendix Table 14 we check whether the gender slant of the panel's senior judge affects the probability of having a specific author (columns 1 and 2) or having a per curiam opinion (columns 3 and 4), and we find no effect. We also show that the slant of the most senior judge on the panel does not impact the probability of unanimous decisions, that is, having dissents or concurrences (columns 5 and 6).

[29] Female judges are, on average, assigned opinions at the same rate as male judges. The baseline probability shown here is higher than 0.33 because we include panels with one, two or three female judges.

[30] Online Appendix Figure 16 shows a binned scatterplot of the main relationship between gender slant and opinion assignment, conditional on demographic controls, and circuit-year fixed effects.

[31] To further put this number in perspective, it is helpful to compare the dynamics we see here with what would have been the opinion assignment process had it been random. If we exclude panels where the most senior judge is female to abstract from possible self-assignment issues (see column 5), we find that with random assignment the share of opinions authored by a female judge would have been 0.359. In the data we find that the observed proportion is lower at 0.347. The gap is larger for assigning judges with slant above the median (0.344) than for those with slant below the median (0.349). In other words, having a senior judge in the bottom half rather than in the top half of the slant distribution closes 40 percent of the gap.

Table 6—Slanted Judges and Opinion Assignment

Dependent variable	Author is female				
	(1)	(2)	(3)	(4)	(5)
Gender slant	−0.028	−0.017	−0.017	−0.018	−0.016
	(0.011)	(0.008)	(0.008)	(0.010)	(0.008)
Democrat		−0.001	−0.001	−0.015	0.002
		(0.014)	(0.014)	(0.017)	(0.014)
Female		0.134	0.133	0.158	0.137
		(0.016)	(0.016)	(0.017)	(0.016)
Observations	32,052	32,052	32,052	32,052	32,052
Clusters	125	125	125	125	125
Outcome mean	0.383	0.383	0.383	0.383	0.383
Circuit-year FE	X	X	X	X	X
Additional controls		X	X	X	X
Year of appointment			X		
Exposure FE				X	
No gender-related cases					X

Notes: The table shows the effect of slanted judges on the probability of the majority opinion being assigned to a female judge. We regress an indicator variable equal to 1 if the authoring judge is female on the gender slant of the most senior judge on the panel, demographic controls, and circuit-year fixed effects (equation (4)). The most senior judge on the panel is customarily in charge of assigning the majority opinion. Demographic controls are gender, party of appointing president, race (i.e., whether minority), region of birth, cohort of birth, religion, law school attended, and whether the judge had federal experience prior to being appointed to a circuit court and they refer to the most senior judge on the panel. Column 1 does not include demographic controls. Column 2 estimates the baseline specification. Column 3 controls for year of first appointment of the most senior judge of the panel to a circuit court. Column 4 includes exposure fixed effects, which are indicator variables equal to 1 if the most senior judge on the panel sat on at least one panel in a given circuit over a given 25-year period. In column 5, gender slant is calculated using embeddings trained excluding gender-related cases. Gender slant is the standardized cosine similarity between the gender and career-family dimensions. The dataset is at the case level. The sample is restricted to cases with a specific author, with at least one female judge on the panel, and that were decided unanimously. Standard errors are clustered at the senior judge level.

Robustness Checks.— Online Appendix Table 15 shows that the result is robust to a number of additional robustness checks. As before, in column 1 we control for the share of clerks that are female and find that it limitedly affects the main result: while the coefficient is no longer statistically significant—possibly because we have information on clerks only for a subset of the judges—it is very similar in magnitude. Reassuringly, column 2 shows that the result is not driven by confounded ideology of judges: if we control for a measure of how conservative a judge's voting record is, the coefficient on slant is unchanged.

Given the large effect of the senior judge being female on the probability that the authoring judge is a woman and that female judges have lower gender slant, we might worry that self-assignments are the driver behind the main effect. Column 3 shows that this is not the case: if we re-estimate the main specification excluding cases where the most senior judge on the panel is a woman, we find the same effect. Column 4 shows instead that the main result is, if anything, larger if we include cases that had dissents or concurrences. Columns 5 and 6 shows that the results are robust to restricting the sample to the post-1980 period and to dropping circuits for which Chilton and Levy (2015) doubted quasi-random assignment.

The result is also robust to controlling for differences in how precisely slant is estimated for different judges, for example by using an EB-adjusted measure of slant (online Appendix Figure 18) or giving more weight to judges whose gender slant is more precisely estimated (column 7). Controlling for the size of the corpus (column 8) does not make a difference. Finally, online Appendix Figure ?? shows that the results are robust to using different word sets to identify the gender and career-family dimensions.

Assignment of Different Types of Cases.—This result raises the question of whether slanted judges also assign different types of cases to female judges. In online Appendix Figure 20 and Appendix Table 16, we look but do not find evidence of any differences. The cases assigned to female judges by higher-slant senior judges are not concentrated in specific areas of the law and do not have different expected importance, as proxied by forward citations predicted based on case characteristics, which is potentially an interesting result in light of the literature on discrimination in task assignments.

C. *Slanted Judges and Disparities in Citations*

The last career-related interaction we examine are citations. Law depends on precedent, and deciding which cases to cite in a specific opinion is a nontrivial decision: "Judges [and] lawyers, who brief and argue cases, (…) could all be thought, with only slight exaggeration, to make their living in part by careful citation of judicial decisions" (Posner 2000, 382). Meanwhile, many judges admit to monitoring and caring about whether they are cited by other judges (Posner 2008), and citations to cases are commonly understood as a measure of judge quality (Ash and MacLeod 2015). Citations are therefore relevant to judicial careers, and differential treatment of male and female judges in citation choices presents another potential domain for high-stakes discrimination.

Identification once again relies on quasi-random assignment of cases to judge panels. Conditional on circuit-year fixed effects, cases assigned to different panels are comparable. However, choice of authorship is endogenous. Even if the fact that we control for a number of judge characteristics improves comparability across judges, it is possible that judges with higher gender slant are systematically assigned authorship of cases for which it would be optimal to differentially cite female judges in the first place: the results in this section have to be interpreted especially carefully.

The specification we estimate is:

$$(8) \qquad CitesFemaleJudge_i = \beta GenderSlant_{j(i)} + \mathbf{X}'_{ji}\gamma + \delta_{c(i)t(i)} + \epsilon_i,$$

where $CitesFemaleJudge_{ictj}$ is an indicator variable equal to 1 if the opinion of case i cites at least one opinion authored by a female judge, $GenderSlant_{j(i)}$ is the slant (standardized cosine similarity between the gender and career-family dimensions) of judge j who authored the majority of case i, $\mathbf{X}_{j(i)}$ are demographic characteristics of judge j (gender, party of nominating president, race, religion, region of birth, cohort of birth, law school attended, and whether the judge had federal experience prior to appointment), and $\delta_{c(i)t(i)}$ are circuit-year fixed effects. The dataset includes one observation for each case, and is restricted to cases in which the opinion was authored by a specific judge. Standard errors are clustered at the judge level.

Table 7 reports the estimates from equation (8). Column 1 only controls for the circuit-year fixed effects, while column 2 estimates the baseline specification. Judges with a one standard deviation higher gender slant are 1 percentage point less likely to cite any opinions authored by female judges. The effect is significant at the 10 percent level, and relatively small in magnitude (2.6 percent of the outcome mean), especially if compared to the effect of the author being a woman (12.8 percentage points or 34 percent). Controlling for year of appointment (column 3) does not affect the result, and neither does including career fixed effects (column 4). However, the result is not robust to using gender slant measured on embeddings trained excluding gender-related cases (column 5). As before, it is interesting to note that the effect of higher-slant judges is only present for opinions authored after 2000 (online Appendix Figure 22).[32]

Robustness Checks.— Online Appendix Table 17 elaborates on the robustness of the main effect on citations with different specifications and sample restrictions. The effect on citations is our least robust result. The main effect is not robust to controlling for the share of female clerks (column 1), perhaps because clerks are important when determining citations or simply because of the smaller sample size. Controlling for the circuit judge's vote record in ideologically divisive cases (column 2) limitedly impacts the main coefficient, and the effect is stable to restricting the sample to the post-1980 period (column 3). However, the result is not robust to excluding the second, eighth, ninth, and DC circuits (column 4).

As with opinion assignment, the large effect of being a female judge on citations suggests that the result might be explained by self-citations. Even when we define the outcome excluding self-cites, however, online Appendix Table 17, column 5 shows that gender slant has a negative and statistically significant at the 10 percent level on the probability of citing a female judge. Interestingly, however, when self-cites are excluded, female judges appear to be less likely to cite other female judges.

In addition, the effect is not an artifact of gender slant being more precisely estimated for some judges. Online Appendix Figure 23 shows that the effect is robust to using an EB-adjusted version of slant and Appendix Table 17, column 6 shows that the result is not impacted by weighting by the inverse of the variance of the slant measure across bootstrap sample. However, as shown in column 7, the result is not robust to controlling for log number of tokens. Finally, the effect on citations is robust to using different word sets (online Appendix Figure 24).

Overall, judges with higher gender slant appear to be less likely to cite opinions authored by female judges, although the effect is less robust than other findings presented in the paper and should be interpreted with special caution because of potential endogeneity.

[32] Online Appendix Figure 21 shows a binned scatterplot of the relationship between gender slant and citations, conditional on demographic controls and circuit-year fixed effects.

Table 7—Slanted Judges and Citations

Dependent variable	Cites at least one female judge				
	(1)	(2)	(3)	(4)	(5)
Gender slant	−0.024	−0.010	−0.009	−0.014	−0.005
	(0.007)	(0.005)	(0.005)	(0.007)	(0.005)
Democrat		−0.012	−0.011	−0.020	−0.011
		(0.011)	(0.011)	(0.010)	(0.011)
Female		0.128	0.125	0.142	0.131
		(0.016)	(0.016)	(0.015)	(0.016)
Observations	107,923	107,923	107,923	107,923	107,923
Clusters	139	139	139	139	139
Outcome mean	0.383	0.383	0.383	0.383	0.383
Circuit-year FE	X	X	X	X	X
Additional controls		X	X	X	X
Year of appointment			X		
Exposure FE				X	
No gender-related cases					X

Notes: The table shows the effect of slanted judges on the probability of citing at least one female judge. We regress a dummy equal to 1 if the majority opinion cites at least one case authored by a woman on the gender slant of the author of the majority opinion, demographic controls, and circuit-year fixed effects (equation (5)). Demographic controls are gender, party of appointing President, race (i.e., whether minority), region of birth, cohort of birth, religion, law school attended, and whether the judge had federal experience prior to being appointed to a circuit court. Column 1 does not include demographic controls. Column 2 estimates the baseline specification. Column 3 controls for year of first appointment to a circuit court. Column 4 includes exposure fixed effects, which are indicator variables equal to 1 if the judge sat on at least one panel in a given circuit over a given 25-year period. In column 5, gender slant is calculated using embeddings trained excluding gender-related cases. Gender slant is the standardized cosine similarity between the gender and career-family dimensions. The dataset is at the case level. The sample is restricted to cases in which the opinion was authored by a specific judge. Standard errors are clustered at the judge level.

D. *Other Judge Characteristics besides Gender*

To strengthen the argument that gender slant is indeed proxying for gender attitudes, we explore whether higher-slant judges also interact differently with judges with specific demographic characteristics: political leanings, minority status, and age.

Reversals.—First, online Appendix Table 18 shows that higher-slant judges are not differentially likely to reverse cases of district judges that were appointed by a Democratic president than cases of district judges that were appointed by a Republican president (column 1). Circuit judges with higher slant, however, appear more likely to reverse cases in which the district judge is a minority (column 2).

Opinion Assignment.—Second, online Appendix Table 19 shows that judges with higher slant are not differentially likely to assign the opinion of the court to Democratic judges (column 1), minority judges (column 2), or based on the age of the judge (column 3).

Citations.—Finally, online Appendix Table 20 shows that judges with higher gender slant are less likely to cite opinions authored by Democrat-appointed judges (column 1). Instead, judges with higher gender slant do not differentially cite minority judges (column 2) or judges of different ages (column 3). Interestingly, judges with higher slant are more likely to cite opinions authored by judges with higher slant as well (column 4), and this pattern is not due to judge gender or party.

Overall, these results suggest that gender is the salient characteristic to which judges with higher slant respond: gender slant specifically proxies for attitudes toward women.

VI. Conclusions

This paper investigates the role of gender attitudes in circuit courts. We find that gender attitudes, at least as far as they are expressed in judicial writing, matter. Judges with higher gender slant vote more conservatively on women's rights cases, are more likely to reverse district courts decisions when the district judge is a woman, are less likely to assign opinions to women, and cite fewer opinions authored by female judges.

These findings add to the literature on gender attitudes by showing that they matter even for skilled professionals making high-stakes, public-oriented decisions. Our text-based metric is a proxy for a psychological factor, and so the policy implications of the results should be considered with caution. Although we estimate the causal effect of assigning higher-slanted judges, we do not have evidence, for example, that forcing judges to use less stereotypical language would causally shift their decisions in gender law or their behavior toward female colleagues.

This research can be extended in a number of directions. First, it would be important to know how well text-based measures of gender attitudes correlate with other measures, such as scores on the implicit association test. Second, the text-based metrics could be computed for other decision-makers such as politicians, journalists, and professors. In these domains, as with judges, there are no traditional measures of attitudes, but large corpora of text are available.

REFERENCES

Abrams, David S., Marianne Bertrand, and Sendhil Mullainathan. 2012. "Do Judges Vary in Their Treatment of Race?" *Journal of Legal Studies* 41 (2): 347–83.

Alesina, Alberto, Michela Carlana, Eliana La Ferrara, and Paolo Pinotti. 2018. "Revealing Stereotypes: Evidence from Immigrants in Schools." Unpublished.

Antoniak, Maria, and David Mimno. 2018. "Evaluating the Stability of Embedding-Based Word Similarities." *Transactions of the Association for Computational Linguistics* 6: 107–19.

Arnold, David, Will Dobbie, and Crystal S. Yang. 2018. "Racial Bias in Bail Decisions." *Quarterly Journal of Economics* 133 (4): 1885–1932.

Ash, Elliott, and W. Bentley MacLeod. 2015. "Intrinsic Motivation in Public Service: Theory and Evidence from State Supreme Courts." *Journal of Law and Economics* 58 (4): 863–913.

Ash, Elliott, Daniel L. Chen, and Sergio Galletta. 2022. "Measuring Judicial Sentiment: Methods and Application to US Circuit Courts." *Economica* 89 (354): 362–76.

Ash, Elliott, Daniel L. Chen, and Arianna Ornaghi. 2024. "Replication data for: Gender Attitudes in the Judiciary: Evidence from US Circuit Courts." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E169861V1.

Battaglini, Marco, Jorgen M. Harris, and Eleonora Patacchini. 2023. "Interactions with Powerful Female Colleagues Promote Diversity in Hiring." *Journal of Labor Economics* 41 (3): 589–614.

Bertrand, Marianne. 2013. "Career, Family, and the Well-Being of College-Educated Women." *American Economic Review* 103 (3): 244–50.

Bertrand, Marianne, Dolly Chugh, and Sendhil Mullainathan. 2005. "Implicit Discrimination." *American Economic Review* 95 (2): 94–98.

Bertrand, Marianne, Claudia Goldin, and Lawrence F. Katz. 2010. "Dynamics of the Gender Gap for Young Professionals in the Financial and Corporate Sectors." *American Economic Journal: Applied Economics* 2 (3): 228–55.

Bloomberg Law. 2022. "Court Opinions: Dataset." Bloomberg Law.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. 2019. "The Dynamics of Discrimination: Theory and Evidence." *American Economic Review* 109 (10): 3395–3436.

Bolukbasi, Tolga, Kai-Wei Chang, James Y. Zou, Venkatesh Saligrama, and Adam T. Kalai. 2016. "Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings." *Advances in Neural Information Processing Systems* 29: 4349–57.

Bonica, Adam, Adam Chilton, Jacob Goldin, Kyle Rozema, and Maya Sen. 2019. "Legal Rasputins? Law Clerk Influence on Voting at the US Supreme Court: Dataset." *Journal of Law, Economics, and Organization*.

Bordalo, Pedro, Katherine Coffman, Nicola Gennaioli, and Andrei Shleifer. 2019. "Beliefs about Gender." *American Economic Review* 109 (3): 739–73.

Bowie, Jennifer Barnes, Donald R. Songer, and John Szmer. 2014. *The View from the Bench and Chambers: Examining Judicial Process and Decision Making on the US Courts of Appeals*. Charlottesville, VA: University of Virginia Press.

Boyd, Christina L., and James F. Spriggs II. 2009. "An Examination of Strategic Anticipation of Appellate Court Preferences by Federal District Court Judges." *Journal of Law and Policy* 29: 37–81.

Bursztyn, Leonardo, Thomas Fujiwara, and Amanda Pallais. 2017. "'Acting Wife': Marriage Market Incentives and Labor Market Investments." *American Economic Review* 107 (11): 3288–3319.

Caliskan, Aylin, Joanna J. Bryson, and Arvind Narayanan. 2017. "Semantics Derived Automatically from Language Corpora Contain Human-Like Biases." *Science* 356 (6334): 183–86.

Card, David, Stefano DellaVigna, Patricia Funk, and Nagore Iriberri. 2020. "Are Referees and Editors in Economics Gender Neutral?" *Quarterly Journal of Economics* 135 (1): 269–327.

Carlana, Michela. 2019. "Implicit Stereotypes: Evidence from Teachers' Gender Bias." *Quarterly Journal of Economics* 134 (3): 1163–1224.

Chen, Daniel L., and Jasmin Sethi. 2018. "Insiders, Outsiders, and Involuntary Unemployment: Sexual Harassment Exacerbates Gender Inequality." Unpublished.

Chen, Daniel L., and Susan Yeh. 2014. "The Construction of Morals." *Journal of Economic Behavior and Organization* 104: 84–105.

Chen, Daniel L., and Susan Yeh. 2020a. "Growth under the Shadow of Expropriation? Economic Impacts of Eminent Domain." Unpublished.

Chen, Daniel L., and Susan Yeh. 2020b. "How Do Rights Revolutions Occur? Free Speech and the First Amendment." Unpublished.

Chen, Daniel L., Vardges Levonyan, and Susan Yeh. 2020. "Can Policies Affect Preferences? Evidence from Random Variation in Abortion Jurisprudence." Unpublished.

Chilton, Adam S., and Marin K. Levy. 2015. "Challenging the Randomness of Panel Assignment in the Federal Courts of Appeals." *Cornell Law Review* 101 (1): 1–56.

Cohen, Alma, and Crystal S. Yang. 2019. "Judicial Politics and Sentencing Decisions." *American Economic Journal: Economic Policy* 11 (1): 160–91.

Corno, Lucia, Eliana La Ferrara, and Justine Burns. 2022. "Interaction, Stereotypes, and Performance: Evidence from South Africa." *American Economic Review* 112 (12): 3848–75.

Croft, Jane. 2016. "Law School Admissions Collapse Continues: They are Being Forced to Innovate or Face Being Left Behind." *Financial Times,* November 20. https://www.ft.com/content/4ddb437e-9ace-11e6-8f9b-70e3cabccfae.

DiMaggio, Paul. 1997. "Culture and Cognition." *Annual Review of Sociology* 23 (1): 263–87.

Eisenberg, Theodore. 2004. "Appeal Rates and Outcomes in Tried and Nontried Cases: Further Exploration of Anti-plaintiff Appellate Outcomes." *Journal of Empirical Legal Studies* 1 (3): 659–88.

Epstein, Lee, William M. Landes, and Richard A. Posner. 2013a. *The Behavior of Federal Judges*. Cambridge, MA: Harvard University Press.

Epstein, Lee, William M. Landes, and Richard A. Posner. 2013b. "The Behavior of Federal Judges: Dataset." USC Gould.

Federal Judicial Center. 2022. "Biographical Directory of Article III Federal Judges, 1789–Present." Federal Judicial Center.

Friese, Malte, Matthias Bluemke, and Michaela Wänke. 2007. "Predicting Voting Behavior with Implicit Attitude Measures: The 2002 German Parliamentary Election." *Experimental Psychology* 54 (4): 247–55.

Garg, Nikhil, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. "Word Embeddings Quantify 100 Years of Gender and Ethnic Stereotypes." *PNAS* 115 (16): E3635–E3644.

George, Tracey E., and Albert H. Yoon. 2019. *The Gavel Gap: Who Sits in Judgement on State Courts?* Washington, DC: American Constitution Society for Law and Policy.

Glaser, Barney G., and Anselm L. Strauss. 2017. *Discovery of Grounded Theory: Strategies for Qualitative Research*. London, UK: Routledge.

Glover, Dylan, Amanda Pallais, and William Pariente. 2017. "Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores." *Quarterly Journal of Economics* 132 (3): 1219–60.

Glynn, Adam N., and Maya Sen. 2015a. "Identifying Judicial Empathy: Does Having Daughters Cause Judges to Rule for Women's Issues?" *American Journal of Political Science* 59 (1): 37–54.

Glynn, Adam N., and Maya Sen. 2015b. "Identifying Judicial Empathy: Does Having Daughters Cause Judges to Rule for Women's Issues?: Dataset." *American Journal of Political Science*. https://doi.org/10.7910/DVN/26544.

Green, Alexander R., Dana R. Carney, Daniel J. Pallin, Long H. Ngo, Kristal L. Raymond, Lisa I. Iezzoni, and Mahzarin R. Banaji. 2007. "Implicit Bias among Physicians and its Prediction of Thrombolysis Decisions for Black and White Patients." *Journal of General Internal Medicine* 22 (9): 1231–38.

Gryski, Gerald S., and Gary Zuk. 2008a. "A Multi-user Database on the Attributes of US Appeals Court Judges, 1801–2000." Songer Project. http://www.songerproject.org/attributes.html.

Gryski, Gerald S., and Gary Zuk. 2008b. "A Multi-user Database on the Attributes of US District Court Judges, 1789–2000." Songer Project. http://www.songerproject.org/attributes.html.

Hengel, Erin. 2022. "Publishing While Female: Are Women Held to Higher Standards? Evidence from Peer Review." *Economic Journal* 132 (648): 2951–91.

Jakiela, Pamela, and Owen Ozier. 2021. "Gendered Language." Unpublished.

Kastellec, Jonathan P. 2013. "Racial Diversity and Judicial Influence on Appellate Courts." *American Journal of Political Science* 57 (1): 167–83.

Kozlowski, Austin C., Matt Taddy, and James A. Evans. 2019. "The Geometry of Culture: Analyzing Meaning through Word Embeddings." *American Sociological Review* 84 (5): 905–49.

Kuersten, Ashlyn K., and Susan B. Haire. 2011. "Update to the Appeals Courts Database 1997–2002." Songer Project. http://www.songerproject.org/data.html.

Le, Quoc, and Tomas Mikolov. 2014. "Distributed Representations of Sentences and Documents." *Proceedings of Machine Learning Research* 32 (2): 1188–96.

Levinson, Justin D., Mark W. Bennett, and Koichi Hioki. 2017. "Judging Implicit Bias: A National Empirical Study of Judicial Stereotypes." *Florida Law Review* 69 (1): 63–113.

LexisNexis. 2022a. "Circuit Court Opinions: Dataset." LexisNexis.

LexisNexis. 2022b. "Supreme Court Opinions: Dataset." LexisNexis.

Lowe, Matt. 2021. "Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration." *American Economic Review* 111 (6): 1807–44.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S. Corrado, and Jeff Dean. 2013. "Distributed Representations of Words and Phrases and their Compositionality." *Advances in Neural Information Processing Systems* 26: 3111–19.

National Association of Women Lawyers. 2019. *Survey Report of the Promotion and Retention of Women in Law Firms*. New York, NY: National Association of Women Lawyers.

Oster, Emily. 2019. "Unobservable Selection and Coefficient Stability: Theory and Evidence." *Journal of Business and Economic Statistics* 37 (2): 187–204.

Pennington, Jeffrey, Richard Socher, and Christopher Manning. 2014. "Glove: Global Vectors for Word Representation." In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 1532–43. Doha, Qatar: Association for Computational Linguistics.

Posner, Richard A. 1995. "Judges' Writing Styles (And Do They Matter?)." *University of Chicago Law Review* 62 (4): 1421–49.

Posner, Richard A. 2000. "An Economic Analysis of the Use of Citations in the Law." *American Law and Economics Review* 2 (2): 381–406.

Posner, Richard A. 2008. *How Judges Think*. Cambridge, MA: Harvard University Press.

Rachlinski, Jeffrey J., Sheri Lynn Johnson, Andrew J. Wistrich, and Chris Guthrie. 2009. "Does Unconscious Bias Affect Trial Judges?" *Notre Dame Law Review* 84 (3): 1195–1246.

Rehavi, M. Marit, and Sonja B. Starr. 2014. "Racial Disparity in Federal Criminal Sentences." *Journal of Political Economy* 122 (6): 1320–54.

Rice, Douglas, Jesse H. Rhodes, and Tatishe Nteta. 2019. "Racial Bias in Legal Language." *Research and Politics* 6 (2): 2053168019848930.

Ricoeur, Paul. 1981. *Hermeneutics and the Human Sciences: Essays on Language, Action, and Interpretation*. New York, NY: Cambridge University Press.

Rodriguez, Pedro L., and Arthur Spirling. 2022. "Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research." *Journal of Politics* 84 (1): 101–15.

Rohde, David W., and Harold J. Spaeth. 1976. *Supreme Court Decision Making*. San Francisco, CA: WH Freeman.

Root, Danielle. 2019. *Women Judges in the Federal Judiciary—A Series of Fact Sheets on the Lack of Diversity on the Bench*. Washington, DC: Center for American Progress.

Rooth, Dan-Olof. 2010. "Automatic Associations and Discrimination in Hiring: Real World Evidence." *Labour Economics* 17 (3): 523–34.

Sarsons, Heather. 2019. "Interpreting Signals in the Labor Market: Evidence from Medical Referrals." Unpublished.

Songer, Donald R. 2008. "Appeals Courts Database 1925–1996." Songer Project. http://www.songer-project.org/us-courts-of-appeals-databases.html.

Starr, Sonja B. 2014. "Estimating Gender Disparities in Federal Criminal Cases." *American Law and Economics Review* 17 (1): 127–59.

Sunstein, Cass R., David Schkade, Lisa M. Ellman, and Andres Sawicki. 2006. *Are Judges Political?: An Empirical Analysis of the Federal Judiciary*. Washington, DC: Brookings Institution Press.

US Census Bureau. 1990. "Frequently Occurring Surnames from Census 1990 Names Files: Dataset."

Washington, Ebonya L. 2008. "Female Socialization: How Daughters Affect Their Legislator Fathers." *American Economic Review* 98 (1): 311–32.