Speaking of Gender: Superiors' Gender Attitudes and Women's Careers

Sebastian Hager

May 6, 2024

Abstract

This paper investigates superiors' revealed preference for diversity. I study the effect of ministers who differ in their gender attitudes on the career trajectories of women in German ministries. I develop a novel measure of gender attitudes: the extent to which a person speaks gender-inclusively. Based on thirty years of parliamentary speeches, I construct a politician-level score measuring the degree of gender-inclusive language. I calculate each minister's score and relate it to self-collected employee-level data from German ministries. Superiors' gender attitudes matter for women's career outcomes: ministers with one standard deviation higher gender attitudes increase women's probability of being promoted. This effect is driven by both male and female ministers.

Ludwig-Maximilians-Universität München, sebastian.hager@econ.lmu.de. I am grateful to Luca Braghieri, Adam Brzezinski, Davide Cantoni, David Card, Florian Englmaier, Guido Friebel, Amory Gethin, Ingrid Haegele, Elena Heller, Jon Piqueras, Steven Raphael, Friederike Reichel, Fabian Waldinger, and Guo Xu for insightful comments and suggestions. Julia Geltinger provided excellent research assistance. I gratefully acknowledge financial support from the Austrian Economic Association (NOeG Dissertation Fellowship) and from the Evidence-Based Economics program of the Elite Network of Bavaria.

I Introduction

When society recruits its leaders from a limited pool of candidates, it loses efficiency (Bertrand, 2018; Hsieh et al., 2019). Despite the progress of the past decades, women are still underrepresented in leadership positions. In Germany in 2023, only 24% of DAX executives were women (Ernst and Young, 2024). In the German public sector, only one in every three civil servants with a leadership role was female (Destatis, 2024). To address the underrepresentation of women in the leadership of organizations, understanding its causes is crucial. Recent research has emphasized the role of superiors in differentially shaping the careers of women and men (e.g., Cullen and Perez-Truglia, 2023; Haegele, 2024). Yet, apart from a few exceptions (e.g., Ronchi and Smith, 2021), little is known about how superiors' gender attitudes affect women's career progression within organizations.

In this paper, I investigate the relationship between superiors' gender attitudes and women's career outcomes. Studying the effects of gender attitudes on labor market outcomes is empirically challenging: gender attitudes are generally unobservable. I address this challenge by introducing a novel measure of gender attitudes based on speech data. I construct this measure for all German ministers and link it to self-collected data on civil servants in leadership positions in German ministries. This allows me to estimate the effect of ministers with different gender attitudes on the probability that female civil servants in the respective ministry get promoted.

In the first part of this paper, I develop a new method for measuring gender attitudes. I construct a score that captures the degree to which a person speaks gender-inclusively. This score relies on the gender-marked nature of the German language: almost every word referring to a person has a female and a male form. While in English, words with gender-specific forms (such as waitress and waiter) are the exception, in German they are the rule. For example, the word for a female doctor (*Ärztin*) is different from the word for a male doctor (*Arzt*). Traditionally, male-specific words are considered generic, i.e., women can be referred to using the male form of a person-specific noun. In recent decades, however, gender-inclusive language has emerged with the explicit goal of making women more visible in language. Instead of referring to a group of male and female doctors by the generic masculine plural (*Ärzte*), the gender-inclusive way of denoting such a mixed-gender group involves both forms (e.g., *Ärztinnen und Ärzte*).

I define the gender inclusivity score as the ratio between the female-specific words and all person-specific words an individual uses. Next, I apply this method to data on German parliamentary speeches and calculate individual-level scores for over two thousand politicians. In several validation checks, I show that the gender inclusivity score is a meaningful way to measure gender attitudes.

I present a series of facts: first, gender-inclusive language in the German parliament has increased by 50% between 1993 and 2018. This is driven by members of all parties and genders. Second, women across the political spectrum use more gender-inclusive language than men. Moreover, male parliamentarians in the 2010s speak as gender-inclusively as did women in the 1990s. Third, left-wing politicians use more gender-inclusive language than right-wing politicians. For example, Greens thirty years ago spoke more gender-inclusively than did Conservatives in 2018. Fourth, there is considerable variation in individuals' gender inclusivity scores. This holds within genders, within parties, and over time.

Further, I show that politicians with a higher gender inclusivity score are more likely to vote in female-friendly ways. In 1997, several Conservative and Liberal members of parliament voted against the criminalization of marital rape. Politicians who voted against the law had a markedly lower gender inclusivity score than those voting in favor. This holds also within parties. These findings add to the validity of using the gender inclusivity score as a measure of gender attitudes.

In the second part of the paper, I use this measure to analyze whether female employees benefit from superiors with higher gender attitudes. I digitize data on civil servants in leadership positions in all German federal ministries over the past four decades. Since I observe civil servants' leadership ranks over time, I can infer whether an individual was promoted in a given year. This allows me to reconstruct around five thousand employees' career trajectories through the German federal administration. In particular, I can observe gender differences in promotion outcomes.

I link the data on civil servants with ministers and their individual gender inclusivity scores. Thus, for each civil servant at each point in time, I observe the score for the gender attitudes of the minister in charge. Regular changes in ministers introduce variation within ministries over time in ministers' gender attitudes. My empirical strategy estimates how female employees' probability of being promoted is affected by ministers who differ in their gender attitudes.

I find that ministers with higher gender attitudes promote more women. At the same time, I find no such effect on men's promotion probabilities. A minister with one standard deviation higher gender attitudes increases women's promotion probability by around two percentage points (around 30% relative to the mean). In my analysis, I hold constant potential selection of ministers into specific ministries and potential effects of governments by including individual and time fixed effects. Across specifications, the result is stable.

Last, I explore heterogeneities of this result. I find that the effect is driven by ministers at the high end and the low end of the gender attitudes distribution. Ministers

in the upper quartile in the gender attitudes distribution increase women's promotion probability by four percentage points relative to ministers in the middle half of the distribution. Further, I investigate whether there are differences in the effect between male and female ministers. I find that, if anything, the effect is driven by male ministers with high gender attitudes.

My study contributes to research at the intersection of three strands of literature. The first literature focuses on careers within organizations, and especially reasons for the underrepresentation of specific groups in leadership. Various channels have been suggested for promotion and representation gaps in organizations, ranging from social interactions with superiors (e.g., Cullen and Perez-Truglia, 2023) to outright discrimination (e.g., Aneja and Xu, 2022). Since gender quotas have become a common policy tool, increased attention has been paid to the effects of female leaders. Overall, the literature finds limited effects of female leaders on the careers of women in their organizations (e.g., Bertrand et al., 2019; Maida and Weber, 2022). However, other characteristics of superiors—such as their attitudes—might be more important for closing an organization's representation gap. For example, Ronchi and Smith (2021) use the birth of a daughter as a positive shock to gender attitudes and thereby show that superiors' gender attitudes impact their female employees' careers. This is where I contribute to the literature: rather than relying on a categorical characteristic of a superior (such as gender or having a daughter), I measure superiors' gender attitudes on a continuum and relate them to rich personnel data on German civil servants over time.

Second, I contribute to a literature that seeks to identify the contribution of taste-based discrimination to observed employment gaps. Becker (1957) modeled discrimination in the labor market as a result of biased employers who differ in their discriminatory preferences. Empirical analyses of Becker's model are complicated by the challenge that measures of such discriminatory preferences are often not available, especially at the firm or manager level (e.g., Charles and Guryan, 2008). An experimental strand in this literature manipulates employers' perceptions of applicants by randomly assigning different characteristics to job applications (e.g., Bertrand and Mullainathan, 2004; Oreopoulos, 2011). While these audit studies can measure discriminatory outcomes, they do not measure discriminatory preferences. To address this limitation, Kline et al. (2022) sent applications to multiple jobs at the same firms and, thus, measured discriminatory bias at the company level. I contribute to the literature on taste-based discrimination by measuring the distribution of gender attitudes of individual superiors. I then relate these individual-level scores to women's career outcomes.

Last, I contribute to a literature that develops methods to measure attitudes and investigates the effects of attitudes on decision-making. Since people can misrepresent

their views in surveys, psychological measures for attitudes such as implicit association tests (IATs) have become prevalent in economics and other social sciences. For example, Glover et al. (2017) use IATs to study the influence of attitudes towards minorities on performance in the workplace, and Carlana (2019) shows that teachers with stereotypical gender attitudes are harmful to girls' school achievements. However, IAT scores have been criticized for being noisy, easy to manipulate, and unstable over time within an individual (e.g. Gawronski et al., 2017; Schimmack, 2021). Thus, other work has focused on individuals' observed behavior to measure revealed attitudes. For example, Ash et al. (2024) construct a score of judges' gender-stereotyped language. They find that more stereotyping judges interact differently with female judges. While I also measure attitudes based on text data over many years, my method is more tractable and relies on no assumptions about word meanings. Moreover, I document a series of facts on the evolution of gender-inclusive language.[1]

II Measuring Gender Attitudes

II.A The Gender Inclusivity Score

I develop a new measure of gender attitudes based on individual-level text data. Given the salience of gender in the grammar of many languages, attitudes are likely to be revealed in people's speech (Sczesny et al., 2015). I measure how German speakers choose to represent women in their language.

Languages differ in the way gender enters their grammar. For example, German and French assign grammatical gender to every noun, including inanimate objects such as sausage or croissant. English, in contrast, does not have grammatical gender, but only gender-specific pronouns (she and he). As a result, gender-specific words for people (such as *waitress* and *waiter*) are the exception. In German they are the rule: there is a gender-specific noun for a female doctor (*Ärztin*) and for a male doctor (*Arzt*), for a female colleague (*Kollegin*) and for a male colleague (*Kollege*), and so on. Almost every noun denoting a person has a female-specific and a male-specific form.[2]

Traditionally, male-specific nouns for people are considered generic. The so-called generic masculine means that the male-specific term can also be used for women. For example, the word for protester (*Demonstrant*) can refer to both men and women, even though there is a female-specific form (*Demonstrantin*). Especially in the plural, this

[1] In studying the social norm of gender-inclusive language over time, I also contribute to a literature on the evolution of culture and social norms (e.g., Young, 2015; Giuliano and Nunn, 2020).

[2] A few person-specific words in German are grammatically neutral, and thus can refer to both men and women. For example, a member of parliament is always a *Mitglied* (member) of the Bundestag.

leads to an oddity: when 99 women gather in a protest, they are *99 Demonstrantinnen*; should only one man join, they suddenly become *100 Demonstranten*.

Feminist linguists and cognitive scientists have criticized the generic masculine as a sexist convention (e.g., Hofstadter, 1985; Saul et al., 2022). Thus, feminists have proposed that equality of the genders must not stop at grammar: women, too, should be visible in language. They have suggested gender-inclusive language as an alternative to the convention of the generic masculine (Journalistinnenbund, 2024b). From the feminist or gender-inclusive point of view, a mixed-gender group of protesters should be denoted as what they are: *Demonstrantinnen und Demonstranten* (i.e., female and male protesters).[3] Thus, gender inclusivity is a salient feature in German; more so than it is in English, for example. Over the past decades, gender-inclusive language has become widely used in public discourse. Politicians, in particular, have started to use gender-inclusive language—even among the conservative end of the spectrum (see Figure 3).

Using individual-level speech data, I can measure the extent to which a person chooses gender-inclusive terms. This is a continuous measure since speaking gender-inclusively is not a binary choice. One person might constantly stick to gender-inclusive language, while another person only occasionally uses gender-inclusive terms. I define the gender inclusivity score as the ratio of female-specific words to all person-specific words an individual uses:

$$Inclusivity_i = \frac{\#(\text{female-specific nouns})_i}{\#(\text{female-specific nouns})_i + \#(\text{male-specific nouns})_i}. \tag{1}$$

Drawing on a list of female-specific and male-specific words (see Appendix Appendix A.2), I count how many times individual i uses female and male forms for words referring to people; i.e., this excludes other nouns such as apple or philosophy. I then calculate the proportion of female-specific nouns relative to all person-specific nouns (i.e., the sum of female-specific and male-specific nouns). Thus, for each individual i, I construct the ratio $Inclusivity_i$, ranging from zero to one. For example, consider a sentence referring to citizens, voters, and taxpayers; it uses the gender-inclusive form only for citizens but not for voters and taxpayers. Out of the four person-specific words in this sentence only one is female-specific; the ratio is thus 1/4.[4]

This measure has some important advantages: first, it is tractable and intuitive. While other text-based measures of attitudes rely on assumptions about word meanings to capture the association between gender and certain stereotypes (e.g. Ash et al., 2024), my score does not. Instead, it counts gender-specific words. This makes the score easy

[3]Alternative gender-inclusive ways to refer to this mixed-gender group in writing are, for example, *DemonstrantInnen*, *Demonstrant:innen*, and *Demonstrant*innen* (see Journalistinnenbund, 2024a).

[4]See Appendix Appendix A.1 for an example from a German parliamentary speech.

to interpret: a person speaking gender-inclusively all the time has a score of 0.5, whereas someone who exclusively speaks in the generic masculine has a score of zero. Thus, I can compare individuals' scores, both qualitatively and quantitatively. Hypothetically, a person who speaks in the unconventional generic feminine would have a score of one. However, in practice, this does not happen and thus the score is capped at around 0.5 (see Table 1).

Second, with a large corpus of text data spanning many years, the gender inclusivity score is constructed from many speeches over time. Thus, it is a less noisy measure than other scores, such as IATs. Yet, this already hints at the major practical challenge of constructing this score: the availability of comprehensive speech data on individuals.

II.B Gender Inclusivity Among German Politicians

I illustrate my gender inclusivity score using data from Rauh and Schwalbach (2020) on all speeches in the German parliament, the Bundestag, between March 1991 and December 2018. To make the score more tractable, I compile a list of the 100 most used person-specific words in German parliamentary speeches (see Appendix Table A.1 for the ten most used words). These words amount to 95% of all instances a person-specific word was used in parliamentary speeches. I count for each individual or group how often they use the female-specific and male-specific versions of these words. I then calculate the ratio between the female-specific words and all person-specific words (see Appendix Appendix A.2 for more details on the construction of the score). I report summary statistics on the politicians in Table 1.

Table 1: Summary Statistics on German MPs

Variable	Mean	Std. Dev.	Median	Min.	Max.	Obs.
Gender Inclusivity Score	0.11	0.10	0.09	0	.54	2,144
Years in Data	8.45	6.26	7.00	1	28	2,156
Share Female	0.31	0.46				2,156
Share Linke	0.08	0.28				2,156
Share Grüne	0.08	0.27				2,156
Share SPD	0.30	0.46				2,156
Share FDP	0.11	0.31				2,156
Share CDU/CSU	0.37	0.48				2,156
Share AfD	0.04	0.20				2,156

Notes: The table reports descriptive statistics on German MPs between 1991 and 2018. *Years in Data* reports the number of years in which an individual politician has given a speech. *Gender Inclusivity Score* reports the measure of gender attitudes as suggested in Equation (1). To account for a small fraction of MPs who switch parties over time, I assign each MP their modal party affiliation.

In this first set of findings, I present four facts about gender-inclusive language in the German parliament. The first fact is that gender inclusivity has increased over time (see Figure 1). The average level of gender inclusivity in the Bundestag has almost doubled from 0.08 in 1991 to around 0.15 in 2018. This means that in 2018 out of 100 person-specific nouns used in all speeches in the German parliament 15 are female-specific. This is in line with the evolution of feminism and the social norm of using gender-inclusive language in the public sphere.

Figure 1: Gender Inclusivity of German MPs Over Time

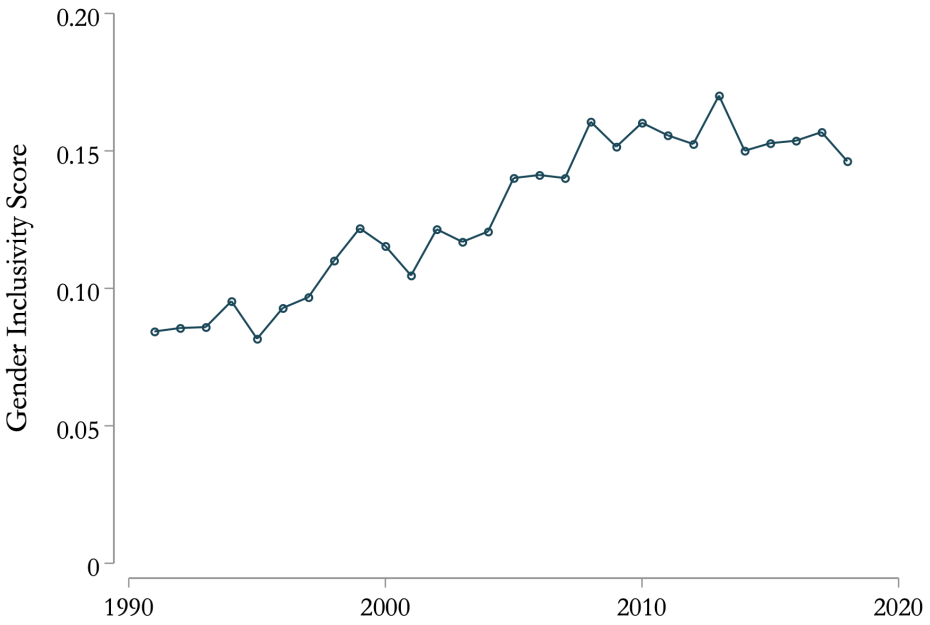

Notes: This figure plots the yearly average of the gender inclusivity score over all speeches in the German parliament.

The second fact is that women speak more gender-inclusively than men (see Figure 2). While politicians of both genders have increased their use of gender-inclusive language, men have consistently been lagging behind women. Male MPs in the 2010s spoke as gender-inclusively as did female MPs in the 1990s.

Third, politicians from left-wing parties speak more gender-inclusively than politicians from right-wing parties (see Figure 3).[5] The conservative CDU/CSU and the liberal FDP have continuously increased their use of gender-inclusive language. Nevertheless, conservative MPs in 2018 spoke less gender-inclusively than the Greens in 1991. Moreover, while the center-left SPD in the early 1990s was closer to the right-wing parties, it is

[5]Throughout this paper, I denote the Greens (Grüne), the Left (Linke), and the Social Democrats (SPD) as left-wing parties, and the Liberals (FDP), the Conservatives (CDU/CSU), and the Alternative for Germany (AfD) as right-wing parties. I exclude the AfD, a xenophobic far-right party, from most analyses, because they only entered parliament in 2017.

Figure 2: Gender Inclusivity of German MPs Over Time, by Gender

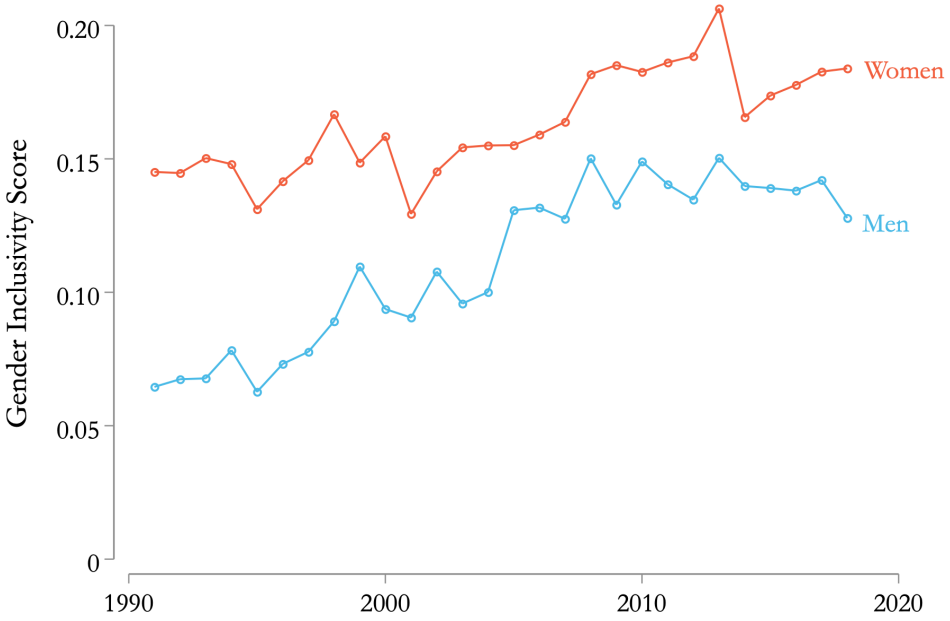

Notes: This figure plots the yearly average of the gender inclusivity score for men (blue) and for women (orange) in the German parliament.

Figure 3: Gender Inclusivity of German MPs Over Time, by Party

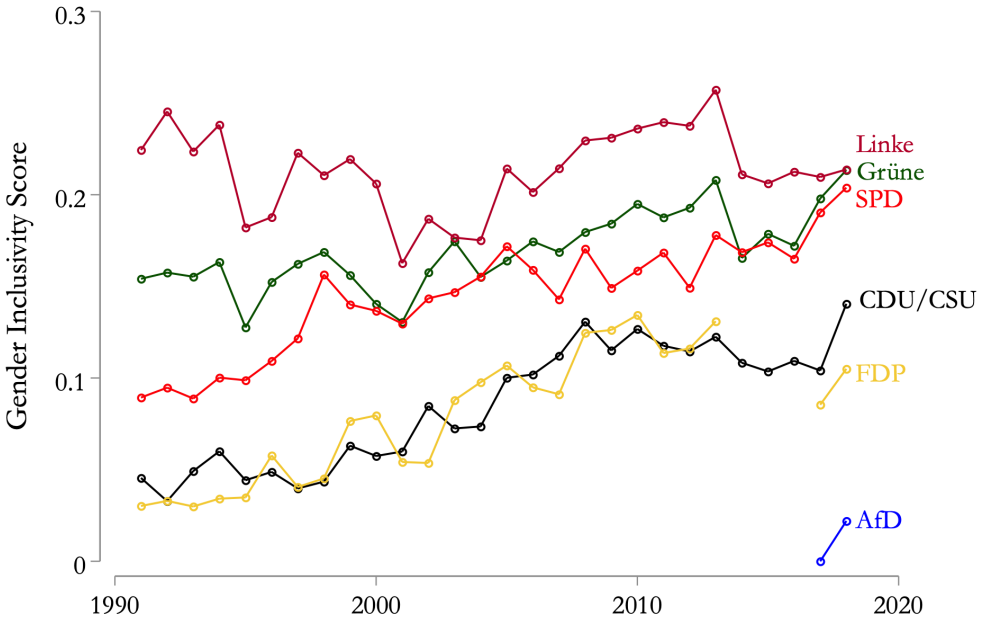

Notes: This figure plots the yearly average of the gender inclusivity score for each party in the German parliament: the Left (Linke, dark red), the Greens (Grüne, green), the Social Democrats (SPD, red), the Liberals (FDP, yellow), the Conservatives (CDU/CSU, black), and the Alternative for Germany (AfD, blue).

nowadays indistinguishable from the other left-wing parties in its use of gender-inclusive language. The far-right and reactionary AfD almost never speak gender-inclusively. In fact, they speak less gender-inclusively than did Conservatives over three decades ago. These findings are not driven by more women being in left-wing parties. I also show that these cross-party patterns hold for men and women separately (see Appendix Figures A.1 and A.2.)

Figure 4: Distribution of Gender Inclusivity Scores, by Party

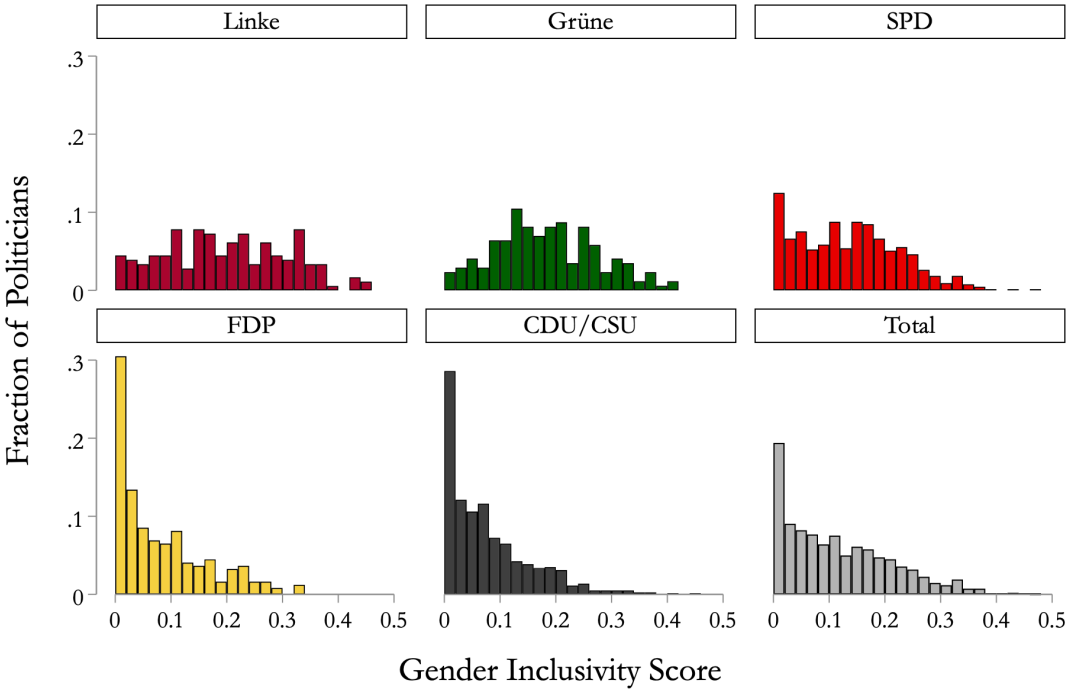

Notes: This figure plots the distribution of individual politicians' gender inclusivity scores within the five main parties. From left to right, these are: the Left (Linke, dark red), the Greens (Grüne, green), the Social Democrats (SPD, red), the Liberals (FDP, yellow), the Conservatives (CDU/CSU, black). I also show the distribution over all politicians in the bottom right panel. I construct the politician-level gender inclusivity score based on all speeches of a politician.

Fourth, there is considerable variation in individuals' gender inclusivity scores.[6] Above, I have shown that gender inclusivity meaningfully tracks political ideology. Yet, Figure 3 only shows the party-level means of politicians' gender inclusivity. Instead, I now construct a time-invariant politician-level score and plot the distribution of scores for each party (see Figure 4). The further one moves to the right on the political spectrum, the more skewed is the distribution toward gender-uninclusive language. While the Left party has an almost uniform distribution, the Greens and the SPD exhibit more bell-like

[6]I plot the distribution of gender inclusivity scores for all politicians in Appendix Figure A.3. I also show that this distribution persists over time in Appendix Figure A.5.

distributions. The CDU/CSU and the FDP have a high proportion of individuals who virtually never use female-specific words in their speech. Yet, even within political parties, there is considerable variation in individuals' gender inclusivity scores. Regardless of the shape of the distribution, in every party there are some gender-inclusive and some gender-uninclusive politicians.

II.C Validity of the Measure

These findings suggest that my gender inclusivity score is a proxy for gender attitudes. This interpretation is in line with studies in social psychology that show that attitudes affect gender-inclusive language (e.g., Sczesny et al., 2015). To further strengthen the validity of my score as a measure for gender attitudes, I show that politicians who speak more gender-inclusively also vote more female-friendly.

In the German political system, MPs usually vote in line with the party leadership. In rare cases, however, MPs within a party are split and can vote against party lines. This was the case in 1997 when the German parliament voted on a law criminalizing marital rape. While the left-wing parties unanimously voted in favor of this law, the Conservatives and the Liberals voted both ways. Thus, I can use politicians' voting decisions on this law as a validation check for my measure of gender attitudes.

I collect data from plenary protocols of the German parliament, which list MPs' voting decisions (Deutscher Bundestag, 1997). I then link politicians' voting decisions and their gender inclusivity scores. I show that those voting in favor of criminalizing marital rape also speak more gender-inclusively (see Figure 5). Among the CDU/CSU and the FDP, those voting in favor of the law had around 30% higher gender inclusivity scores than those voting against it. I also show that this pattern persists within parties (see Appendix Figure A.6). This finding adds to the validity of using the gender inclusivity score as a proxy for gender attitudes.

Figure 5: Voting on the Criminalization of Marital Rape

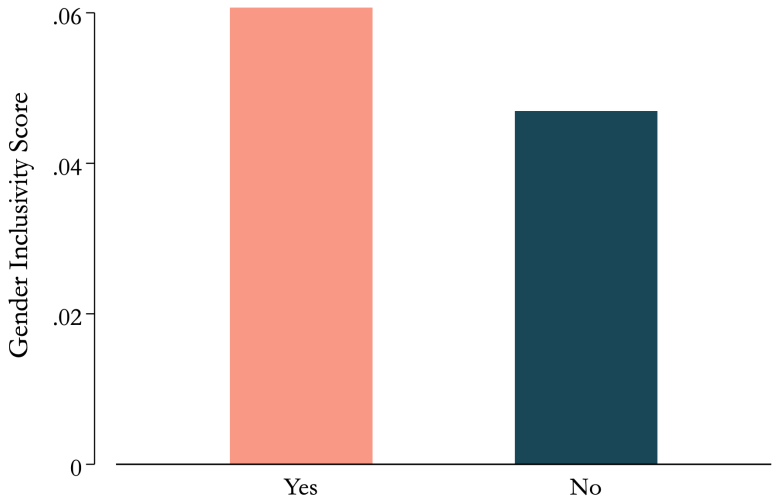

Notes: This table shows the average gender inclusivity score of MPs from the Conservatives (CDU/CSU) and the Liberals (FDP) by whether they voted in favor of a 1997 law criminalizing marital rape.

III Data on German Civil Servants

In the second part of my paper, I use this measure to investigate the role of superiors' gender attitudes in shaping women's careers. To study whether ministers with high gender attitudes are more likely to promote women, I collect employment data on civil servants in German ministries and I construct gender inclusivity scores for all ministers.

Ministerial Employment Data

I digitize data on the population of civil servants in leadership positions in German federal ministries. At least once a year, every German ministry publishes an official organization chart, listing all civil servants who hold a leadership position (see Appendix Figure A.8 for an example). These charts detail the organizational structure of the ministry: each section (*Referat*; henceforth: low level leadership) belongs to a subdivision (*Unterabteilung*; middle level), which in turn belongs to a division (*Abteilung*; high level). At the top level, ministries are led by a small number of state secretaries (*Staatssekretär:in*) and the minister. Every civil servant above and including low level leadership is listed in these organization charts.

I collect organization charts for all ministries from records at the German Federal Archives (*Bundesarchiv*). I digitize a total of 138 organization charts at ten points in time from 1982 to 2016. I always choose the year before an election, so I observe civil

servants' career outcomes at the end of a minister's appointment. These data contain 20,103 observations at the individual-year level.[7]

One advantage of working with these data is that the organizational structure within a ministry remains stable over time. Moreover, all ministries have the same hierarchy of state secretaries, high level, middle level, and low level leaders. This institutional setting is crucial for my main analysis because it allows me to systematically track female employees' career trajectories over time. Moreover, using these data I can compare female representation in different leadership levels across ministries and over time.

Figure 6: Share of Women in Different Leadership Levels

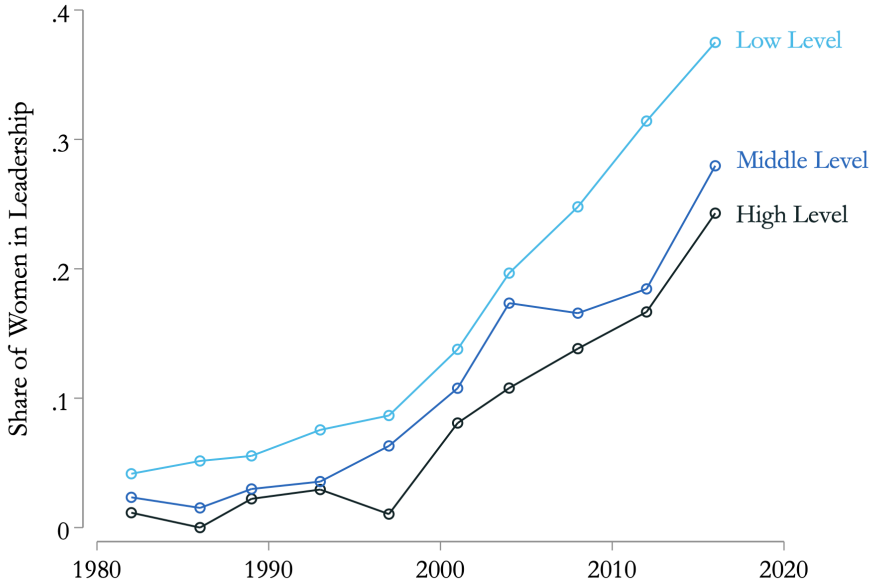

Notes: This figure plots the share of women in different leadership ranks in German ministries: low level (light blue), middle level (medium blue), and high level (dark blue).

My study period coincides with the entry of women into higher leadership ranks of the German civil service. The share of women among low level leaders rose from 5% in 1982 to over 35 40% in 2016 (see Figure 6). The share of female leaders is lower in higher ranks. Yet, there was a considerable increase of women in high level leadership: from virtually no women in 1997 to around 23% in 2016. This increase was not uniform across ministries (see Appendix Figure A.7). Some ministries (e.g., the Ministry for International Cooperation and Development) have appointed many women to leadership ranks. At the same time, other ministries have been lagging behind. For example, in the Finance Ministry in 2016, 23% of its low level leaders were women. At higher levels, this share is even lower.

[7]My final dataset excludes civil servants at the Ministry of Defence. The Ministry of Defence does not indicate the names of its civil servants in organization charts.

Final Dataset on Civil Servants and Ministers' Gender Attitudes

My dataset contains 4,973 civil servants who appear at least twice. For each civil servant in each year, I observe their name, gender, their position, their pay band, and who their direct superior is. I then link individuals over time and thus construct for each civil servant a history of their career steps. Thus, I can observe, at the individual level, which civil servant was promoted in a given year. Importantly for my analysis, I construct an indicator variable $\mathbb{1}[Promotion]_{imt}$ that is equal to one if individual i working in ministry m was promoted from low level leadership to middle level leadership between $t-1$ and t (usually a period of four years).[8]

Among the civil servants who I link over time, 16% are female (see Table 2). 79% of all civil servants in my data were at some point low level leaders, whereas only nine percent were ever promoted to high level leadership. 28% of all civil servants were promoted at least once.

Table 2: Summary Statistics

Variable	Mean	Std. Dev.	Median	Min.	Max.	Obs.
Panel A: Civil Servants						
Number Times in Data	3.25	1.30	3	2	10	4,973
Share Female	0.16	0.36				4,973
Share Ever Promoted	0.28	0.45				4,973
Share Ever State Secretary	0.01	0.12				4,973
Share Ever High Level	0.09	0.28				4,973
Share Ever Mid Level	0.18	0.39				4,973
Share Ever Lower Level	0.79	0.41				4,973
Panel B: Ministers						
Number Times in Data	1.61	0.93	1	1	6	79
Gender Inclusivity Score	0.11	0.08	0.10	0	.34	79
Share Female	0.28	0.45				79
Share Grüne	0.04	0.19				79
Share SPD	0.33	0.47				79
Share FDP	0.16	0.37				79
Share CDU/CSU	0.46	0.50				79

Notes: Panel A reports summary statistics at the employee level in the ministerial employment data. I include only those employees who appear more than once. Panel B reports summary statistics for all ministers for whom I have speech data and can thus construct the gender inclusivity score.

For each employee-year observation, I observe the minister in charge of the ministry at the time. I calculate the lifetime gender inclusivity score for nearly all ministers who appear in the parliamentary speech dataset. For 79 of these ministers, I calculate their

[8] I can only construct this variable for the second time I observe an individual in the data. Hence, I lose 4,973 observations from my dataset.

gender inclusivity score, covering 127 ministry-year pairs.[9] Just as individual politicians vary in their gender inclusivity scores, so do ministers. For example, among the cabinet in 2020, ministers' gender inclusivity scores ranged from 0.04—i.e., four in every one hundred person-specific words were female-specific—to 0.35 (see Appendix Table A.2). This introduces variation within ministries over time in ministers' gender attitudes (see Figure 7 for an example). Thus, for each civil servant at each point in time, I observe the score for the gender attitudes of the minister in charge.

Figure 7: Example: Gender Inclusivity in a Ministry Over Time

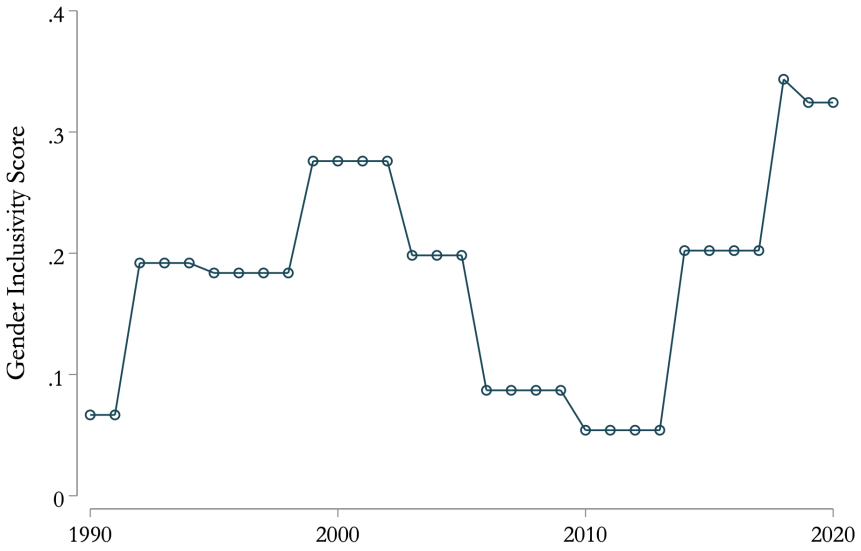

Notes: This figure plots the variation in the gender inclusivity score within a ministry over time using the example of the Ministry for Family Affairs, Senior Citizens, Women, and Youth (BMFSFJ). For each year, it shows the gender inclusivity score of the minister in charge of the ministry. Thus, the level of the score changes when a new minister is appointed. The gender inclusivity score for each minister is calculated using all their speeches in the German parliament between 1991 and 2018.

[9]Only for 14 ministers, covering 18 ministry-year pairs, I cannot construct the gender inclusivity ratio. All but one of these ministers served in the 1980s, which is unsurprising since the parliamentary speech data covers speeches only after 1991.

IV Ministers' Gender Attitudes and Women's Careers

IV.A Empirical Strategy

I study whether ministers with different gender attitudes affect female civil servants' careers.[10] Over time, civil servants in the same ministry experience ministers with different gender attitudes. I use this variation at the ministry level to estimate the effect of ministers with different gender attitudes on the probability that low level leaders are promoted to a higher position in the ministerial hierarchy.

Using my panel dataset on civil servants, I estimate the regression:

$$\begin{aligned}\mathbb{1}[Promotion]_{imt} = \beta_1 \cdot Inclusivity_{mt} + \beta_2 \cdot Inclusivity_{mt} \times FemaleEmployee_i \\ + X'_{mt} \cdot \theta + \alpha_t + \lambda_i + \epsilon_{imt}\end{aligned} \tag{2}$$

where the dependent variable $\mathbb{1}[Promotion]_{imt}$ is an indicator variable equal to one if individual i in ministry m has been promoted from the low level to middle level leadership between $t-1$ and t. $Inclusivity_{mt}$ is the gender attitudes score of the minister in ministry m at time t; its value changes when a new minister gets appointed to ministry m. $FemaleEmployee_i$ is an indicator variable equal to one if i is female. X_{mt} is a vector of control variables for the minister in charge (e.g., minister's gender and party), α_t is a full set of time fixed effects, and λ_i is a full set of individual fixed effects. To account for potential correlation of shocks within ministries across time, I cluster standard errors at the ministry level.[11]

The coefficient on $Inclusivity_{mt}$ captures the effect of ministers' gender attitudes on male civil servants' promotion probability. Gender attitudes are unlikely to predict men's promotion probabilities. Thus, I expect this coefficient to be near zero. The coefficient on the interaction term $Inclusivity_{mt} \times FemaleEmployee_i$ captures the differential effect of ministers with high gender attitudes on female civil servants' promotion probability. If the estimate for this coefficient is positive, this would indicate that, indeed, women's careers benefit from female-friendly ministers. Under the identifying assumption that ministers with higher gender attitudes do not get appointed to ministries in which women were already more likely to be promoted, this effect is causally estimated.

[10] Ministers cannot fire civil servants at will, since civil servants have a high degree of job protection. Ministers do, however, have discretionary powers in the appointment of civil servants to leadership positions (see Goetz, 2007; Jann and Veit, 2010).

[11] To avoid a bias in estimated standard errors due to the small number of clusters (16), I implement a cluster-bootstrap to calculate standard errors.

IV.B Main Findings

I report estimates of Equation (2) in Table 3. The first column includes individual fixed effects and time fixed effects. The point estimate on $Inclusivity_{mt}$ is 0.02 and is not statistically significant. This is in line with the idea that gender attitudes do not affect the promotion probability of men in a ministry. The point estimate on $Inclusivity_{mt} \times FemaleEmployee_i$ is 0.23 and is statistically significant at the 10%-level. An increase in the gender inclusivity score by 0.01 is associated with a 0.23 percentage point increase in women's promotion probability. In other words, ministers with one standard deviation higher gender attitudes increase the promotion probability of women by around two percentage points ($2 \approx 0.23 \times 8.3$; i.e., the point estimate times the standard deviation of ministers' gender inclusivity scores, see Table 2).

Table 3: Main Results

	Dependent Variable: Promotion			
	(1)	(2)	(3)	(4)
Gender Inclusivity	0.02	0.01	0.03	
	(0.06)	(0.06)	(0.06)	
Gender Inclusivity × Female Employee	0.23*	0.23*	0.25**	0.28**
	(0.13)	(0.13)	(0.13)	(0.13)
Individual Fixed Effects	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	
Female Minister Control		Yes	Yes	
Party Minister Controls			Yes	
Ministry × Time Fixed Effects				Yes
Observations	7,751	7,751	7,751	7,751
R^2	0.706	0.706	0.706	0.715
Dependent Variable Mean	6.54	6.54	6.54	6.54

Notes: This table reports estimates of Equation (2). The dependent variable is an indicator equal to one if individual i working in ministry m was promoted from a low level leadership position to middle level leadership at time t. The explanatory variable *Gender Inclusivity* is the gender inclusivity score of the minister in charge of ministry m at time t (on a scale from 0 to 1). The explanatory variable *Gender Inclusivity* × *Female Employee* is the interaction between the variable *Gender Inclusivity* and an indicator equal to one if individual i is female. Column (1) reports estimates from a regression on these two explanatory variables, and individual and time fixed effects. Column (2) reports estimates from a regression which additionally controls for the gender of the minister in charge of ministry m at time t. Column (3) additionally controls for the party of the minister in charge of ministry m at time t. Column (4) includes ministry×time fixed effects. Standard errors are clustered at the ministry level and calculated using the cluster-bootstrap. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1.

One might be concerned that the gender inclusivity score of the minister is correlated with other characteristics of the minister. Thus, I control for the minister's gender

(column (2)) and additionally for the minister's party (column (3)).[12] The estimated coefficients remain stable and statistically significant at the 5% or 10%-level. In column (4), I control for ministry-year fixed effects. These hold fixed all aspects within a ministry at a specific point in time, e.g., characteristics of a minister. Thus, I can no longer estimate the baseline effect of the gender inclusivity score. Yet, I can still estimate the coefficient of interest: the differential effect of ministers with higher gender attitudes on women. The estimated coefficient is 0.28 and significant at the 5%-level.

Across specifications, the effect of ministers with one standard deviation higher gender attitudes on women's promotion probability is around two percentage points. This is a meaningful and sizeable effect, considering that the overall promotion probability from low level to middle level leadership is around 6.5%.

Identification of the impact on women's careers of ministers who differ in their gender attitudes relies on the assumption that female-friendly ministers do not get appointed to ministries in which women were already more likely to be promoted. Of course, ministers are not randomly appointed to lead specific ministries. Yet, ministers are generally appointed as a result of political considerations. To address potential government-level effects (e.g., a shift of the government along the left-right spectrum or a potential shift to an explicitly female-friendly government), I control for time fixed effects. Moreover, ministry fixed effects would control for selection of ministers with higher gender attitudes into specific ministries. Since I observe individual civil servants only within ministries, individual fixed effects absorb ministry fixed effects. Thus, my specification controls for time-invariant ministry-level confounders.

A limitation of my analysis is that I cannot identify the precise mechanism by which superiors with higher gender attitudes affect women's career outcomes. The estimated effect could be explained by both demand and supply forces. For example, a sexist minister might be less likely to choose a woman from the pool of candidates. But, at the same time, female employees might also be less likely to apply for a promotion under a sexist minister. Hence, my empirical strategy estimates the reduced-form effect of ministers with higher gender attitudes on women's promotions.

IV.C Heterogeneous Impact of Ministers

I now turn to an additional analysis and ask: which ministers matter? I study the heterogeneous effects (1) of ministers with particularly high or low gender attitudes scores, and (2) of male relative to female ministers.

[12] In Appendix Table A.3, I show that this result is stable to the inclusion of further control variables.

Ministers With High and Low Gender Attitudes

First, I investigate if ministers with higher or lower gender attitudes drive the effect. I estimate a non-parametric version of my main regression:

$$\begin{aligned}\mathbb{1}[Promotion]_{imt} = \beta_1 \cdot HighIncl_{mt} + \beta_2 \cdot HighIncl_{mt} \times FemaleEmployee_i \\ + \gamma_1 \cdot LowIncl_{mt} + \gamma_2 \cdot LowIncl_{mt} \times FemaleEmployee_i \\ + \alpha_t + \lambda_i + \epsilon_{imt}\end{aligned} \tag{3}$$

where $HighIncl_{mt}$ and $LowIncl_{mt}$ are indicator variables equal to one if the minister in ministry m at time t is in the highest quarter or the lowest quarter in the distribution of ministers' gender inclusivity scores. I interact these indicator variables with $FemaleEmployee_i$. The remaining variable definitions are identical to Equation (2).

I report estimates of this specification in Table 4. In column (1), I report estimates from a regression where I only estimate the effects of ministers who are in the top quartile of the gender inclusivity score distribution, relative to ministers in the lowest 75 percent of the distribution. I again find that the baseline effect of a minister with high gender attitudes on men is insignificantly different from zero. However, the differential effect on women, i.e., the coefficient on the interaction $HighIncl_{mt} \times FemaleEmployee_i$, is positive and statistically significant at the 10%-level. The differential effect of a minister with high gender attitudes on the probability that a woman gets promoted is around four percentage points. In column (2), I report estimates from an analogous regression in which I estimate the effect of ministers with low gender attitudes. The reported estimates are symmetric: the baseline effect on male employees is near zero, but the differential effect on female employees is negative. However, this effect is not statistically significant.

In column (3), I report estimates from a regression with indicator variables on both ministers with high gender attitudes and ministers with low gender attitudes, i.e., of Equation (2). The estimated effects of these variables are relative to ministers between the 26^{th} and 75^{th} percentiles. For both sets of coefficients, I find similar results to the ones in columns (1) and (2). I find null effects on the baseline promotion probability, indicating that under ministers with higher gender attitudes promotions are not more likely for men. They are, however, more likely for women: I find that the differential effect of a minister with high gender attitudes on women's promotion probability is around 3.9 percentage points, which is statistically significant at the 10-% level. The differential effect of ministers with low gender attitudes is -4.5 percentage points. This would indicate that the effect is symmetric, i.e., driven by both ministers with high and low gender attitudes. However, this effect is not statistically significant.

I replicate these results using ministry-time fixed effects, which control for party

Table 4: Non-Linear Effect of Gender Attitudes

	Dependent Variable: Promotion			
	(1)	(2)	(3)	(4)
Top Quarter Inclusivity	0.007		0.007)	
	(0.009)		(0.009)	
Top Quarter Inclusivity × Female Employee	0.041*		0.039*	0.045*
	(0.022)		(0.022)	(0.024)
Bottom Quarter Inclusivity		0.005	0.005	
		(0.031)	(0.031)	
Bottom Quarter Inclusivity × Female Employee		-0.060	-0.045	-0.047
		(0.048)	(0.049)	(0.045)
Individual Fixed Effects	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	
Ministry × Time Fixed Effects				Yes
Observations	7,751	7,751	7,751	7,751
R^2	0.706	0.705	0.706	0.715
Dependent Variable Mean	6.54	6.54	6.54	6.54

Notes: This table reports estimates of Equation (3). The dependent variable is an indicator equal to one if individual i working in ministry m was promoted from a low level leadership position to middle level leadership at time t. The explanatory variable *Top Quarter Inclusivity* is an indicator equal to one if the minister in charge of ministry m at time t is in the top quartile in the distribution of ministers' gender inclusivity scores. The explanatory variable *Bottom Quarter Inclusivity* is an indicator equal to one if the minister in charge of ministry m at time t is in the bottom quartile in the distribution of ministers' gender inclusivity scores. *Top Quarter Inclusivity* × *Female Employee* and *Bottom Quarter Inclusivity* × *Female Employee* are the interactions between these indicator variables and the indicator $FemaleEmployee$, which is equal to one if individual i is female. Columns (1)-(3) include individual and time fixed effects. Column (1) reports estimates from a regression on the variable *Bottom Quarter Inclusivity* and its interaction with *Female Employee*. Column (2) reports estimates from a regression on the variable *Top Quarter Inclusivity* and its interaction with *Female Employee*. Column (3) reports estimates from a regression with both explanatory variables, *Bottom Quarter Inclusivity* and *Top Quarter Inclusivity*, and their interactions with the indicator *Female Employee*. Column (4) reports estimates from a regression with the same explanatory variables as in column (3), except that it includes individual and ministry×time fixed effects. Standard errors are clustered at the ministry level and calculated using the cluster-bootstrap. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1.

and gender of the minister as well as potential government effects over time. In this regression, I can no longer estimate the baseline coefficients β_1 and γ_1 because they vary at the ministry-time level. Estimated coefficients remain broadly unchanged.

Male and Female Ministers

Second, I analyze if the effect of ministers with high gender attitudes is driven by male or female ministers. In this specification, I estimate the effect of gender attitudes on

female employees separately for male and female ministers:

$$\begin{aligned}\mathbb{1}[Promotion]_{imt} = \beta_1 \cdot Inclusivity_{mt} \\ + \beta_2 \cdot Inclusivity_{mt} \times FemEmployee_i \\ + \beta_3 \cdot Inclusivity_{mt} \times FemMinister_{mt} \\ + \beta_4 \cdot Inclusivity_{mt} \times FemMinister_{mt} \times FemEmployee_i \\ + \beta_5 \cdot FemMinister_{mt} \times FemEmployee_i \\ + \beta_6 \cdot FemMinister_{mt} + \alpha_t + \lambda_i + \epsilon_{imt}\end{aligned} \tag{4}$$

where $FemMinister_{mt}$ is an indicator equal to one if the minister in ministry m at time t is a woman. As in the main analysis, I include the variables $Inclusivity_{mt}$ and $Inclusivity_{mt} \times FemEmployee_i$. Additionally, I estimate the differential effect of female ministers with high gender attitudes on female employees, i.e., I include $Inclusivity_{mt} \times FemMinister_{mt} \times FemEmployee_i$. I control for ministry and time fixed effects, and cluster standard errors at the ministry level.

I report the results of this regression in Table 5. Column (1) reports estimates from the main regression, i.e., without the interactions for female ministers, and is included for reference. In column (2), I report estimates of Equation (4). The differential effect of ministers with high gender attitudes on the baseline probability of promotion is near zero (0.04), as is the differential effect of female ministers (-0.02). This again shows that gender attitudes are not driving men's promotions in a ministry, under both male and female ministers. Also, the differential effect of ministers with high gender attitudes on female employees is positive (0.37) and significant at the 10%-level.

If female ministers with high gender attitudes promote more women than male ministers with high gender attitudes, I expect the differential effect of female ministers with high attitudes to be positive. I find that this is not the case. The point estimate is -0.20, which is large relative to the estimated coefficient on $Inclusivity_{mt} \times FemEmployee_i$ (0.37). However, this effect is not statistically significant. I replicate the findings in column (3), where I include ministry-year fixed effects. I conclude that the effect of ministers with high gender attitudes is not driven by female ministers but by ministers of both genders. If anything, the point estimate (-0.20) indicates that male ministers with high gender attitudes are more likely to promote women than female ministers with high gender attitudes.

Table 5: Female and Male Ministers

	Dep. Var.: Promotion		
	(1)	(2)	(3)
Gender Inclusivity	0.03	0.04	
	(0.06)	(0.08)	
Gender Inclusivity × Female Employee	0.25**	0.37*	0.38*
	(0.13)	(0.22)	(0.22)
Gender Inclusivity × Female Minister		-0.02	
		(0.17)	
Gender Inclusivity × Female Minister × Female Employee		-0.20	-0.23
		(0.25)	(0.26)
Individual Fixed Effects	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	
Party Minister	Yes	Yes	
Female Minister	Yes	Yes	
Female Minister × Female Employee Control		Yes	Yes
Ministry × Time Fixed Effects			Yes
Observations	7,751	7,751	7,751
R^2	0.706	0.706	0.715
Dependent Variable Mean	6.541	6.541	6.541

Notes: This table reports estimates of Equation (4). The dependent variable is an indicator equal to one if individual i working in ministry m was promoted from a low level leadership position to middle level leadership at time t. The explanatory variable *Gender Inclusivity* is the gender inclusivity score of the minister in charge of ministry m at time t (on a scale from 0 to 1). The explanatory variable *Gender Inclusivity* × *Female Employee* is the interaction between the variable *Gender Inclusivity* and indicator equal to one if i is female. The explanatory variable *Gender Inclusivity* × *Female Minister* is the interaction between the variable *Gender Inclusivity* and an indicator equal to one if the minister in charge of ministry m at time t is female. The explanatory variable *Gender Inclusivity* × *Female Minister* × *Female Employee* is the interaction between the variable *Gender Inclusivity*, and an indicator equal to one if the minister in charge of ministry m at time t is female, and an indicator equal to one if individual i is female. Column (1) reports estimates from a regression on the first two explanatory variables, individual and time fixed effects, and indicator variables for the party and gender of the minister in charge of ministry m at time t. It is equivalent to column (3) in Table 3 and is included for reference. Column (2) reports estimates from a regression which additionally includes *Gender Inclusivity* × *Female Minister*, *Gender Inclusivity* × *Female Minister* × *Female Employee*, and *Female Minister* × *Female Employee*. Column (4) additionally includes ministry×time fixed effects and therefore omits time fixed effects and minister controls. Standard errors are clustered at the ministry level and calculated using the cluster-bootstrap. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1.

V Conclusion

In this paper, I introduced a new method to measure gender attitudes: the extent to which a person speaks gender-inclusively. I then presented a series of facts about German politicians' use of gender-inclusive language and argued that my score is a meaningful measure of gender attitudes. I used this measure to estimate the effect of ministers who differ in their gender attitudes on female employment outcomes. I find that ministers with one standard deviation higher gender attitudes increase the promotion probability of women by about two percentage points. This effect holds both for male and female ministers.

This finding is important for policy: while many countries have introduced quotas

for women in leadership (e.g., boards of advisors or public bodies), research has shown that the representation of women at higher levels often does not trickle down to lower levels (e.g., Bertrand et al., 2019; Maida and Weber, 2022). My findings indicate that female-friendly superiors—as opposed to female superiors—might be more beneficial for women's careers.

This paper leaves two aspects of the role of superiors' gender attitudes unanswered. First, I estimate a reduced-form effect, which could be explained by both demand and supply. Second, my findings do not show that an intervention that makes superiors speak more gender-inclusively or have more female-friendly gender attitudes (e.g., diversity training) would have a causal effect on their likelihood of promoting women. Yet, my findings do suggest that in the absence of female-friendly superiors women's underrepresentation in leadership may persist.

References

Aneja, Abhay and Guo Xu (2022). "The Costs of Employment Segregation: Evidence from the Federal Government Under Woodrow Wilson". *The Quarterly Journal of Economics* 137(2), 911–958.

Ash, Elliott, Daniel L. Chen, and Arianna Ornaghi (2024). "Gender Attitudes in the Judiciary: Evidence from US Circuit Courts". *American Economic Journal: Applied Economics* 16(1), 314–350.

Becker, Gary S. (1957). *The Economics of Discrimination*. Chicago: University of Chicago Press.

Bertrand, Marianne (2018). "Coase Lecture – The Glass Ceiling". *Economica* 85(338), 205–231.

Bertrand, Marianne, Sandra E. Black, Sissel Jensen, and Adriana Lleras-Muney (2019). "Breaking the Glass Ceiling? The Effect of Board Quotas on Female Labour Market Outcomes in Norway". *Review of Economic Studies* 86(1), 191–239.

Bertrand, Marianne and Sendhil Mullainathan (2004). "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination". *American Economic Review* 94(4), 991–1013.

BMFSFJ, (2024). *Organigramm*. URL: https://www.bmfsfj.de/bmfsfj/ministerium/ministerin-hausleitung/organigramm (visited on 03/07/2024).

Bundesregierung, (2020). *Regierungserklärung von Bundeskanzlerin Merkel*. URL: https://www.bundesregierung.de/breg-de/service/archiv/alt-inhalte/regierungserklaerung-von-bundeskanzlerin-merkel-1807160 (visited on 03/07/2024).

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

Charles, Kerwin Kofi and Jonathan Guryan (2008). "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination". *Journal of Political Economy* 116(5), 773–809.

Cullen, Zoë and Ricardo Perez-Truglia (2023). "The Old Boys' Club: Schmoozing and the Gender Gap". *American Economic Review* 113(7), 1703–1740.

Destatis, (2024). *Gleichstellungsindex 2023*. URL: https://www.destatis.de/DE/Themen/Staat/Oeffentlicher-Dienst/Publikationen/_publikationen-innen-gleichstellungsindex.html (visited on 02/22/2024).

Deutscher Bundestag, (1997). *Stenographischer Bericht, 175. Sitzung, Donnerstag, den 15. Mai 1997*. URL: https://dserver.bundestag.de/btp/13/13175.pdf (visited on 03/07/2024).

Ernst and Young, (2024). *Neuer Rekordwert: Vorstände deutscher Top-Konzerne so weiblich wie noch nie*. URL: https://www.ey.com/de_de/news/2024/01/ey-mixed-leadership-barometer-januar-24 (visited on 02/22/2024).

Gawronski, Bertram, Mike Morrison, Curtis E. Phills, and Silvia Galdi (2017). "Temporal Stability of Implicit and Explicit Measures: A Longitudinal Analysis". *Personality and Social Psychology Bulletin* 43(3), 300–312.

Giuliano, Paola and Nathan Nunn (2020). "Understanding Cultural Persistence and Change". *The Review of Economic Studies* 88(4), 1541–1581.

Glover, Dylan, Amanda Pallais, and William Pariente (2017). "Discrimination as a Self-Fulfilling Prophecy: Evidence from French Grocery Stores". *The Quarterly Journal of Economics* 132(3), 1219–1260.

Goetz, Klaus (2007). "German Officials and the Federal Policy Process: The Decline of Sectional Leadership". In: *From the Active to the Enabling State: The Changing Role of Top Officials in European Nations*. Ed. by Edward C. Page and Vincent Wright. London: Palgrave Macmillan UK, 164–188.

Haegele, Ingrid (2024). "Talent Hoarding in Organizations". *[Working Paper]*.

Hofstadter, Douglas (1985). "Changes in Default Words and Images, Engendered by Rising Consciousness". In: *Metamagical Themas: Questing for the Essence of Mind and Pattern*. New York: Basic Books, 136–158.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow (2019). "The Allocation of Talent and U.S. Economic Growth". *Econometrica* 87(5), 1439–1474.

Jann, Werner and Sylvia Veit (2010). "Politicisation of Administration or Bureaucratisation of Politics? The Case of Germany". *Potsdamer Diskussionspapiere zur Verwaltungswissenschaft* 6.

Journalistinnenbund, (2024a). *Wie geht Gendern beim Schreiben?* URL: https://www.genderleicht.de/schreiben/ (visited on 03/08/2024).

—— (2024b). *Wie geht Gendern beim Sprechen?* URL: https://www.genderleicht.de/sprechen/ (visited on 03/08/2024).

Kline, Patrick, Evan K. Rose, and Christopher R. Walters (2022). "Systemic Discrimination Among Large U.S. Employers". *The Quarterly Journal of Economics* 137(4), 1963–2036.

Maida, Agata and Andrea Weber (2022). "Female Leadership and Gender Gap within Firms: Evidence from an Italian Board Reform". *ILR Review* 75(2), 488–515.

Oreopoulos, Philip (2011). "Why Do Skilled Immigrants Struggle in the Labor Market? A Field Experiment with Thirteen Thousand Resumes". *American Economic Journal: Economic Policy* 3(4), 148–171.

Rauh, Christian and Jan Schwalbach (2020). *The ParlSpeech V2 data set: Full-text corpora of 6.3 million parliamentary speeches in the key legislative chambers of nine representative democracies*. Harvard Dataverse, Version 1. URL: `https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/L4OAKN` (visited on 03/03/2024).

Ronchi, Maddalena and Nina Smith (2021). "Daddy's Girl: Daughters, Managerial Decisions, and Gender Inequality". *[Working Paper]*.

Saul, Jennifer, Esa Diaz-Leon, and Samia Hesni (2022). "Feminist Philosophy of Language". In: *The Stanford Encyclopedia of Philosophy*. Ed. by Edward N. Zalta and Uri Nodelman. Fall 2022. Metaphysics Research Lab, Stanford University.

Schimmack, Ulrich (2021). "The Implicit Association Test: A Method in Search of a Construct". *Perspectives on Psychological Science* 16(2), 396–414.

Sczesny, Sabine, Franziska Moser, and Wendy Wood (2015). "Beyond Sexist Beliefs: How Do People Decide to Use Gender-Inclusive Language?" *Personality and Social Psychology Bulletin* 41(7), 943–954.

Young, H. Peyton (2015). "The Evolution of Social Norms". *Annual Review of Economics* 7(1), 359–387.

Appendix

This appendix presents details on data collection and additional results:

- Appendix A provides details on the construction of the gender inclusivity score.
- Appendix B provides additional figures.
- Appendix C reports additional findings.

A Further Details on the Gender Inclusivity Score

A.1 Example

This example illustrates the construction of the gender inclusivity score using a speech by Angela Merkel in the German parliament on 29 October 2020 (Bundesregierung, 2020). Below is a part of the speech in its German original. Female-specific words are colored in orange and male-specific words in light blue. In the English translation of the text, I indicated the use of gender-specific words by adding these terms in the respective colors.

> "Liebe **Kolleginnen** und **Kollegen**! ... Am 18.10. waren es 769 betreute **Patienten**, am 28.10., zehn Tage später, 1 569. ... Deshalb haben sich die **Regierungschefinnen** und **Regierungschefs** des Bundes und der Länder gestern zu einer weiteren Konferenz getroffen und weitere Vereinbarungen beschlossen. ... Wir haben also gemeinsam mit den **Ministerpräsidentinnen** und **Ministerpräsidenten** Folgendes beschlossen ..."

> "Dear **colleagues** and **colleagues**! ... On the 10$^{\text{th}}$ of October, 769 **patients** were in intensive care; on the 29th of October, 10 days later, this number has reached 1,569. ... Hence, the federal and state **heads of government** and **heads of government** have met yesterday and made further agreements. ... Together with the **minister-presidents** and **minister-presidents** we have come to the following conclusions ..."

In this example, Angela Merkel uses four different person-specific words: colleagues, patients, heads of government, and minister-presidents. For three of these words she uses the gender-inclusive form, i.e., she uses a male-specifc and a female-specific word to describe these people. For one of these words, patients, she only uses the male word. Hence, she has used 7 person-specific words, of which 3 were female-specific. Her gender inclusivity score in this example is thus 3/7.

A.2 Details on Constructing the Gender Inclusivity Score

I calculate the gender inclusivity score for all German politicians in my dataset by counting the number of female-specific and male-specific words they use in their speeches. Hence, I need a list of person-specific words (e.g., baker, terrorist, expert) with their female and male forms. Instead of relying on an arbitrary selection of person-specific words, I draw these person-specific words from the data itself. I construct the gender inclusivity score in three steps:

(1) I compile a list of person-specific words with gender-specific forms. I parse through all speeches and extract the word after the substring "innen und." For example, the word for pilot in German is "Pilot" for men and "Pilotin" for women. The gender-inclusive plural is "Pilotinnen und Piloten." The term after "innen und" is "Piloten," the plural for the male-specific word for pilot. This process gives me a list of 6,000 person-specific words that were used at least once in their gender-inclusive form. (This list contains some mistakes which I manually clean. For example, in the case of the words "gewinnen und verlieren" (in English: "win and lose"), I would extract the string "verlieren," which of course is not a noun.)

(2) For each of these terms, I count how often they appear in all speeches. I rank terms by their frequency and extract a list of the 100 most used person-specific words. Out of all instances a person-specific word was used in a speech (based on my list compiled in step 1), in 95% of these instances the word was among the top 100 words. (I report the ten most used person-specific nouns in speeches in the German parliament in Table A.1.) I compile a list of the respective female and male forms of these 100 words.

(3) For all speeched, I count the number of times each of these 100 words was used in their male form and in their female form.

(4) I then calculate, within a politician, the sum of all female-specific words and the sum of all male-specific words they used in all their speeches. Then I construct the ratio, as defined in Equation (1).

When I construct scores for genders or parties, I do not perform step 4 at the level of the politician, but rather I count the sum of all male-specific and all female-specific words within all speeches of politicians of that specific gender or party. Likewise, when I construct time-varying gender inclusivity scores I count the sum of these words at the gender-year or the party-year level.

Table A.1: Ten Most Used Person-Specific Words

Rank	Male-specific word	Female-specific word	English translation
1	Kollege	Kollegin	Colleague
2	Bürger	Bürgerin	Citizen
3	Soldat	Soldatin	Soldier
4	Arbeitnehmer	Arbeitnehmerin	Employee
5	Mitarbeiter	Mitarbeiterin	Co-worker
6	Verbraucher	Verbraucherin	Consumer
7	Schriftführer	Schriftführerin	Clerk/minute taker
8	Patient	Patientin	Patient
9	Rentner	Rentnerin	Pensioner
10	Wähler	Wählerin	Voter

Notes: This table reports the ten most used person-specific words in speeches in the German parliament.

B Additional Figures

B.1 Gender Inclusivity Among German Politicians

Figure A.1: Gender Inclusivity of Male MPs Over Time, by Party

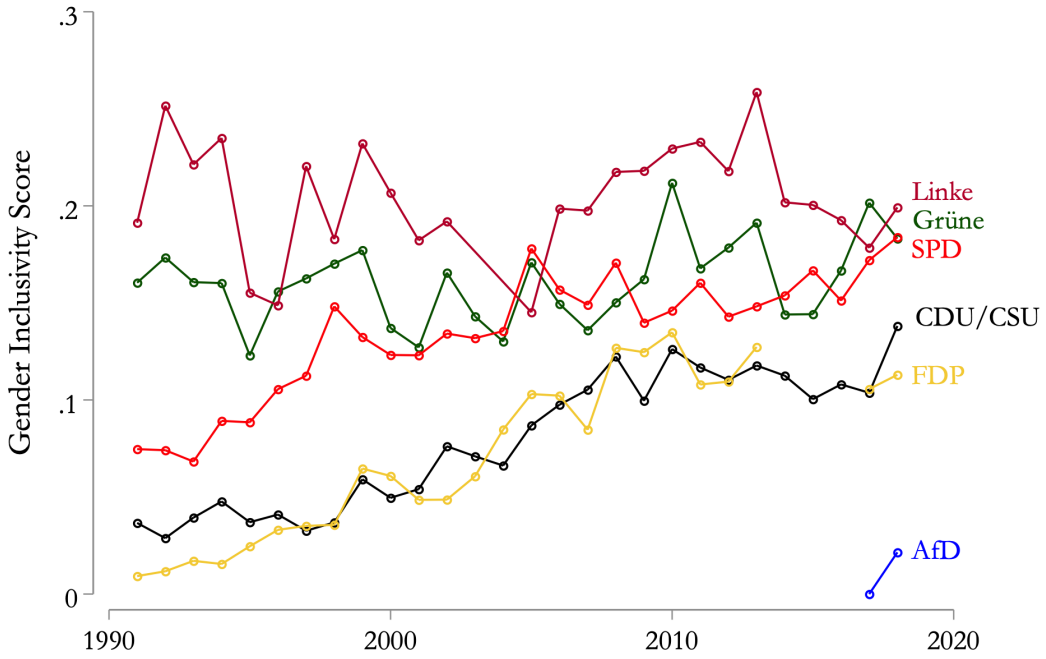

Notes: This figure plots the yearly average of the gender inclusivity score for male MPs in each party: the Left (Linke, dark red), the Greens (Grüne, green), the Social Democrats (SPD, red), the Liberals (FDP, yellow), the Conservatives (CDU/CSU, black), and the Alternative for Germany (AfD, blue).

Figure A.2: Gender Inclusivity of Female MPs Over Time, by Party

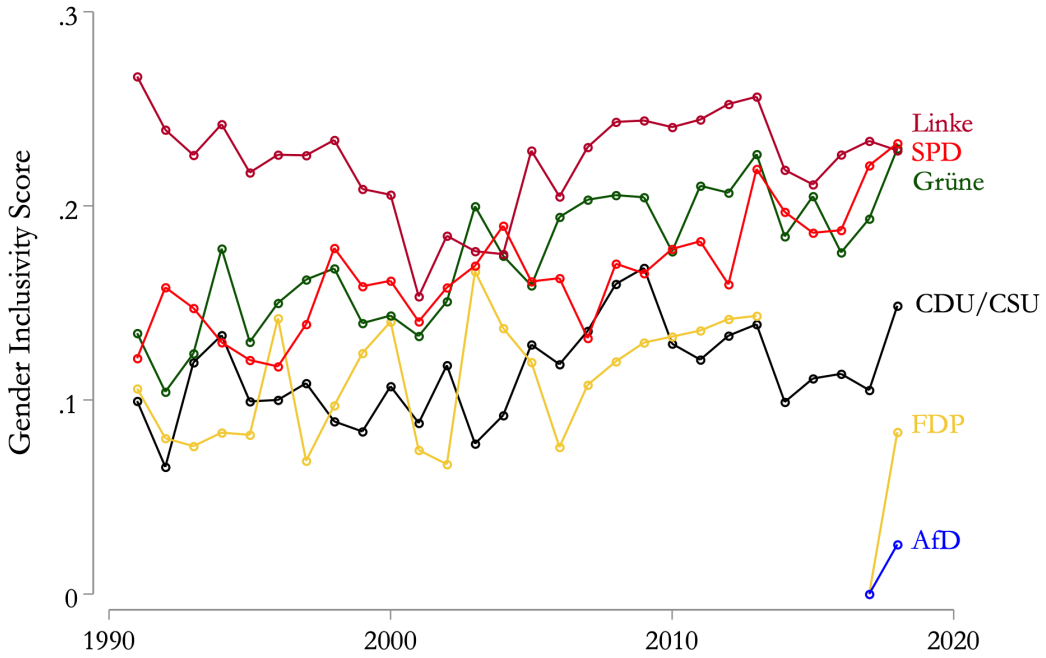

Notes: This figure plots the yearly average of the gender inclusivity score for female MPs in each party: the Left (Linke, dark red), the Greens (Grüne, green), the Social Democrats (SPD, red), the Liberals (FDP, yellow), the Conservatives (CDU/CSU, black), and the Alternative for Germany (AfD, blue).

Figure A.3: Distribution of Gender Inclusivity Scores

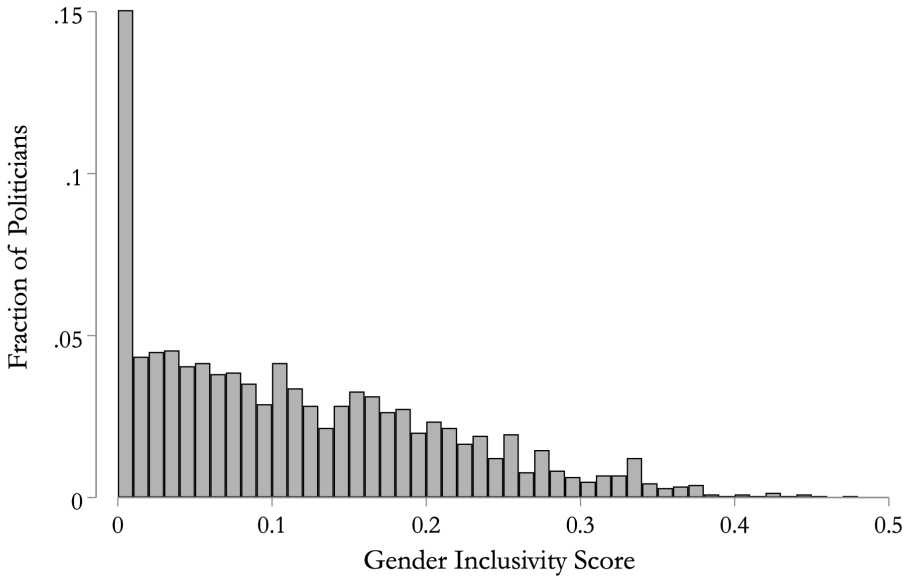

Notes: This figure plots the distribution of individual politicians' gender inclusivity scores.

Figure A.4: Distribution of Gender Inclusivity Scores, by Gender

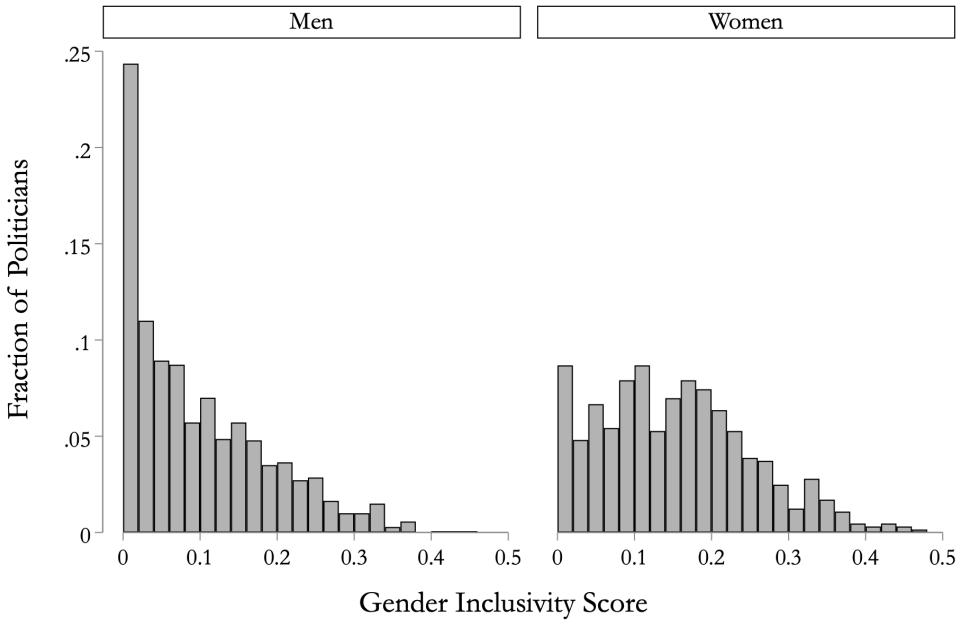

Notes: This figure plots the distribution of male and female politicians' gender inclusivity scores.

Figure A.5: Distribution of Gender Inclusivity Scores, By Decade

Notes: This figure plots the distribution of individual politicians' gender inclusivity scores by decade (1990s, 2000s, 2010s).

Figure A.6: Voting on the Criminalization of Marital Rape, by Party

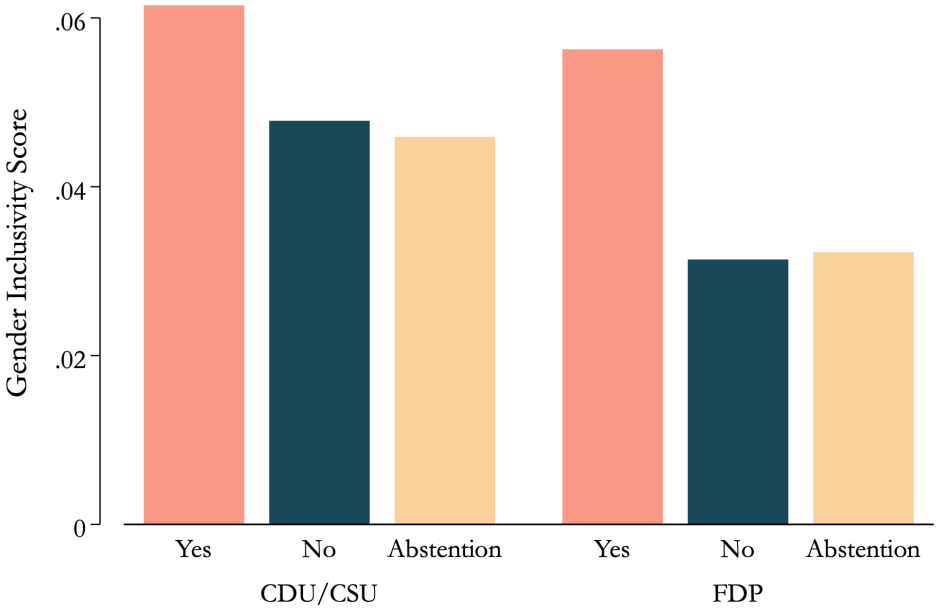

Notes: This table shows, separately for the Conservatives (CDU/CSU) and the Liberals (FDP), the average gender inclusivity score for politicians who voted in favor of a 1997 law criminalizing marital rape, for those who voted against the law, and for those who abstained.

Table A.2: Gender Inclusivity Score Among Ministers in 2020

Minister	Ministry	Gender	Party	Inclusivity
Svenja Schulze	Environment	F	SPD	0.35
Franziska Giffey	Family & Women	F	SPD	0.32
Andreas Scheuer	Transport	M	CSU	0.27
Olaf Scholz	Finance	M	SPD	0.26
Hubertus Heil	Labour	M	SPD	0.23
Julia Klöckner	Agriculture	F	CDU	0.19
Angela Merkel	Chancellor	F	CDU	0.19
Peter Altmaier	Economy	M	CDU	0.19
Heiko Maas	Foreign Affairs	M	SPD	0.18
Gerd Müller	Development	M	CSU	0.14
Christine Lambrecht	Justice	F	SPD	0.13
Helge Braun	Chancellery	M	CDU	0.10
Anja Karliczek	Education	F	CDU	0.07
Jens Spahn	Health	M	CDU	0.06
Horst Seehofer	Interior	M	CSU	0.04
Average				0.18

Notes: This table lists the individual-level gender inclusivity score for each minister in 2020, with information on their ministry, their gender, and party. The Minister of Defence in 2020, Annegret Kramp-Karrenbauer, is omitted from this table due to missing speech data.

Figure A.7: Share of Women in Low Level Leadership

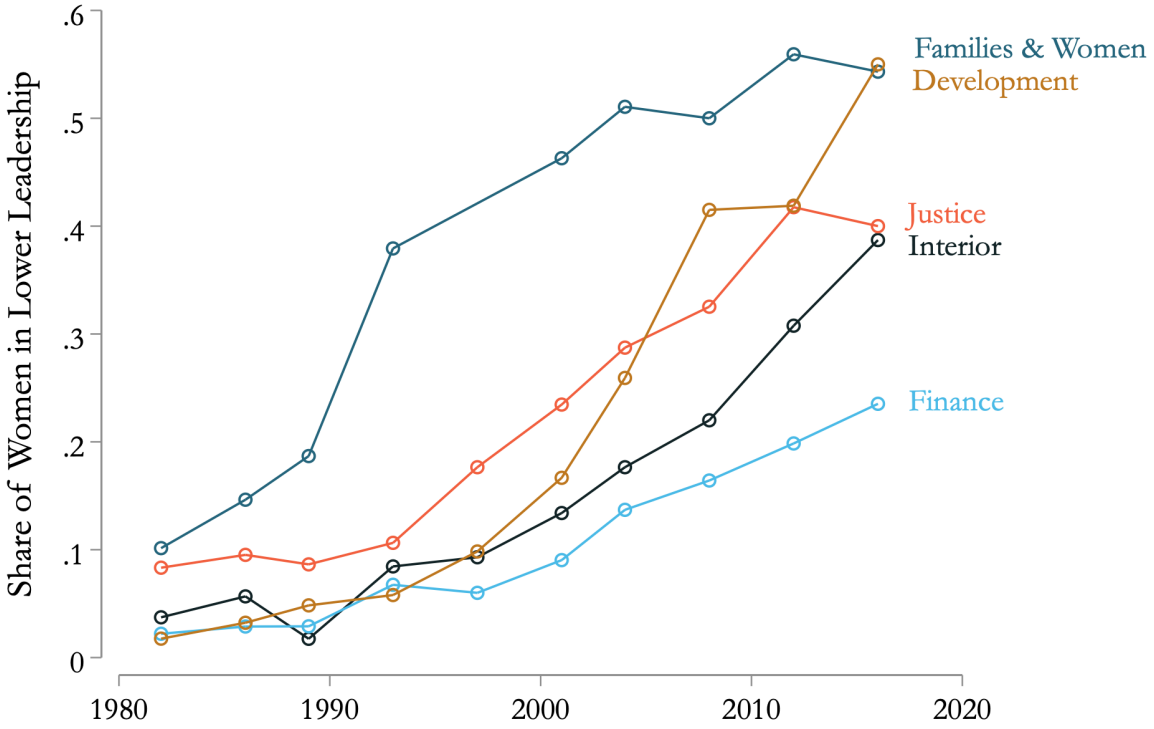

Notes: This figure plots the share of women in low leadership ranks between 1982 and 2016 in five ministries: the German Federal Ministry of Family Affairs, Senior Citizens, Women, and Youth; the Ministry of Finance; the Ministry for International Cooperation and Development; the Ministry of the Interior; and the Ministry of Justice.

Figure A.8: Example Organization Chart

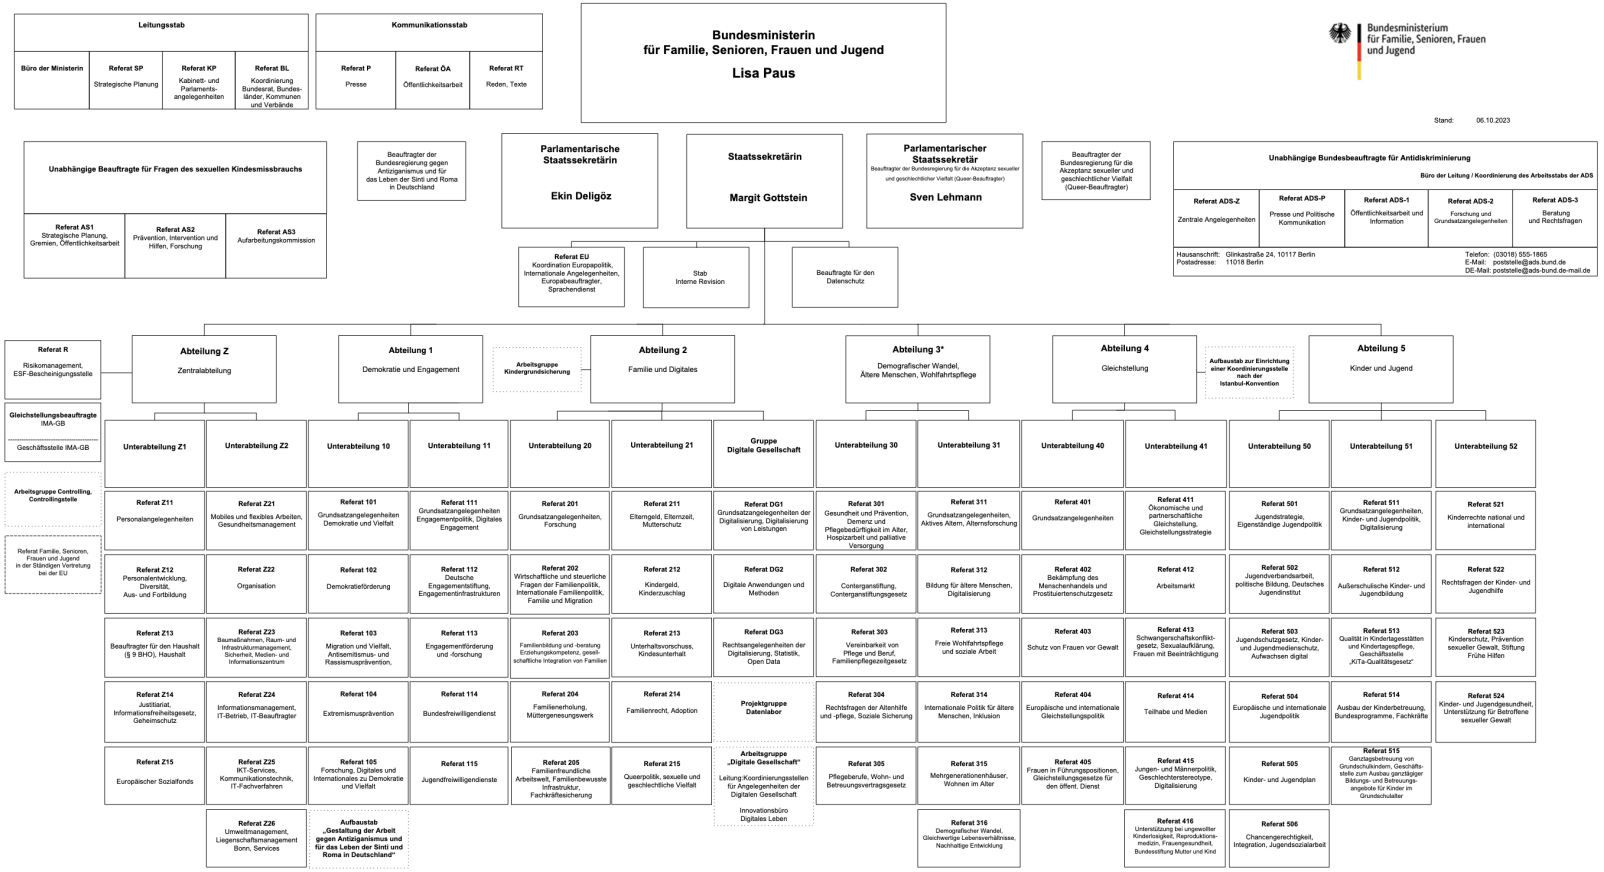

Notes: This figure shows the organization chart of the German Federal Ministry of Family Affairs, Senior Citizens, Women, and Youth in 2023. For data protection, I show a chart without names (unlike the ones I have digitized). Source: BMFSFJ (2024).

C Additional Findings

Table A.3 reports estimates from regressions using additional ministry-level and individual-level controls. The results in the main analysis (here columns (1) and (4), included for reference) are robust to controlling for the share of lower-level leaders who are female at $t-1$ (column (2)). The results are also robust to controlling for indicator variables that capture individual i's years of experience, i.e., how many times I observe them in the data (columns (3) and (5)).

Table A.3: Robustness: Adding Additional Controls

	Dependent Variable: Promotion				
	(1)	(2)	(3)	(4)	(5)
Gender Inclusivity	0.03	0.03	0.00		
	(0.06)	(0.07)	(0.07)		
Gender Inclusivity × Female Employee	0.25**	0.25*	0.30*	0.28**	0.33**
	(0.13)	(0.14)	(0.16)	(0.13)	(0.15)
Individual Fixed Effects	Yes	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes		
Minister Female Control	Yes	Yes	Yes		
Minister Party Controls	Yes	Yes	Yes		
Share of Lower-Level Leaders Control		Yes	Yes		
Years of Experience Controls			Yes		Yes
Ministry × Time Fixed Effects				Yes	Yes
Observations	7,751	7,707	6,042	7,751	6,086
R^2	0.706	0.707	0.720	0.715	0.728
Dependent Variable Mean	6.54	6.55	6.17	6.54	6.16

Notes: This table reports estimates of Equation (2). The dependent variable is an indicator equal to one if individual i working in ministry m was promoted from a low level leadership position to a middle level leadership at time t. The explanatory variable *Gender Inclusivity* is the gender inclusivity score of the minister in charge of ministry m at time t (on a scale from 0 to 1). The explanatory variable *Gender Inclusivity* × *Female Employee* is the interaction between the variable *Gender Inclusivity* and whether individual i is female, i.e., it is equal to zero for men and equal to *Gender Inclusivity* for women. Column (1) reports estimates from a regression on these two explanatory variables, individual and time fixed effects, and controls for the gender and party of the minister in charge of ministry m at time t. This column is equivalent to column (3) of Table 3 and is included for reference. Column (2) additionally controls for the share of women among low level leaders in ministry m at time $t-1$. Column (3) additionally controls for a set of indicator variables for the years of experience of individual i at time t. Column (4) is equivalent to column (1) except that it includes ministry×time fixed effects. This column is equivalent to column (4) of Table 3 and is included for reference. Column (5) additionally controls for a set of indicator variables for the years of experience of individual i at time t. Standard errors are clustered at the ministry level and calculated using the cluster-bootstrap. Significance levels: *** p<0.01, ** p<0.05, and * p<0.1.