Political Sorting in the U.S. Labor Market: Evidence and Explanations*

Sahil Chinoy† and Martin Koenen‡

August 20, 2026

Abstract

We study political sorting in the labor market and examine its sources. Merging voter file data and online résumés to create a panel of 34.5 million people, we show that Democrats and Republicans choose distinctive career paths and employers. This leads to marked segregation at the workplace: a Democrat or Republican's coworker is 10% more likely to share their party than expected based on local partisan shares. Then, we ask whether segregation arises because jobs shape workers' politics or because workers' politics shape their job choices. To study the first, we use a quasi-experimental design leveraging the timing of job transitions. We find that uncommitted workers do adopt the politics of their workplace, but not workers who were already registered Democrats or Republicans. The average effect is too small to generate the segregation we document. To study the second, we measure the intensity of workers' preferences for politically compatible jobs using two survey experiments motivated by the observational data. Here, we find that the median Democrat or Republican would trade off 3% in annual wages for an ideologically congruent version of a similar job. These preferences are strong enough to generate segregation similar to the observed levels.

*First draft: October 2024. For invaluable feedback, we thank Desmond Ang, Jacob Brown, Shani Cohen, Ryan Enos, Ed Glaeser, Benjamin Goldman, Jeff Gortmaker, Jamie Gracie, Larry Katz, Namrata Narain, Amanda Pallais, Vincent Pons, Jesse Shapiro, Advik Shreekumar, Stefanie Stantcheva, and Wilbur Townsend. We also thank Ryan Enos for sharing TargetSmart data and the Harvard Digital Lab for the Social Sciences for assistance with the second survey. We are grateful to seminar and workshop participants at Harvard, MPSA, the BSE Summer Forum, Polmeth, APSA, and the Toronto Political Behaviour Workshop for helpful comments. The survey experiments were approved by the Harvard Committee on the Use of Human Subjects (IRB24-0817) and pre-registered on the AEA registry (AEARCTR-0014137). We gratefully acknowledge funding from the Kenneth C. Griffin Economics Research Fund.

†Stanford University, schinoy@stanford.edu

‡IIES, Stockholm University, martin.koenen@iies.su.se

1 Introduction

American society exhibits widespread sorting along political lines: Democrats and Republicans live in different neighborhoods, date and marry amongst themselves, and get news from like-minded sources.[1] But individuals might be less able to select their coworkers than their friends, partners, and news outlets. The workplace could be a rare site where Democrats and Republicans interact and therefore might facilitate discourse between them (Estlund, 2003; Mutz and Mondak, 2006).

On the other hand, partisanship is a powerful identity (Green et al., 2004; Mason, 2018) that could inform the jobs that workers seek, producing political segregation at work that mirrors other divides. The labor market is unlikely to promote mixing if Democrats and Republicans are willing and able to trade off higher wages for copartisan workplaces. There is scant evidence, however, about how much individuals value politics at work. While some say they would leave a job over political differences (Indeed, 2024), many employees report that they don't even know the parties of their coworkers (Hertel-Fernandez, 2020).

This paper studies the extent and the sources of political sorting in the labor market. First, we use résumé data merged with voter records to document the schools, jobs, and companies that lean Democratic or Republican and show how these differences give rise to marked partisan segregation at the workplace. Second, we ask to what extent segregation arises because individuals adopt the political preferences of their coworkers. Third, we use two survey experiments to quantify the wage trade-offs that workers would make for ideologically compatible employers. Our findings indicate that the pronounced segregation we observe is less because jobs change workers' political behavior and more because, for some workers, politics plays an important role in shaping which jobs they seek.

We begin by building what is, to our knowledge, the largest and most comprehensive dataset that combines worker and political information for the United States. Labor market surveys do not typically collect information about politics and political surveys do not typically ask about an individual's employer. Therefore, we merge online résumé data from LinkedIn, which has information about jobs and education, with comprehensive voter file data, which has information about the party that a voter registered with. Starting from the universe of public U.S. LinkedIn profiles, we match 40% to a voting record to assemble a panel of 34.5 million American workers containing their education, employment, and party affiliation from 2012 to 2021. We estimate that the sample covers 28% of the voting-eligible employed population in the areas we study, which include 30 states and the District of Columbia. Although the sample over-represents white, more educated, higher-income individuals, it has broad coverage of working and voting Americans.

Our first contribution is to use this data to characterize political sorting and segregation in the labor market. We start with sorting: the systematic allocation of workers with different politics to different groups. At five key junctures in a worker's labor market trajectory – the choice of college, major, occupation, industry, and employer – Republicans and Democrats diverge. These different paths are not due simply to the different demographics of the two parties or the political geography of the United States. Moreover, sorting persists within parties and among independents. To show this, we use a within-party measure of ideology inferred from an individual's donations to political candidates and committees (Bonica, 2023). Though we measure this for only 10% of workers in our

[1] For residential segregation and its recent increase, see Glaeser and Ward (2006), Brown and Enos (2021), and Brown et al. (2022); for the contribution of geographic mobility, see Mummolo and Nall (2017) and Martin and Webster (2020). For dating and marriage, see Huber and Malhotra (2017) and Iyengar et al. (2018). For news and social media, see Gentzkow and Shapiro (2011), Levy (2021), Nyhan et al. (2023), and Mangold et al. (2024).

sample, we find that even among those registered as Republicans, more conservative workers work with at companies with a higher share of Republicans. Likewise, more liberal Democrats work at companies with a higher share of Democrats.

Motivated by evidence of extensive sorting, we next quantify segregation: the extent to which Democrats encounter other Democrats at work and Republicans encounter other Republicans. We define a workplace as individuals with the same employer within a commuting zone and measure segregation with the overexposure ratio, or the excess probability that a Democrat or Republican has a coworker of the same party after predicting each worker's affiliation using background characteristics. Since Democrats and Republicans are not evenly distributed across the U.S., our preferred estimate accounts for the shares of partisans in each commuting zone. This ratio is 1.100, which means that the probability that an individual's coworker shares their party is 10% higher than if their coworkers were randomly drawn from the local population. Even after incorporating detailed predictors of partisanship – each worker's demographics, college, major, occupation, industry, as well as the exact Census block group where they live – individuals still face 3.7% higher chances of encountering a co-partisan at work.

To interpret these estimates, we compare them to two benchmarks: residential partisan segregation and other types of workplace segregation. First, the overexposure ratio at the workplace (1.100) is similar to overexposure at the ZIP code level (1.113). Second, we find that partisan segregation is more pronounced than racial segregation at the workplace (1.076), but that workplaces are more segregated by gender (1.151).

After documenting segregation, we analyze its sources. The first mechanism we consider is conformity: segregation could arise because individuals change their political behavior to match that of their coworkers, similar to how they adopt the politics of their neighbors (Perez-Truglia, 2018; Brown, 2024).

Our second contribution is to decompose the relationship between worker and workplace politics into a causal component, which reflects the workplace changing an individual's partisanship, and a residual component due to selection. To do this, we isolate quasi-random variation in *when* a worker was exposed to a partisan workplace by focusing on individuals who switched jobs. For example, we compare two people who made similar moves from workplace A to workplace B, one who moved in 2015 and one who moved in 2017, and then measure their party registration in 2016. The correlation of the second person's partisanship and the partisanship of workplace B reflects selection – unobservable characteristics that make certain types of workers switch to certain types of companies – since the workplace cannot affect them before they arrive. The first person captures the combined selection and causal effect, and, under the assumption that the selection effect is the same for both, the difference between the two isolates the causal effect of the workplace. Identification does not assume that workers' choices of *which* job to take are exogenous, rather that *when* they transition is exogenous, and we show evidence to support this assumption.

We deploy this research design on about 1.1 million people who changed jobs once between 2013 and 2021. For the half who were neither Democrats nor Republicans at the beginning of our sample period, we find that going from a workplace where none of one's coworkers are Republicans (or Democrats) to one where all are them are makes the individual 5.0 p.p. more likely to register as a Republican (or Democrat) themselves. This is a relatively small effect, and the estimated effect of the workplace on partisanship for the other half of workers who start as Democrats or Republicans is essentially zero. We show that about 20% of the overall relationship between workplace and worker partisanship can be attributed to the causal effect of the workplace, and the other 80% to

selection.

Despite high levels of political homophily in the labor market, these results point to the limited role for the workplace in shaping partisanship in the short run – especially in changing the minds of those who already know where they stand. By contrast, selection is strong: even those who haven't *yet* registered as a Republican move to Republican-leaning firms, and vice versa for Democrats. This leads us to study the effect of workers' politics on their job choices.

Our third contribution is to measure workers' preferences for politically compatible workplaces. Job ads do not contain the shares of Democrats and Republicans at the company, so instead of measuring whether workers prefer to work with co-partisans holding all else constant, we study job *characteristics* that Democrats and Republicans value differently. These "ideological amenities" might lead them to sort into different jobs and generate differences in pay. Indeed, we see suggestive patterns in the cross-sectional data: despite attending more selective schools, Democrats are found in lower-paying majors, occupations, and industries. We test whether these differences could reflect preferences for ideological amenities using a framing experiment.

Using our merged LinkedIn-voter file data, we first analyze the text of nearly 300,000 company descriptions, showing that even within the same industry, companies with more Democrats use signaling language like "community," "empower," and "impact" relative to companies with more Republicans. Then, we use a large language model to generate about 30,000 job ads, half that feature these signals and half that do not. We deploy these ads in a personalized online survey, showing respondents in hundreds of occupations ads for job titles they are interested in and randomly varying their salary offers.

We find that Democrats are willing to trade off 2.9% of their salary ($1,700 per year) for job ads that emphasize liberal values, while Republicans are willing to trade off 2.4% ($1,600) to avoid them. These effects for the median Democrat and Republican mask important heterogeneity, which we show in two ways. First, the gap between those who identify as strong Democrats and strong Republicans in their WTP for the liberal framing is 7.1 p.p., about 33% larger. Second, we find that 25% of Democrats would trade off more than 6.0% of their salary for the liberal framing, and 25% of Republicans would trade off more than 6.9% to avoid it.

To relate these estimates back to political segregation, we fit a machine learning model using the relationship between company descriptions and partisanship in the LinkedIn data, then we predict the partisanship of our synthetic job ads. This allows us to interpret our estimates as workers' willingness to pay for a company whose description signals a particular share of partisans. Using this scaling, we show that Democrats would trade off 4.7% ($2,700) of their salary for a job that signals a Democratic share 1 standard deviation (about 24 p.p.) higher, while Republicans would pay 4.3% ($2,800) to go in the opposite direction. These estimates are similar in magnitude to workers' WTP for other amenities such as the ability to telecommute (4.5%) or a relaxed versus fast pace of work (4.4%) from previous studies (Maestas et al., 2023). The difference is that both Democrats and Republicans have positive valuations for these "standard" amenities, while their preferences for ideological amenities go in opposite directions, creating segregation.

While the framing experiment estimates workers' preferences for ideological amenities in a more natural setting, it cannot pinpoint exactly what respondents care about. So, we run a second survey with a conjoint design to disentangle workers' preferences for co-partisan coworkers from other coworker demographics and company characteristics. We again present pairs of job profiles specific to the respondent's current occupation, varying nine job, coworker, and company attributes, and estimate willingness to pay by party for each. Here, we find that Democrats are willing to

trade off 3.3% of their salary for a company with a Democratic share 1 standard deviation higher, while Republicans would pay 4.0% to go in the opposite direction. But we also find that workers have a strong distaste for companies that donate to the out-party, and would give up 8 to 10% of their salary to avoid this.

The two surveys target different estimands that together are informative about not just the strength but also the nature of workers' preferences. The first bundles everything the respondent infers from the job ad: not only coworker political characteristics, but other coworker demographics, company culture, corporate donations, etc. The second experiment isolates the effect of each of these attributes separately. Combined, they show that some workers have strong preferences for ideological amenities – which can be signaled simply by emphasizing different values in the way the same job is described – but that these preferences reflect a notion of politics broader than just whether workers expect to have colleagues who vote like them.

Finally, we interpret the estimates using the canonical Rosen (1986) model of job amenities and compensating differentials. In the model, the difference in Democrats' and Republicans' preferences for ideological amenities leads them to sort into different workplaces, but heterogeneity in their preferences for these amenities limits the segregation that this creates. For example, even if *some* Democrats value these amenities highly, the ones who care more about their salary than their ideological compatibility with their employer will take jobs where they encounter Republicans at work, and overall segregation will be low. Thus, the extent to which this channel can drive segregation depends on the full distribution of worker preferences, not just the average.

Our experiment identifies this distribution, and a simple calibration of the model using our survey estimates generates segregation similar to the level we see in the observational data. This implies that worker preferences are powerful, but it also underscores the key mechanism: the overlap in Democrats' and Republicans' preferences prevents runaway levels of segregation. The framework illuminates how – like other settings in which ideology and market forces collide – the labor market can check the expression of political identity.

This paper contributes firstly to an evolving literature in political science, economics, and finance on politics in the workplace. One branch of this literature shows the effect of partisanship on employers' hiring decisions and employees' job choices. On the employer side, Colonnelli et al. (2024) show that business owners in Brazil are more likely to employ copartisan workers, and Gift and Gift (2015) show that job-seekers with minority partisan affiliations are less likely to be called back. On the employee side, McConnell et al. (2018) show that workers in an online labor market request lower wages when the employer shares their political stance. Hurst and Lee (2024) and Adrjan et al. (2023) show that employers' politically charged statements lead to differential interest in their job postings from Republicans or those living in Republican states. Bermiss and McDonald (2018) and Bonica et al. (2020) show that ideological misfit leads to a higher likelihood of job departure among investment professionals and physicians, respectively.[2] Our paper contributes descriptive evidence on the distinctive choices made by millions of Democratic and Republican workers across industries and occupations, and survey-experimental evidence that quantifies how much workers would trade off for ideologically compatible jobs, which enables us to assess the contribution of these preferences to partisan segregation at work.[3]

[2]Rice and Schiller (2024) show that corporate philanthropy reduces employee turnover, particularly among demographic groups more likely to be Democrats.

[3]Related work studies the political preferences of executives (Fos et al., 2022; Steel, 2024), corporate boards (Hoang et al., 2022), and entrepreneurs (Engelberg et al., 2022). Partisanship also predicts how credit analysts (Kempf and Tsoutsoura, 2021), police officers (Ba et al., 2023), and doctors (Kim, 2024) behave on the job. In

A second branch of the literature studies the effect of jobs on political beliefs and behavior, particularly campaign contributions (Stuckatz, 2022; Matzat and Schmeißer, 2023). Our paper includes a more demanding test of the impact of jobs on politics: we measure the short-run effects of the workplace on party registration, an outcome that changes relatively infrequently, but that we can measure for a large group of workers.[4]

Our descriptive analysis of segregation is most similar to contemporaneous work by Frake et al. (2024), who also measure partisan segregation in the labor market using LinkedIn profiles merged with voter registration data. Our paper dissects the *drivers* of workplace segregation, using a quasi-experimental design and survey experiments to tease apart two causal mechanisms.[5]

Finally, this paper contributes to the economics literature on job amenities. We connect this literature to American politics by measuring how much workers value a job attribute that has been studied less: ideological valence. Our survey designs are modeled after Mas and Pallais (2017) and Maestas et al. (2023), who study workers' willingness to pay for "canonical" amenities like flexible work arrangements and paid time off, as well as Colonnelli et al. (2023), who study workers' preferences for job ads that advertise environmental, social, and governance (ESG) practices in Brazil.[6] We extend these designs by leveraging recent advances in large language models to offer tailored job choices. Then, we use a machine learning algorithm to collapse the high-dimensional text that respondents see to a measure that we can relate back to the observational data.

The rest of this paper proceeds as follows: Section 2 describes our LinkedIn and voter file data, outlines our merging procedure, and presents summary statistics about our sample. In Section 3, we use this data to describe the extent of political sorting in the labor market, then we measure segregation at individual workplaces. We analyze the effects of jobs on partisanship in Section 4. Then, we ask how politics shape job choice in Section 5. Section 6 concludes.

2 Data and sample

To study political sorting in the labor market, we need combined information about workers' careers and their political behavior. We merge online résumé data from LinkedIn, which has information about jobs and education, with administrative voter file data, which has information about partisanship, to create a panel of 34.5 million individuals in 31 states. This section summarizes the data and merging procedure and assesses the representativeness of the merged sample; Appendix B contains more details.

addition, we draw from recent work on how consumers respond to partisan firms (e.g., Panagopoulos et al., 2020; Conway and Boxell, 2024; Wang and Lu, 2022), but we ask how *employees* respond to these partisan values.

[4] Examining a specific occupation, Goncalves and Tuttle (2024) find that becoming a police officer increases Republican party affiliation and contributions to Republican campaigns.

[5] Our paper also relates to an extensive literature on workplace segregation and occupational sorting of women and racial minorities (e.g., Carrington and Troske, 1997; Hellerstein and Neumark, 2008; Bayard et al., 2003; Alonso-Villar et al., 2012) and adds to work measuring segregation in the U.S. more broadly, such as racial residential segregation (Cutler et al., 1999) and consumption segregation (Davis et al., 2019).

[6] A related literature studies workers' preferences for amenities such as "mission-driven" versus "immoral" work (e.g., Burbano et al., 2024; Schneider et al., 2020); dignity at work (Dube et al., 2022); and corporate social responsibility (e.g., Burbano, 2016; Hedblom et al., 2019). While our notion of ideological amenities may encompass some of these features, this literature has not linked labor supply preferences to partisanship, a form of identity that plays a powerful role in organizing American society.

2.1 LinkedIn data

We use nationally comprehensive LinkedIn data scraped between 2016 and 2022, which includes 159 million individual and 15.3 million company pages. This data only has information that individuals post publicly about their education and employment – akin to a résumé – and does not include their connections. Below, we briefly describe the key variables for our analysis.

Education, college, and major: We first infer an individual's highest level of education: if they report any institution associated with a graduate degree, then it is graduate; likewise for bachelor's and associate's.[7] If they report only a high school name, it is high school. If they do not report any education, then it is missing. To obtain college characteristics such as the average SAT score of its admitted class, we match the name of associate's and bachelor's degree schools to the IPEDS database. We also use string matching to assign the majors that individuals report to a "detailed field of degree" (e.g., Petroleum Engineering) as recorded in the American Community Survey.

Industry: Both individual profiles and company pages can report an industry (e.g., Insurance or Real Estate). These follow LinkedIn's own taxonomy, and there are 161 categories in our sample. Our baseline analysis uses these labels without modification, and we crosswalk them to NAICS codes to assess representativeness and for use in robustness checks.[8]

Occupation: We use string matching between the job titles that individuals report and the O*NET database to assign each position a six-digit SOC code (e.g., Graphic Designers).

2.2 Voter file data

We use nationally comprehensive voter file snapshots spanning 2012 to 2021 from the vendor TargetSmart.[9] These records contain 311 million people. While the voter file does not record whom an individual voted for, it does record whether they voted and – critically, for our study – the party that they registered with. Thirty states and the District of Columbia record party affiliation in their voter data; our analysis is limited to these areas.[10] In Section 3, we measure partisanship using the most recent occasion that each individual registered in one of these locations, but in Section 4, we allow partisanship to vary by year.

The primary measure of political behavior in this paper is whether a voter is registered as a Democrat or a Republican, which is highly predictive of vote choice and of ideology defined on a liberal-conservative spectrum (e.g., Bartels, 2000; Levendusky, 2009; Abramowitz and Saunders, 1998). However, we also merge the voter file with campaign finance data from the Database on Ideology and Money in Politics (Bonica, 2023) by name and street address.[11] To study intra-party differences, we use the "DIME score," a continuous measure of political ideology based on an

[7] We do not count an institution if the individual reported that they "attended," "took classes," "did not complete," etc.

[8] The taxonomy has four levels of hierarchy. We map 21 of the 161 categories to a 6-digit NAICS code, 53 to a 5-digit code, 85 to a 4-digit code, 57 to a 3-digit code, and 30 to a 2-digit code.

[9] This data was cleaned and de-duplicated in Brown et al. (2023); see Appendix C of that paper for data quality checks.

[10] Party affiliation is recorded in Alaska, Arizona, California, Colorado, Connecticut, Delaware, the District of Columbia, Florida, Iowa, Idaho, Kansas, Kentucky, Louisiana, Massachusetts, Maryland, Maine, North Carolina, Nebraska, New Hampshire, New Jersey, New Mexico, Nevada, New York, Oklahoma, Oregon, Pennsylvania, Rhode Island, South Dakota, Utah, West Virginia, and Wyoming. The remaining 20 states do not offer voters the option of affiliating with a party when they register.

[11] We successfully match 63% of individuals in DIME to the voter file; see Appendix B for more details.

individual's donations to candidates and committees.[12]

The voter file also provides some demographic information: gender is recorded in many states and imputed or sourced from commercial data by TargetSmart otherwise; race is directly recorded for some or all years in eight states,[13] and sourced from commercial data or imputed by TargetSmart based on name and local demographics otherwise. We observe the exact date of birth of most individuals, and we observe their exact registration address.

2.3 Merging LinkedIn and voter file data

To merge the LinkedIn and voter file data, we mainly use name, location, and (inferred) year of birth, achieving a match rate of 40.4%. We first assign the user's reported location to a Census place if possible, or a county, CBSA, or state if not. Very few people report their birth year on LinkedIn, so we attempt to infer it from their high school or college years of attendance if possible, or from their first reported job if not.[14]

Our merging procedure uses various combinations – 72 in all – of name, birth year, and four geographic identifiers. Roughly speaking, we "zoom out," starting at the Census place level, then the county, then the CBSA, and finally the state. We also use progressively less restrictive criteria for name matches – first the exact name and middle initial, then first and last name only, then nicknames and diminutives, then the first four letters of the first and last name. We first try to match using birth years within a ten year window, then we relax this restriction.[15] Across all methods, we match 45.7% of LinkedIn profiles to at least one individual in the voter file. Finally, we keep only profiles that are matched to *exactly* one individual in the voter file. This leaves 64.3 million individuals, for a match rate of 40.4%. Most individuals are matched by first name, last name, and approximate birth year within a Census place.

2.4 Sample and representativeness

To arrive at our analysis sample, we keep individuals whose latest voter registration is in an area that records party affiliation, leaving 34.5 million people. This is about 41% of the size of the voting-eligible labor force in these areas (84.6 million).[16] However, our sample includes individuals who are not necessarily currently in the labor force; some are students, retired, or not looking for work. Approximately 22.0 million individuals in our sample list an employer as of 2022, which is 27.6% of the number of voting-eligible *employed* individuals in these areas (79.7 million). Our analysis of workplace segregation requires that we observe at least two individuals at a workplace; this "coworker sample" (which excludes self-employed people) contains 20.3 million people.

[12] This donation-based measure aligns with a scaling based on Congressional roll call votes (Bonica, 2014) and predicts individual survey responses to policy questions (Bonica, 2019).

[13] These are Alabama, Florida, Georgia, Louisiana, North Carolina, South Carolina, Tennessee, and Texas.

[14] Specifically, we prioritize lexicographically (1) high school graduation year minus 18, (2) high school start year minus 14, (3) bachelor's graduation year minus 23, (4) bachelor's start year minus 19, (5) start year of first job listed minus 23.

[15] However, if the birth year is reported by the user or inferred from high school or college start or end date, then we always enforce the 10-year birth year window. We also always enforce that middle initials do not conflict.

[16] According to the BLS Local Area Unemployment Statistics annual averages, in 2021, there were 92.3 million people in the areas that record party affiliation. The CPS Voting and Registration Supplement (averaged from 2012 to 2022) indicates that 92.2% of the individuals in the labor force in these areas were eligible to vote; we use the state-specific eligibility rates to estimate that there were approximately 84.6 million voting-eligible workers in the labor force in these areas.

Table 1: Summary Statistics for Analysis Sample

	Mean	Std. dev.	Num. of obs.
Birth year	1974.6	16.6	34,411,975
Male	0.491	0.500	33,788,193
White	0.799	0.400	33,453,568
Black	0.076	0.266	33,453,568
Hispanic	0.081	0.273	33,453,568
Asian	0.037	0.189	33,453,568
Democrat	0.399	0.490	34,533,903
Republican	0.286	0.452	34,533,903
High school	0.066	0.247	20,389,197
Associate's degree	0.106	0.307	20,389,197
Bachelor's degree	0.570	0.495	20,389,197
Graduate degree	0.259	0.438	20,389,197
Missing field of degree	0.622	0.485	34,533,903
Block group HH income	90,288	46,024	33,893,527
Num. connections	131.9	175.7	34,533,903
Num. jobs	3.1	3.2	34,533,903
Missing industry	0.340	0.474	34,533,903
Missing occupation	0.275	0.446	34,533,903
Lists employer in 2022	0.637	0.481	34,533,903
At least one coworker	0.588	0.492	34,533,903
Share Dem. coworkers	0.379	0.316	28,236,416
Share Rep. coworkers	0.276	0.304	28,236,416
Num. coworkers	321.1	1105.7	28,236,416
Has DIME score	0.110	0.313	34,533,903

Notes: "Democrat" and "Republican" refers to registering with the Democratic or Republican party in the latest observed registration. Number of connections is topcoded at 500. DIME is the Database on Ideology and Money in Politics, a database of contributions to political campaigns. Block group-level median household income is from the 2016-2021 ACS. Appendix Table A4 shows these summary statistics separately by party and for other subsets of the analysis sample.

Although our sample is large and broadly representative of the U.S. working and voting population, it is skewed towards young, white, more-educated, higher-income individuals. We present summary statistics in Table 1.[17] The average individual in our sample was born in 1975. About half of the people in our sample are male and 80% are white. While we only observe the highest level of education for about 60% of the sample, about 84% of individuals who list their education report a bachelor's or graduate degree, and fewer than 7% report only a high school degree. We link 11% of the sample to a DIME campaign finance record with an ideology score; for these individuals, we obtain a continuous measure of ideology in addition to party affiliation.

Next, we compare our sample to voters in the labor force sampled in the 2022 Cooperative Election Study who live in an area that records party affiliation (see Appendix Table A1). Compared to this group, our sample is similar in age and gender, but disproportionately white (80% vs. 70%). On partisanship, our sample, with 40% Democrats and 29% Republicans, is similar to the CES,

[17] Appendix Table A4 shows these summary statistics separately by party and for other subsets of the analysis sample.

which has 39% and 27%, respectively. Our sample is also more highly educated than the CES: we over-represent those with a bachelor's degree, who comprise 34% of our sample vs. 30% of the CES sample. Those with a graduate degree comprise 15% of our sample vs. 18% of the CES.[18]

Finally, we compare the industry and occupation coverage in our sample to the 2017-2022 ACS sample in the relevant states (see Appendix Tables A2 and A3). Our merged sample over-represents the finance, insurance, and information sectors; it under-represents educational services and healthcare. It also over-represents those in management occupations. These differences become less pronounced upon re-weighting the ACS sample to match the educational attainment of our analysis sample, but they remain substantial.

While our sample is skewed towards individuals who are more likely to use a professional social network and more likely to be registered to vote, it does not include exclusively white-collar workers: the top employer in our sample (as in the United States) is Walmart.

3 Sorting and segregation in the labor market

This paper studies both *sorting*, the systematic allocation of Democrats and Republicans to different groups, and *segregation*, the degree to which Democrats and Republicans encounter each other in a setting such as the workplace or the neighborhood. In this section, we document the distinctive labor market choices associated with each party, then we show that these different paths lead to substantial segregation at the workplace.

3.1 Partisan pathways

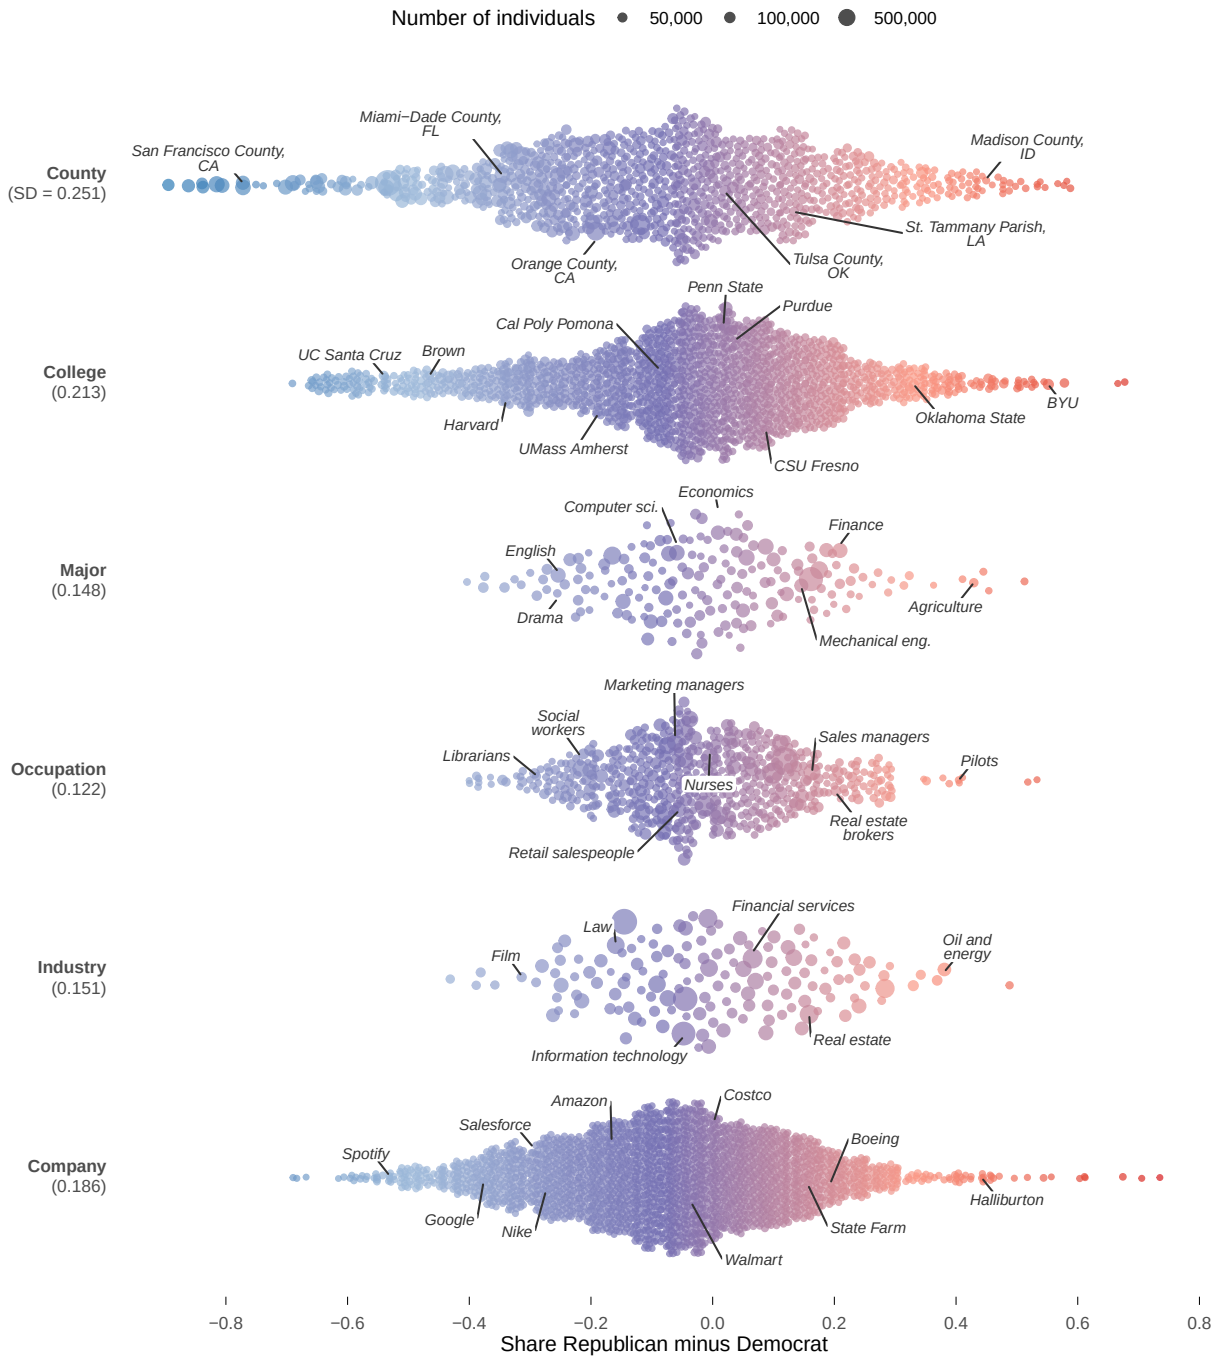
We first describe partisan differences at five key junctures in an individual's trajectory through the labor market: college, major, occupation, industry, and employer. To characterize the partisanship of group j (e.g., a particular college major or a particular industry), we use the partisan gap, g_j: the share of Republicans in the group minus the share of Democrats. This measure accounts for the 32% of our sample who are independents or registered with a third party, though all of our results are qualitatively similar if we use alternatives.[19] To measure sorting, we compute the standard deviation of the partisan gap across groups, e.g., across all college majors, weighted by the number of individuals in each group, which we refer to as σ_g.

We begin by visualizing the spread in the partisan gap across colleges, majors, occupations, industries, and employers in Figure 1. For comparison, we also show the distribution of g across counties (among individuals in our sample) at the top of the figure, and we include estimates of σ_g for each of these categories. The distributions are centered at zero to more easily assess their spread. We now discuss each stage in turn.

[18] Although 4% of our sample reports only a high school degree compared to 27% of the CES, about 41% of our sample does not list any college or high school on their profiles, and we think that many of these individuals are likely high school graduates.

[19] We prefer this approach over two common alternatives. First, we could use the share of Republicans alone, but a group with a share of Republicans equal to 0.5 could be 50% Republicans, 50% independents, and 0% Democrats ($g = 0.5$) or 50% Republicans, 0% independents, and 50% Democrats ($g = 0$). Second, we could use the share of Republicans among partisans (the two-party share), but a group with a two-party share of Republicans equal to 1 could be 100% Republicans, 0% independents, and 0% Democrats ($g = 1$) or 1% Republicans, 99% independents, and 0% Democrats ($g = 0.01$). Thus, to summarize the patterns with a single measure, we characterize partisanship using the share of Republicans minus the share of Democrats among all individuals in a group, including independents.

Figure 1: Sorting Into Colleges, Majors, Occupations, Industries, and Companies



Notes: This figure shows the partisan gap – the share of Republicans minus the share of Democrats – in each county, college, major, occupation, industry, and company. Each distribution is centered at zero. College refers to the IPEDS identifier of the associate or bachelor's degree school; major refers to the detailed field of degree as recorded in the ACS; occupation refers to a 6-digit SOC code; industry refers to the industry as reported on the individual's profile; and company refers to the name of the employer that the individual lists on their profile. We show only groups with more than 1,000 individuals in our sample.

Colleges: The second row of Figure 1 shows the partisan gap across colleges in our sample, as well as examples of distinctively Democratic and Republican schools. The correlation between partisanship and educational attainment is well-documented (e.g., Zingher, 2022), and this trend is apparent in our sample: for example, 29.2% of Democrats have a post-graduate degree, but only 22.1% of Republicans do. But even among those who graduated from college, the *types* of schools that Democrats and Republicans attend are very different. The standard deviation of the partisan gap is $\sigma_g = 21.5$, roughly the same as the standard deviation across counties, 25.3 p.p. In other words, going to a college with a partisan gap one standard deviation more Republican – e.g., from the University of Massachusetts, Amherst (−27.4 p.p.) to Penn State (−6.4 p.p.) – is like going from Orange County, California (−2.0 p.p.) to Tulsa County, Oklahoma (19.3 p.p.).

This variation is not due to idiosyncrasies of small colleges. Among the 100 colleges with the most graduates in our sample (with an average size of about 66,000 individuals), the most Democratic, U.C. Santa Cruz, has $g = -52.9$ p.p., and the most Republican, Brigham Young University, Idaho, has $g = 58.9$ p.p. Of course, the variation in both colleges and counties is due mostly to the partisan geography of the U.S. But we can further show that these differences cannot be explained just by whether these colleges are in red or blue states, or by the age, race, or gender composition of their graduates. In Appendix Figure A1, we present estimates of g for the most Republican and Democratic among this group of large colleges, controlling for each individual's current location and interactions of birth year, race, and gender, and show that these large gaps persist. And, in Appendix Table A5, we compute the standard deviation across colleges after residualizing partisanship on these controls, showing that one third of the variance remains unexplained.

Beyond location, college partisanship is clearly correlated with selectivity and prestige: a number of elite schools, including Harvard, Columbia, and Cornell, make the list of the most Democratic colleges, while the most Republican colleges do not include any. In fact, the share of Democrats among Ivy-Plus graduates in our sample is 57.7%, compared to the share among all other college graduates of 41.4%.[20] In Appendix Figure A2, we show the partisan gap for graduates of the top 30 schools in the U.S. News and World Report college rankings, including these Ivy-Plus schools.[21]

Majors, occupations, and industries: Next, we show that Republicans and Democrats make different decisions about what to study in college and about which occupations and industries to enter afterwards. Like we did for colleges, we show in Figure 1 the spread of the partisan gap across these groups and highlight distinctive examples in each category. We find that σ_g is 14.8, 12.2, and 15.1 p.p. for majors, occupations, and industries. This is substantially less than the variation across colleges and counties, but it is striking because these are large groups that cut across geographies: the average industry, for example, has about 159,000 individuals in our sample. In Appendix Figure A3, we show the most partisan majors, occupations, and industries (among the 50 largest in each group) controlling for location and demographics.

[20] The 12 Ivy-Plus schools are Brown, Columbia, Cornell, Dartmouth, Harvard, Princeton, the University of Pennsylvania, Yale, the University of Chicago, Duke, MIT, and Stanford.

[21] To examine college selectivity and partisanship more systematically, we regress two college characteristics – the average SAT score among admitted students and the admission rate – on the partisan gap at each school. These estimates (shown in Appendix Table A6) indicate that colleges with a partisan gap 100 p.p. more Republican have an admission rate 42.5 p.p. (2.1 standard deviations) higher and an average SAT score 163 points (1.2 standard deviations) lower, on average. The patterns are even stronger if we measure the partisan gap after accounting for CZ and demographic fixed effects. Although these are correlations, Firoozi (2023) shows that admission to selective University of California campuses reduces Republican registration and increases registration as an independent or a Democrat.

Among the largest majors, the difference between the most Republican – Finance – and the most Democratic – Drama and Theater – is almost 60 p.p. Overall, Republicans pursue majors in business, finance, engineering, and technology, while Democrats tend to study the arts, social sciences, and humanities. Among the largest occupations, the most Republican are real estate brokers, chief executives, and financial managers, and the most Democratic are writers, social workers, professors, and counselors. Again, there is more than a 40 p.p. gap between the most Democratic occupation and the most Republican. Finally, the most Republican industries include oil and energy, construction, and the military, while the most Democratic are mental health care, research, and entertainment. We find an even larger gap between the distinctively partisan industries than we do for occupations and majors: the difference between mental health care and oil and energy, for example, is more than 60 p.p.[22]

Companies: Lastly, we demonstrate that individual employers have distinctive partisan leans. Figure 1 shows that the standard deviation of partisanship across companies with more than 1,000 employees in our sample is $\sigma_g = 18.6$ p.p. Even among these large employers, we estimate sizeable gaps in partisanship: for example, Boeing has $g = 22$ p.p., while Google has $g = -35$ p.p. In Appendix Figure A4, we show estimates of partisanship for the largest companies, with a mean size of about 20,000 employees in our sample. These generally replicate the industry patterns seen before – defense companies, banks, and financial firms lean Republican, while technology companies lean Democratic. But we also present estimates that control for industry (in addition to location and demographics), which reveal some surprising patterns: Costco Wholesale, for example, rises from the 23rd to the 7th most Republican, while Google, the 4th most Democratic employer in the raw estimates, is ranked 24th with controls. This suggests that there are companies that attract Democrats and Republicans even within each labor market area and within each industry.

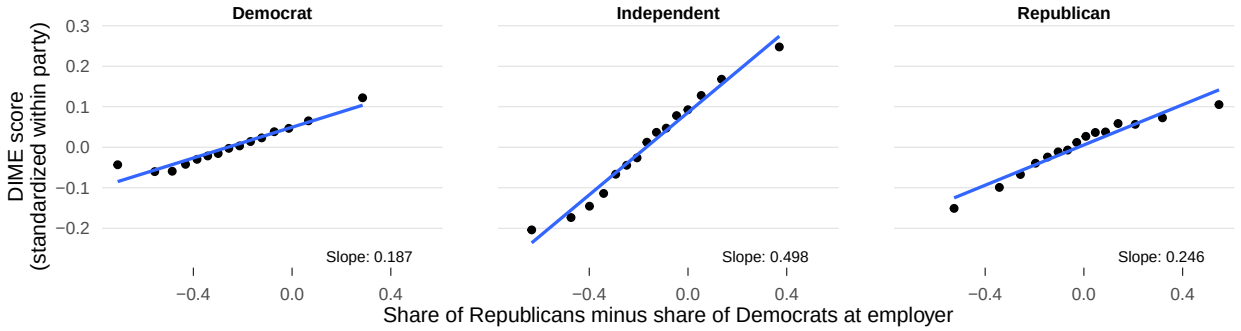
For the subset of individuals we link to campaign contributions, we can also examine a measure of *within*-party sorting. Figure 2 shows the relationship between the individual's DIME score, standardized within party, and the share of Republicans minus the share of Democrats among other individuals with the same employer. The Republicans at companies with a higher share of Republican coworkers are more ideologically extreme than their co-partisans at companies with fewer Republicans, and the same for Democrats. For Republicans, going from a company with all Democratic coworkers to one with all Republican coworkers is associated with a 0.50 standard deviation change in the DIME score, while for Democrats, it is 0.37. However, for independents or those registered with a third party, an individual's DIME score correlates even more strongly with their employer's partisanship. Going from an employer with all Democratic coworkers to one with all Republican coworkers is associated with a 1.00 standard deviation change in the DIME measure of ideology for this group. This pattern suggests that some independents, though nominally not registered with a party, have an ideology that leads them to select partisan workplaces.[23] In Section 4, we assess the magnitude of selection on "latent" or unobserved ideology correlated with partisanship.

Finally, we show that each of these stages represents a *distinct* form of sorting, not reducible to earlier decisions. It is possible, for example, that differences in major choice explain the occupations

[22]The industry patterns are broadly similar to the estimates in Bonica (2014), which rely on donations to political campaigns rather than voter registration to measure the ideological orientation of occupations and industries. Our approach measures partisan lean for a broader set of workers, not just those who donate, and allows for individual-level demographic controls. While we also find, for example, that academia leans to the left and construction to the right, our estimates indicate that the legal profession is distinctly Democratic, while Bonica (2014) finds that, according to donation behavior, it is more mixed.

[23]This is especially likely to be true among the group of independents who make campaign contributions.

Figure 2: Within-Party Sorting Across Employers



Notes: These binned scatter plots show the relationship between the DIME measure of political ideology (where greater values indicate more conservative ideology) on the vertical axis and partisan gap (share of Republicans minus share of Democrats) on the horizontal axis. DIME scores are standardized within all individuals in our sample registered with the given party. We include fixed effects for the individual's CZ and interactions of year of birth, gender, and race and exclude observations where the share of Democratic or Republican coworkers is zero or one.

that workers subsequently choose – i.e., that real estate brokers are Republicans because Republicans tend to study business. Table A5 shows how σ_g changes when, for example, we account for the choice of college and major in measuring sorting across occupations. These predictors are important: major choice drastically curtails the set of occupations that an individual is likely to work in, and σ_g drops to 6.7 p.p, from 8.5 p.p. with CZ and demographic predictors. Still, about 60% of the variance remains. Similarly, relative to estimates with CZ and demographics, about 40% of the variance across industries remains even after controlling for college, major, and occupation. Lastly, even after accounting for college, major, occupation, and industry, about one third of the variation in partisanship across employers is unexplained. These results indicate that politics predict an individual's decisions at each stage of their path through the labor market.

3.2 Measuring segregation

The previous section shows that partisan workers take different paths through the labor market. Now, we ask what this implies for the degree of political segregation at the workplace: how likely is a Democrat or a Republican to have a coworker who shares their party, and how does this compare to other commonly studied forms of segregation?

Workplace segregation could arise because labor markets are local and Democrats and Republicans live in different regions of the United States. It could also arise because, for example, young people tend to work with other young people, and young people tend to be Democrats. Thus, we want to study segregation conditional on covariates such as location and age. We adapt the approach in Åslund and Skans (2009), measuring segregation with the *overexposure ratio*, which is the probability that a Democrat or Republican's coworker shares their partisanship beyond that expected by observable characteristics.

This statistic has two components. We first compute *own-party exposure*, i.e., the exposure of a Republican to Republican coworkers at the workplace or the exposure of a Democrat to Democratic coworkers. This is a leave-out measure in that we measure partisans' exposure to partisan coworkers,

not including the focal employee themselves.[24]

Index individuals by i and their workplace by $w(i)$. We allow partisanship to take one of three values: Democrat ($D_i = 1$), Republican ($R_i = 1$), and independents or members of another party ($D_i = R_i = 0$). We compute own-party exposure only for Democrats and Republicans, but we include all individuals when calculating N_w, the number of workers at workplace w. Then, the share of individual i's coworkers who share their party is:

$$E_i = \frac{\left(\sum_{j:w(j)=w(i)} D_i D_j + R_i R_j\right) - 1}{N_{w(i)} - 1}.$$

Next, we compute the average *predicted* exposure for each individual:

$$\mathcal{E}_i(\mathbf{X}) = \frac{\left(\sum_{j:w(j)=w(i)} D_i \hat{D}_j^{-w(i)}(\mathbf{X_j}) + R_i \hat{R}_j^{-w(i)}(\mathbf{X_j})\right) - 1}{N_{w(i)} - 1}.$$

where $\hat{D}_i^{-w(i)}(\mathbf{X_i})$ and $\hat{R}_i^{-w(i)}(\mathbf{X_i})$ are the predicted values from regressions of the party indicators D_i and R_i on covariates $\mathbf{X}_i$ for all individuals at workplaces $w \neq w(i)$. In our baseline measure without controls, $\mathbf{X}_i = 1$, thus $\hat{D}_i^{-w(i)}(\mathbf{X_i})$ and $\hat{R}_i^{-w(i)}(\mathbf{X_i})$ are the population shares of Democrats and Republicans at all workplaces other than i's. When $\mathbf{X}_i$ includes just CZ fixed effects, then $\hat{D}_i^{-w(i)}(\mathbf{X_i})$ and $\hat{R}_i^{-w(i)}(\mathbf{X_i})$ are the shares of Democrats and Republicans among all individuals in i's CZ at workplaces other than i's. When we add other covariates to $\mathbf{X}_i$, individuals working at places where the characteristics of their coworkers predict that those coworkers are likely to be Republicans – say, because they are older, or went to a conservative college – will have higher predicted exposure to Republicans. Crucially, these characteristics are measured for each of i's coworkers, not individual i themselves.[25]

In practice, estimating these regressions leaving one workplace out at a time is computationally infeasible. To make the procedure tractable, we assign each workplace to one of 10 folds, $f(w(i)) \in \{1, \ldots, 10\}$, and compute the predictions $\hat{D}_i^{-f(w(i))}(\mathbf{X_i})$ and $\hat{R}_i^{-f(w(i))}(\mathbf{X_i})$ leaving one fold out at a time.

Our estimand is the population ratio of exposure and predicted exposure, or the *overexposure ratio*, $\mathbb{E}[S_i(\mathbf{X})] = \mathbb{E}[E_i/\mathcal{E}_i(\mathbf{X})]$, which measures the relative probability that a Republican or Democrat has a co-partisan coworker given the distributions of Republicans and Democrats in the population according to the characteristics $\mathbf{X}$. When $\mathbb{E}[S_i] = 1$, a partisan is no more likely to work with a co-partisan than would be expected from the population shares. If partisans were perfectly segregated by industry, then while $\mathbb{E}[S_i] > 1$, segregation conditional on industry fixed effects $\mathbf{I}$ would be $\mathbb{E}[S_i(\mathbf{I})] = 1$.

[24] Most workplaces are small, and many (e.g., Carrington and Troske, 1997; Ransom, 2000; d'Haultfoeuille and Rathelot, 2017; Gentzkow and Shapiro, 2011; Gentzkow et al., 2019) have pointed out that naive segregation indices are subject to a small-sample bias. We address this by computing the partisan share of each individual's *coworkers*, not the share at each workplace.

[25] For example, in analyzing the contribution of gender to workplace segregation, we do not compare two men and ask whether their coworkers share their party; rather, we compare two individuals whose *coworkers* have the same share of men and ask whether those coworkers share the focal worker's party.

We estimate the overexposure ratio using its sample analog $\hat{S}(\mathbf{X})$, and we compute standard errors using 50 iterations of a nonparametric bootstrap: we resample workplaces, predict partisanship using 10-fold leave-out regressions, and re-compute $\hat{S}$.

Our preferred measure is S_{cz}, which compares the average realized own-party exposure to the predicted own-party exposure based on commuting zone fixed effects.[26] As we have discussed, it is not surprising that Democrats in Los Angeles are exposed to more Democrats at work than would be expected based on national shares. But additional demographic predictors are part of the segregation we study – for example, the fact that Democrats work with more highly-educated workers (who are more likely to be Democrats) is an essential *feature* of political segregation, much like the difference in the occupations that men and women choose is an essential feature of gender segregation at work. We report estimates with a larger set of covariates to examine how much these correlates of partisanship can explain the segregation we observe.

Finally, to operationalize this measure, we need to define whom we consider an individual's coworkers. Ideally, we would observe not just the company that employs an individual, but the establishment at which they work. Since this is not possible in the LinkedIn data, we define workplaces as a company within a commuting zone.[27] This means we we will regard individuals who live in the same CZ but work at different establishments of the same company as coworkers.[28]

3.3 Estimates of workplace partisan segregation

Using this approach, we find substantial partisan segregation that is not attributable to political geography and cannot be explained by detailed covariates. Figure 3 shows the partisan overexposure ratio and explores how it changes when we account for observable measures. At baseline, $S = 1.200$, meaning that the coworker of a Democrat or Republican is about 20% more likely to share their partisanship than a random individual from our sample.

We first predict exposure using CZ fixed effects. The overexposure ratio drops considerably, to 1.100, but remains substantial, indicating that partisan segregation is 50% explained by CZ-level geography. This is our preferred measure, as it accounts for the shares of Democrats and Republicans among people who could reasonably share an individual's workplace because they are part of the same (geographic) labor market.

Next, we add covariates to investigate how much of the residual segregation we can explain with observables. When we predict partisanship using both CZ and fixed effects for the interaction of age, gender, and race, S drops to 1.073. An observer who knows these characteristics for each person – and knows their relationship with party affiliation – would still find that an individual's coworkers are about 7.3% more likely to share their party than expected.

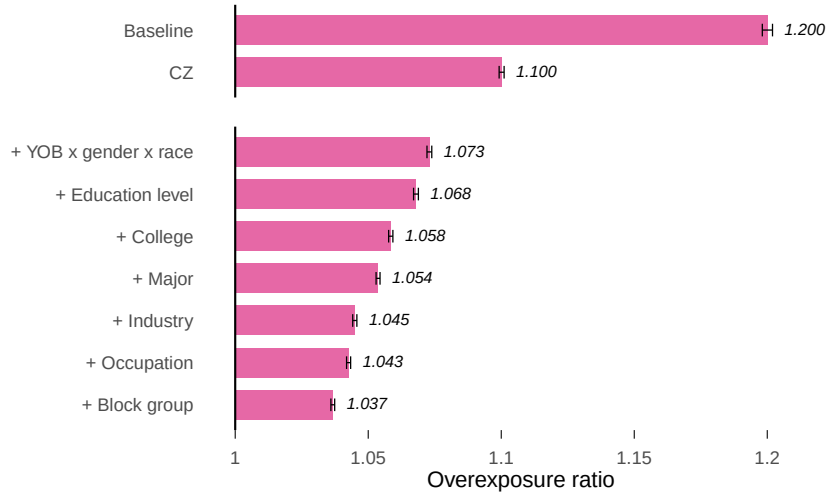
We then add various other predictors of partisanship: the highest level of education we observe on a user's profile, then the exact college they attended, then their major. The overexposure ratio is further diminished, to 1.054, but remains clearly above one. Next, we add industry and occupation

[26] CZs are aggregations of counties based on commuting patterns, and we pick this geographic unit to account both for political geography and the local nature of labor markets. There are 708 CZs in the areas we study. On average, each of them contains 4.4 counties and has 49,000 people in our sample.

[27] We use the CZ recorded on LinkedIn in the relevant year – the location of the user's job, if they report it, or their current location if they do not.

[28] In our main analysis, we use the most recent company that an individual reported working for, so it is possible that two individuals whom we consider coworkers did not work for the same company at the same time. Our results are similar under alternative definitions, for example, using only an individual's position in 2022.

Figure 3: Estimates of Workplace Partisan Segregation

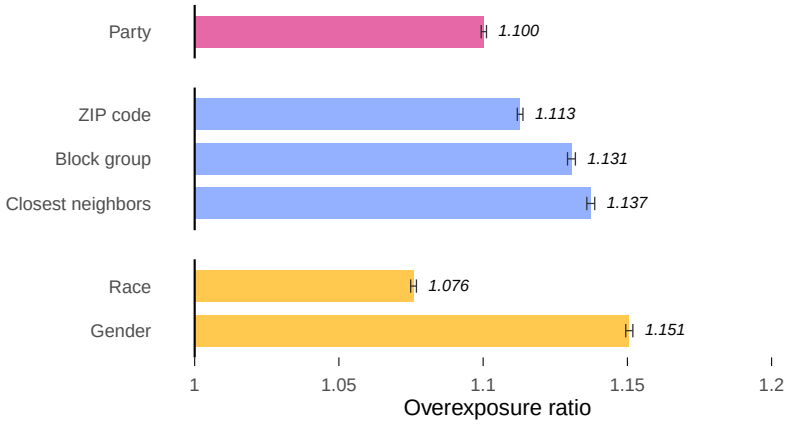


Notes: The figure shows the partisan overexposure ratio, which measures the probability that an Democrat or Republican's coworker shares their party affiliation, relative to the probability expected if their coworkers were drawn randomly from the population. We define a workplace as an employer by commuting zone. The first row shows the baseline estimates, which compare the shares of co-partisan coworkers to the shares of Democrats and Republicans in the full sample. The second row compares the shares of co-partisan coworkers to the expected shares of Democrats and Republicans based on CZ. Proceeding down the vertical axis, we cumulatively add covariates that predict the party of each individual; see the text for details. For all covariates, missing values are a separate category (i.e., we do not drop observations with missing values). We use the industry reported on the individual's profile, and occupation is the 5-digit SOC code associated with their job title. Standard errors are computed using the bootstrap with 50 replications, and horizontal bars are 95% confidence intervals.

fixed effects, and S drops to 1.045. Finally, we predict partisanship using the previous characteristics along with fixed effects for the exact block group where the individual registered to vote. Census block groups typically contain between 600 and 3,000 people, so this is a fine measure of the neighborhood where an individual lives. The overexposure ratio with the full set of predictors for expected exposure is $S = 1.037$. So, conditioning on the distribution of Democrats and Republicans across space at a very detailed level, along with their tendencies to attend particular schools, study particular fields, and work in particular occupations and industries, the probability that an individual's coworker shares their party is still about 4% higher than expected. Overall, we can explain about 82% of partisan overexposure with these covariates, or 63% of overexposure conditional on CZ.[29]

[29] Frake et al. (2024) also study workplace segregation using merged LinkedIn and voter file data. Their sample is different (they have a lower match rate but infer partisanship in states that do not record it) as well as their definition of coworkers (they use MSAs, while we use CZs). Moreover, their main segregation measure is an analog of the isolation index: the average share of a Democrat's coworkers who are Democrats minus the average share of a Republican's coworkers who are Democrats. They compute an isolation index of 14.8 p.p., and when they add fixed effects for the *focal* worker's MSA, occupation, and industry, the estimate drops to 2.3 percentage points. Our segregation measure is symmetric – it averages over Democrats and Republicans – and to incorporate covariates, we predict the partisanship of the focal individual's *coworkers* and measure how their realized own-party exposure differs from this prediction.

Figure 4: Comparing Workplace Partisan Segregation to Neighborhood, Racial, and Gender Segregation



Notes: The figure compares the partisan overexposure ratio at the workplace (S_{cz}) to other forms of segregation. In blue, we show estimates of gender segregation between men and women and racial segregation between white and nonwhite individuals at the workplace, accounting for the shares of each group in each CZ. In yellow, we show estimates of partisan segregation at the ZIP code and Census block group levels, as well as the overexposure ratio among an individual's closest neighbors, again comparing to the share of partisans in each CZ. Standard errors are computed using the bootstrap with 50 replications, and horizontal bars are 95% confidence intervals.

3.3.1 Comparison to residential partisan segregation

Predicting each individual's partisanship using block group fixed effects leaves residual segregation, suggesting that workplace sorting is a distinct phenomenon, not merely a result of residential sorting. Now, we compare the two explicitly.

We first compute S_{cz} using two units commonly used to study residential segregation: the ZIP code and the Census block group. Figure 4 shows that segregation at the workplace (1.100) is comparable to segregation at the ZIP code (1.113) but somewhat less than segregation at the block group (1.131). In Appendix Figure A6, we show that partisan segregation across neighborhoods is explained substantially by the demographics of Democrats and Republicans. After predicting partisanship using age, race, and gender, the overexposure ratio at the workplace, ZIP code, and block group are very similar.

However, while these units capture intuitive notions of the neighborhood, their average sizes are quite different, both from each other and from the workplace. Moreover, the *distributions* of unit sizes are quite different: while some individuals have many more coworkers than block group neighbors, some have far fewer.[30] Even holding fixed unit sizes, drawing boundaries differently could result in different estimates (the modifiable areal unit problem). To sidestep these issues, we adopt an approach inspired by Brown and Enos (2021), leveraging each voter's exact residential address.

[30]The average individual in our sample has 470 coworkers and 433 block group neighbors. However, there are many small workplaces and few large ones, while block groups are designed to be relatively similar in size. This means that at the median, an individual in our sample has just 36 coworkers and 314 block group neighbors.

Consider a Democrat or Republican i with $N_{w(i)} - 1$ coworkers. We rank i's neighbors (among those in our sample) by distance and compute the share of co-partisans among their $N_{w(i)} - 1$ closest neighbors, as well as the expected share of co-partisan neighbors based on their CZ. Then, we compute S_{cz}. This means that the distributions of the number of coworkers and the number of neighbors in this comparison are identical. We find that S_{cz} among neighbors is 1.137, higher than among coworkers. In Appendix Figure A5, we show further that this pattern is heterogeneous: while Democrats experience slightly lower levels of partisan overexposure among their coworkers than their neighbors, Republicans at small workplaces experience far higher overexposure among coworkers, while Republicans at large workplaces experience less.

On average, we find that partisan segregation at the workplace is similar to partisan segregation at the ZIP code, but less than segregation among one's closest neighbors.

3.3.2 Comparison to racial and gender segregation

Next, we benchmark workplace political segregation a different way, showing that it is more pronounced than racial segregation at the workplace, but less pronounced than gender segregation. We again compute the overexposure ratio, but instead of computing average own-party exposure, we compute own-race exposure for white and nonwhite individuals and own-gender exposure for men and women. These estimates account for covariates with the same approach, predicting each individual's race or gender with CZ fixed effects; we thus compare, for example, a white person's share of white coworkers to the share in each CZ.

These results are shown in Figure 4. First, we find that workplace partisan segregation ($S_{\text{cz}} = 1.100$) is greater than racial segregation (1.076). These averages, however, mask a key difference between partisan and racial segregation: the latter is highly asymmetric. A relatively small group of nonwhite individuals experiences very high own-race exposure, while a larger group of white individuals experiences own-race exposure that is not so different from expectation. This is different from partisan segregation, where Democrats and Republicans both experience higher-than-expected own-party exposure.

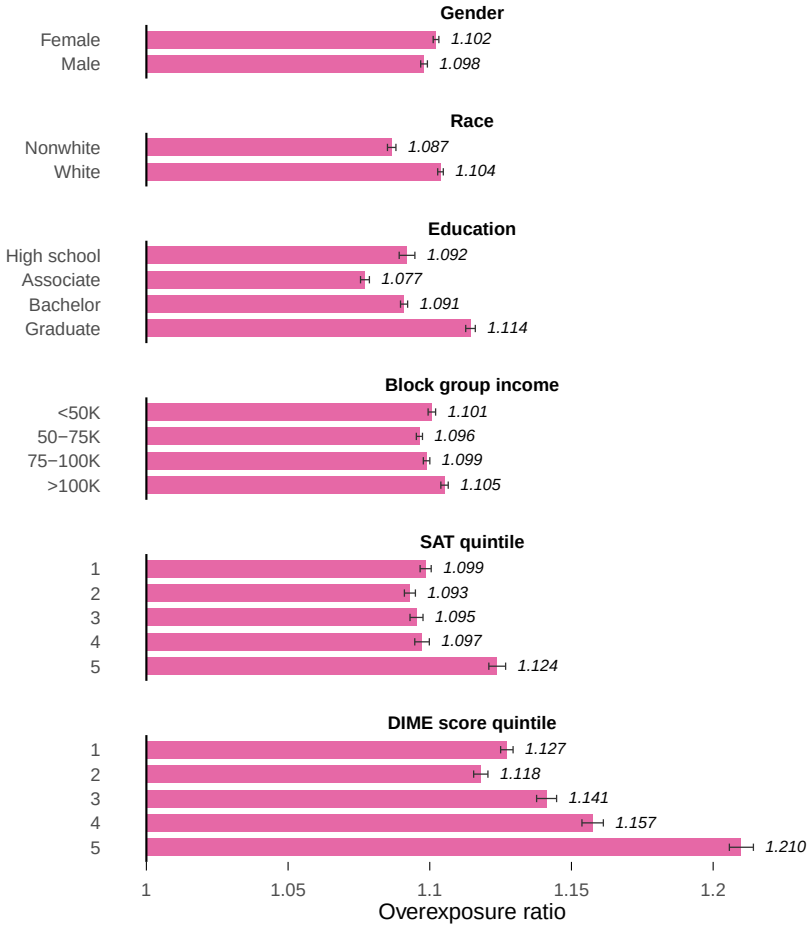
We also find that workplaces are more segregated by gender ($S_{\text{cz}} = 1.151$) than by partisanship. In Appendix Figure A7, we show that this is due to the familiar phenomenon of occupational segregation: using industry and occupation fixed effects to predict coworker gender shrinks the overexposure ratio to 1.048, less than partisan segregation with the same controls.

3.4 Who experiences partisan segregation at work?

Lastly, we ask *who* experiences the greatest workplace segregation, finding that partisan segregation is more correlated with educational attainment than with income.

Figure 5 shows S_{cz} computed separately by groups like race and education. Men and women experience approximately the same level of segregation, but white people experience slightly more (1.104) than nonwhite people (1.087). We find more heterogeneity by educational attainment. Workers with a graduate degree experience more segregation (1.114) compared to others. While those with a bachelor's degree experience a comparable level of segregation to high school graduates, we see a similar pattern when we turn to a different measure of education: the average SAT score at the college that the individual graduated from. Graduates in the top quintile of the college SAT score distribution have an overexposure ratio of 1.124, while even those in the fourth quintile experience $S_{\text{cz}} = 1.097$ only. Turning to income, we split the sample by four levels of block-group

Figure 5: Heterogeneity in Workplace Partisan Segregation



Notes: The figure shows the own-party overexposure ratio computed separately for individuals by gender, race, education, block-group level median household income, quintiles of SAT score associated with the college they attended, and quintiles of the DIME score. To compute overexposure, we compare actual own-party exposure to exposure predicted by coworkers' CZs. Block group income is from the 2016-2021 ACS. Standard errors are computed using the bootstrap with 50 replications, and horizontal bars are 95% confidence intervals.

level median household income – less than $50,000, $50,000 to $75,000, $75,000 to $100,000, and $100,000 or more – which correspond approximately to quartiles of this distribution in our sample. Segregation does not change very much across the income distribution – from $S_{cz} = 1.101$ for the lowest income group to 1.105 for the highest.[31]

Finally, we split the sample by quintiles of the DIME score, a continuous measure of ideology.

[31] Median household income measured at the block group is an inexact measure of individual income, while educational attainment is measured at the individual level. Thus, in Appendix Figure A8, we perform an apples-to-apples comparison of education and income by computing the overexposure ratio separately for deciles of household income and share of individuals with a graduate degree measured at the Census tract level. We reproduce the same pattern: individuals living in more-educated places experience greater partisan segregation at work, while the relationship between overexposure and income is far less striking.

First, the overexposure ratio for *all* individuals with a DIME score – i.e., those who have made political contributions – is above the average in our sample. But we also see evidence of asymmetric polarization: while the most liberal individuals experience slightly more segregation than the second quintile, the most conservative individuals have an overexposure ratio of $S_{\text{cz}} = 1.210$.

In summary, those with graduate degrees, those who attended selective colleges, and the most conservative individuals experience the greatest workplace partisan segregation.

4 The effect of jobs on political behavior

There are two categories of explanations for the extensive political segregation that we have documented. One is that jobs shape political behavior, and workers conform to the politics of their workplace. Work can be a powerful source of meaning and identity (Cassar and Meier, 2018; Akerlof and Kranton, 2005), and employed Americans spend a substantial fraction of their days at the workplace, where their views might be shaped by coworkers and managers or by the nature of the job that they do. It is possible that the workplace causes individuals to form political attachments and register as Democrats or Republicans, or even causes registered Democrats and Republicans to change their party or unaffiliate. If workers tend to register with the majority party among their coworkers, then jobs would contribute to partisan segregation.

Alternatively, the homophily we see in the labor market may not be a consequence of jobs, but largely a matter of selection. In this world, workers end up in jobs that suit their politics. Even if a worker is not officially registered as a Republican or Democrat when they start a new job, the workplace does not *affect* their decision to later register with a party, as they would have done so anyways.

Of course, both explanations could be partially true. In this section, we distinguish between them by focusing on people who switch jobs to identify the causal effect of the workplace on party affiliation.

4.1 Empirical strategy

Our empirical strategy to identify this causal effect isolates quasi-random variation in *when* a worker is exposed to a partisan workplace. We combine elements of the event-study design used by Perez-Truglia (2018) to study whether individuals adopt the political donation behavior of their neighbors and a movers design used to study childhood exposure effects (Chetty and Hendren, 2018; Chetty et al., 2018; Brown et al., 2023).

Intuitively, we compare two people, Individual 1 and Individual 2, who made similar moves from workplace A to workplace B, except Individual 1 made the move before year t and Individual 2 made the move after year t.[32] Then, we regress each individual's probability of registering with the Republican or Democratic party in year t on the partisanship of their coworkers at workplace A and workplace B. If the workplace has a causal effect on political behavior, then Individual 1's partisanship in year t should look "more like" their coworkers at workplace B than Individual 2, since they have already had an opportunity to be influenced by that workplace. That Individual 2 also looks like their coworkers at workplace B even though they haven't moved there yet captures selection – unobservable characteristics that make certain types of workers move to certain types of firms.

[32] We continue to define a workplace as a company within a commuting zone.

Outcomes: The outcome in our regressions is Y_{ipt}, an indicator that is 1 if, in year t, individual i's latest party affiliation is with party $p \in \{D, R\}$. It is 0 if their latest registration is with another party or if they are unregistered but living in a state that records party affiliation (and are at least 18 years old). Our main specification "stacks" two observations (for the Democratic and Republican parties) per individual, and we cluster standard errors at the individual level, though we examine affiliation with each party separately in robustness checks.

Our primary outcome is partisanship in 2016, $Y_{ip,2016}$, and we compare people moving before 2016 to people moving after. In principle, we could measure party affiliation in any year over our sample period. However, individuals update their registration much more frequently in election years than in others, and more frequently still in presidential election years. There are three presidential elections in our sample period: 2012, 2016, and 2020. We only observe moves for one year after the 2020 election, and we do not observe any moves before 2012. For this reason, we focus on 2016 in our main results, although we show that estimates for other election years are similar.

Sample: We split the sample into two mutually exclusive groups: one-time movers, individuals who changed workplaces exactly once between 2012 and 2021; and all others (group S), individuals who did not change workplaces or changed more than once between 2012 and 2021.

Appendix Table A7 shows summary statistics for one-time movers and splits the sample into those who moved before and after 2016. The two are similar on most observable characteristics, including their average birth years, shares of partisans in 2016 and 2020, and shares of partisans at their origin and destination workplaces. However, we find that pre-2016 movers tend to move from smaller to larger workplaces, and post-2016 movers tend to move from larger to smaller workplaces. We also find that the pre-2016 movers are somewhat more likely to have been registered and voted in 2012 than the post-2016 movers. Therefore, our baseline regression specification includes fixed effects for the quintile of the origin workplace size interacted with the quintile of the destination workplace size (measured at the individual level), as well as controls for registration and turnout in 2012.

Exposure: We characterize workplaces by the share of Democrats or Republicans among group S, fixing individuals' partisanship in 2012: we take the average of $Y_{ip,2012}$ at each workplace j among this group over the entire sample period. Let $w(i,t)$ index the workplace where an individual i works at time t. Then

$$\hat{Y}_{jp} = \frac{\sum_{i \in S} \sum_t Y_{ip,2012} \mathbb{1}(w(i,t) = j)}{\sum_{i \in S} \sum_t \mathbb{1}(w(i,t) = j)}.$$

To ensure we have a reliable measure of partisan exposure, we limit our analysis to workplaces where we observe more than 10 people.[33]

4.1.1 Identification

Let $Y_{ip,2016}$ be the individual i's affiliation with party p in 2016, $\hat{Y}_{dp}$ be the partisanship of the destination workplace, and $\hat{Y}_{op}$ be the partisanship of the origin. If the one-time movers were randomly assigned to workplaces, then we would run the following regression on the sample of workers who moved before 2016:

[33] We compute the split-sample reliability of this measure by dividing S randomly at the individual level to obtain two independent estimates, $\hat{Y}^{(1)}_{jp}$ and $\hat{Y}^{(2)}_{jp}$. The reliability of the exposure measure is 0.443, but it is 0.653 for workplaces with more than 10 people and 0.848 for workplaces with more than 50 people. In a robustness check, we use a split-sample instrument to account for measurement error in exposure.

$$Y_{ip,2016} = \beta \hat{Y}_{dp} + \gamma \hat{Y}_{op} + \varepsilon_{ip}$$

and β would recover the causal effect of spending between 1 and 3 years (since moves took place between 2013 and 2015) at a workplace where the share of Republicans or share of Democrats among coworkers is 100 p.p. higher. However, workers are not randomly assigned, and the types of people who move to workplaces with more Republicans are likely different from the types of people who move to workplaces with more Democrats. In other words, the error term is correlated with the destination workplace partisanship, so the observational regression coefficient is

$$b_{\text{bef}} = \beta + \frac{\text{Cov}(\hat{Y}_{dp}, \epsilon_i)}{\text{Var}(\hat{Y}_{dp})} = \beta + \delta_{\text{bef}}$$

where δ_{bef} captures a selection effect. However, we can identify β by using the sample of one-time movers who moved *after* 2016, i.e., after the outcome was measured. To do so, we make two assumptions:

Assumption 1 (No anticipation): There is no causal effect of the workplace on an individual before they start to work there.[34] Under this assumption, if we run the same regression on the sample of people who moved after 2016, we recover

$$b_{\text{aft}} = \delta_{\text{aft}}.$$

We expect $\delta_{\text{aft}} > 0$, reflecting the fact that an individual's partisanship in 2016 is correlated with whether they later move to a more or less Republican workplace between 2017 and 2021.

Assumption 2 (Constant selection): The degree of selection does not vary with whether the individual chooses to change jobs before or after 2016, or

$$\delta_{\text{bef}} = \delta_{\text{aft}}.$$

Under these two assumptions, the causal effect of the workplace on partisanship is identified by the difference between the observational regression coefficients on $\hat{Y}_{dp}$ for movers before versus after 2016:

$$\beta = b_{\text{bef}} - b_{\text{aft}}.$$

Importantly, this strategy identifies the bundled effect of the *workplace* on political behavior, not necessarily the effect of partisan *coworkers* alone. Our estimates also capture changes in a worker's partisan identity that could arise from features that are correlated with the partisanship of the workplace. For example, if unionized workplaces make workers more likely to register as Democrats, then this effect would be included in our estimates. In addition, our estimates include changes in partisanship resulting from the nature of the job itself. For example, if starting a job at a mining company leads someone to reconsider their views on energy policy, then this effect would also be part of our estimates.[35]

[34] This could be violated if, for example, workers change their political affiliation to match a job that they have not yet started.

[35] The interpretation of our estimates is similar to Brown et al. (2023), where the partisanship of "permanent

4.2 Main results

We now use this empirical strategy to show that the workplace has a causal effect on party registration among those who were previously neither Democrat nor Republican, but no detectable effect among committed partisans.

4.2.1 Graphical analysis

Before turning to a more parsimonious regression, we show results from a less parametric specification to check the assumptions underlying our identification strategy. We first plot the correlation between worker partisanship in 2016 and the partisanship of their coworkers at the destination workplace separately for moves taking place each year. Let M_i^y be a set of year-of-move indicators for $y \in \{2013 \dots 2021\}$. We estimate

$$Y_{ip,2016} = \sum_{y=2013}^{2021} \beta_y (M_i^y \times \hat{Y}_{dp}) + \gamma \hat{Y}_{op} + \psi_p X_i + \phi_p (J_o \times J_d) + \varepsilon_{ip} \tag{1}$$

where X_i includes fixed effects for the individual's registration status, party affiliation, and turnout in 2012, as well as the interaction of the individual's age, race, and gender; and J_o and J_d include company CZ, industry, quintile of workplace size, and SOC code of the job at the origin and destination workplaces, respectively.

We split the regression sample into two groups based on their party affiliation in 2012: those who were unregistered, independent, or registered with a third party (52.7% of the sample),[36] and those who were registered Democrats or Republicans in 2012 (47.3%). We think of the former group – who were not already committed to the Democratic or Republican Party – as those who are susceptible to the influence of their workplace. Indeed, we see very different patterns for these two groups.

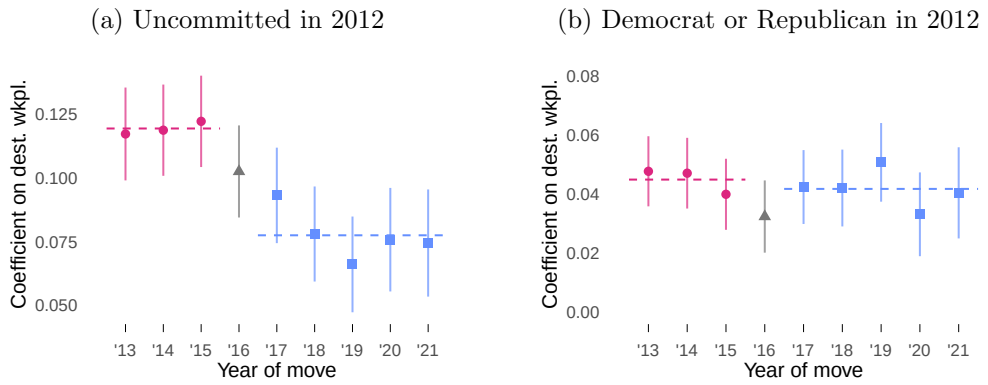
Figure 6 shows the β_y coefficients from estimating Equation 1 on the set of one-time movers, separately for uncommitted voters and committed partisans. In Panel (a), the coefficients for moves before 2016 (shown in magenta) are consistently higher than the coefficients for moves after 2016 (in blue). This drop occurs relatively sharply: with the exception of 2017, there appears to be an additional $\beta > 0$ term that shifts the level of the pre-2016 b coefficients relative to the post-2016 coefficients. The higher level of the 2017 coefficient could be due to anticipation, or to measurement error in the timing of moves. In either case, Assumption 1 appears to be satisfied for moves at least one year before we measure partisanship, and including 2017 in our estimate of the selection term will lead to a conservative estimate of the causal effect. Second, the relatively flat pattern in the post-2016 coefficients suggests that the selection effect is constant, consistent with Assumption 2. This selection effect is large: the partial correlation between the destination workplace partisanship and the worker's partisanship in 2016 is similar in magnitude even for workers who moved *after* 2016. We return to an interpretation of these magnitudes below.

In Panel (b), the levels of the coefficients for moves before versus after 2016 do not exhibit a noticeable difference, indicating that there is no detectable causal effect of the workplace on partisanship

residents" of a neighborhood is used to characterize an individual's exposure to a more or less partisan environment, not because the causal effect of the neighborhood is necessarily due to interactions with those individuals.

[36] But who were over 18 years old and living in a state that records party, i.e., for whom we *could* have observed party registration.

Figure 6: Effect of the Workplace on Partisanship in 2016



Notes: The figure shows estimates corresponding to Eq. 1 of the relationship between the destination workplace partisanship (share of Republicans or Democrats) and the worker's own partisanship measured in 2016. Each coefficient corresponds to individuals who moved in a particular year. The magenta coefficients are estimated for those who moved before partisanship is measured and the blue coefficients are for those who moved after. In Panel (a), we show estimates for those who were "uncommitted" (independent, unregistered, or members of a third party) in 2012, and in Panel (b), we show this for those who were registered Democrats or Republicans in 2012. Vertical bars are 95% confidence intervals computed using standard errors clustered by individual.

for those who were already registered Democrats or Republicans in 2012. The levels of the coefficients on both destination and origin workplace partisanship are smaller for this group because we control for their party affiliation in 2012, and it is rare for partisans to switch affiliation or to unaffiliate.

4.2.2 Parametric specification

To maximize power, our preferred specification groups all pre-2016 moves together and post-2016 moves together. Let $K_i^{2016} = \mathbb{1}(M_i < 2016)$ be an indicator that is 1 if the individual moved before 2016, 0 if they moved after 2016, and missing if the individual moved in 2016. We estimate

$$Y_{ip,2016} = \beta(K_i^{2016} \times \hat{Y}_{dp}) + \delta\hat{Y}_{dp} + \gamma\hat{Y}_{op} + \psi_p X_i + \phi_p(J_o \times J_d) + \varepsilon_{ip}. \tag{2}$$

The coefficient of interest is β, which corresponds to the difference in the correlation between worker and destination workplace partisanship for workers moving before versus after 2016. Under the assumptions described previously, this identifies the causal effect of the workplace on partisanship. The coefficient δ corresponds to the selection term.

Results from estimating Equation 2 are in Table 2. Column (1) shows the estimates for the sample of uncommitted voters, where we find $\hat{\beta} = 0.050$. The magnitude of this estimate implies that going from a workplace where none of one's coworkers are members of party p to one where all of them are makes an uncommitted individual about 5.0 p.p. more likely to register with that party. We also find, consistent with Figure 6, that $\hat{\alpha} = 0.074$, which we interpret as selection. The correlation between worker and destination partisanship is positive and significant for those moving after 2016, but it is about 68% larger for movers before 2016. Put differently, about 40% of the observed relationship between workplace and worker partisanship – for workers spending one to three years at that workplace who were previously uncommitted – can be attributed to the causal effect of the workplace, with the rest due to selection.

Table 2: Effect of the Workplace on Partisanship in 2016

	Uncommitted			Committed	All
	(1)	(2)	(3)	(4)	(5)
Moved before × dest. exposure	0.050***	0.041***	0.035***	0.001	0.017***
	(0.006)	(0.006)	(0.012)	(0.004)	(0.004)
Dest. workplace exposure	0.074***	0.077***	0.077***	0.042***	0.061***
	(0.005)	(0.005)	(0.009)	(0.003)	(0.003)
Origin workplace exposure	0.102***	0.101***	0.093***	0.040***	0.070***
	(0.004)	(0.004)	(0.008)	(0.003)	(0.003)
Individual controls	✓	✓	✓	✓	✓
Origin × dest. controls	✓	✓	✓	✓	✓
Drop CZ switchers		✓			
Drop industry switchers			✓		
Observations	1,089,264	937,414	308,458	990,076	2,131,962
R^2	0.168	0.164	0.169	0.812	0.637
Dependent variable mean	0.133	0.131	0.130	0.472	0.295
Dependent variable std. dev.	0.340	0.338	0.336	0.499	0.456

Notes: This table shows results from estimating Equation 2. We restrict the sample to one-time movers between workplaces where we observe at least 10 individuals who are *not* one-time movers. Columns (1) through (3) include only individuals who were unregistered, independent, or registered with a third party in 2012. Columns (2) and (3) drop those who move between CZs and industries, respectively. Column (4) includes only those who were registered Democrats or Republicans in 2012, and column (5) includes the full sample. Standard errors are shown in parentheses and clustered by individual. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

In Column (4), we estimate the same regression on individuals who were already Democrats or Republicans in 2016, finding $\hat{\beta} = 0.001$. The 95% confidence interval for this estimate is $[-0.006, 0.009]$, allowing us to rule out even small effects. Column (5) shows coefficients for the full sample, where we find $\hat{\beta} = 0.017$ and $\hat{\alpha} = 0.061$, implying that 22% of the relationship between workplace and worker partisanship in the full sample is causal.

4.2.3 Robustness

Next, we show that these results are not driven by individuals who switch CZs or industries, that they are not specific to registration with one party or to the 2016 election, and that they are not significantly attenuated by noise in our measure of workplace partisanship. These robustness checks focus on the sample of uncommitted voters, for whom we estimate a nonzero causal effect.

Dropping CZ and industry switchers: We account for origin-CZ-by-destination-CZ fixed effects in our regression specification, so the hypothetical comparison we make is between two individuals moving between workplaces in the same two CZs. However, not all individuals switch CZs when they change jobs, and a plausible concern is that our results are driven by those who do. The neighborhood where an individual lives has an independent causal effect on their partisanship (Cantoni and Pons, 2022), so we might attribute changes in partisanship due to someone's new residential context to their new workplace. In column (2) of Table 2, we restrict the sample to people who do not change CZs when they move between workplaces. We find that the effect is about 80% as large as the baseline estimates.

Similarly, while we account for origin-industry-by-destination-industry fixed effects, it could be

that the effects are driven by those who make the more dramatic decision to switch industries, rather than those who change jobs within the same industry. While this would not invalidate our empirical strategy, it would imply that the effect of jobs on political behavior is driven more by large changes in the nature of the individual's job, rather than interactions with coworkers or workplace characteristics that vary within an industry. In column (3) of Table 2, we restrict the sample to people who do not change industries when they move between workplaces, finding that the effect is about 70% as large as the baseline estimates.

Effects on registration with each party: In Appendix Table A8, we show that the effect is not driven by registration with either party alone by estimating Equation 2 using each outcome ($Y_{iR,2016}$ and $Y_{iD,2016}$) separately. The estimates are similar, though they are somewhat larger for registration with the Republican Party (0.055) than with the Democratic Party (0.044).

Other elections: These estimates are also not specific to partisanship measured in 2016. In Appendix Table A9, we show results for partisanship measured in 2014, 2018, and 2020, finding similar estimates. As discussed previously, these are not our preferred outcomes for two reasons. First, they do not occur near the middle of our sample period, so we observe substantially more moves before or after the outcome is measured. Second, presidential elections are more salient than midterm elections, and it is unclear whether we should expect workers to update their party registration at the same rates in midterm versus presidential election years.[37]

In columns (5) and (6) of that table, we "stack" rows for the same individual observed in different years, pooling over these elections. The regression we estimate is, for election years t,

$$Y_{ipt} = \beta(K_i^y \times \hat{Y}_{dp}) + \delta_t \hat{Y}_{dp} + \gamma_t \hat{Y}_{op} + \psi_{pt} X_i + \phi_{pt}(J_o \times J_d) + \mathbb{1}(T_{ipt} = t) + \varepsilon_{ipt} \tag{3}$$

where $K_i^y = \mathbb{1}(M_i < t)$ is an indicator that is 1 if individual i moved before year t, zero if they moved after t, and missing if they moved in t. We include election year fixed effects $\mathbb{1}(T_{ipt} = t)$ and allow the coefficients to vary by election in order to account for the different levels of correlation between worker and workplace partisanship across the different election years.[38] We continue to cluster standard errors by individual. We find that $\hat{\beta}$ pooling over all four general election years is 0.048, and pooling over the two presidential elections in our sample is 0.050 – the same as the baseline estimate for 2016 alone.

Measurement error in exposure: Although our baseline estimates include only workplaces where we observe more than 10 people, noise in our measure of workplace partisan exposure could still attenuate the estimates. To account for this possibility, we use a split-sample instrumental variables approach following Chetty et al. (2018). We divide the sample of those who are *not* one-time movers (i.e., not in our regression sample) randomly into two groups, and compute partisanship of each workplace j separately for each group to obtain two independent estimates, $\hat{Y}_{jp}^{(1)}$ and $\hat{Y}_{jp}^{(2)}$. Then we estimate Eq. 2 replacing $\hat{Y}_{dp}$ with $\hat{Y}_{dp}^{(1)}$ and $\hat{Y}_{op}$ with $\hat{Y}_{op}^{(1)}$. We instrument for $\hat{Y}_{dp}^{(1)}$ with $\hat{Y}_{dp}^{(2)}$, and likewise for $\hat{Y}_{op}^{(1)}$.

These estimates are in Appendix Table A10. In column (1), we show that the coefficient using the

[37] In Appendix Figure A9, we show plots corresponding to Figure 6 for the other election years in our sample. We find similar patterns: larger estimates for moves made pre-election versus post-election, and relatively constant estimates for the latter.

[38] Specifically, the correlation between these two measures grows over time as individuals who were previously unregistered have successive opportunities to register with a party.

split-sample instrument on the baseline sample is somewhat higher (0.071 versus 0.050). However, when we include workplaces with fewer than 10 employees, the SSIV estimate in column (2) falls to 0.069 while the OLS estimate in column (4) falls to 0.023, consistent with noise attenuating the estimates when we include smaller workplaces. While the point estimates are larger, the SSIV estimates imply that about 25 percent of the relationship between worker and workplace partisanship for this group of uncommitted voters is causal.

5 The effect of politics on job choice

Despite strong political sorting in the labor market, the results in the previous section point to the limited role of the workplace in shaping political behavior in the short run – especially in changing the minds of those who already know where they stand. By contrast, selection is remarkably strong: even those who haven't yet registered as a Republican are likely to move to Republican-leaning firms, and vice versa for Democrats. But our estimates of selection do not tell us *who*, if anyone, is doing the selecting. It could be that workers have a preference for politically compatible workplaces. Alternatively, it could be that employers want to hire partisan workers, either explicitly because of their political views or based on attributes that are correlated with being a Democrat or a Republican. Or, it could be that politics do not enter worker or employer preferences at all, yet matching frictions (e.g., employers hiring from homogeneous alumni networks) generate segregation across workplaces.

All three of these forces likely play a role in the substantial selection that we observe, and therefore in generating segregation at the workplace. In this section, we focus on the first component – the extent to which *workers'* political preferences affect their job choices – leaving the other two components to future study.

How could worker preferences generate partisan segregation? One class of explanations in the tradition of Schelling (1971) highlights the role of social interactions: Democrats might explicitly prefer to mix with other Democrats than with Republicans, and vice versa. Even if these preferences are relatively weak, these models can lead to "tipping point" equilibria – own-gender preference can lead to occupations that are male-dominated (Pan, 2015), or own-race preference can lead to neighborhoods that are dominated by white people (Card et al., 2008). However, we see few examples of all-Democrat or all-Republican groups in Figure 1. Moreover, it may be difficult for workers to act on preferences over their colleagues' political views, which are less readily observed than race and gender and not reported in, e.g., diversity scorecards.[39]

A different class of explanations highlights how workers' preferences over non-monetary job characteristics can lead to sorting, an insight summarized by Rosen (2002): "markets sort buyers and sellers appropriately to take advantage of heterogeneous talents and tastes." In this framework, Democrats and Republicans do not necessarily value the partisanship of their coworkers explicitly, rather, their different preferences over job amenities lead them to sort into different workplaces. These preferences also generate pay differences between the jobs that Democrats tend to take and the jobs that Republicans tend to take. Indeed, we find suggestive evidence in the observational data of such differences.

[39] Indeed, diversity scores or disclosures do impact job selection decisions (Choi et al., 2023; LaViers and Sandvik, 2023).

5.1 Partisanship and pay

The partisan pathways that we document in Section 3 are not random – they are correlated with prestige and pay. The colleges with more Democrats tend to be more selective. But Republicans go from finance major to financial manager in the insurance industry, for example, while Democrats go from English major to writer at a public relations firm. These two careers have significantly different economic returns.

To generalize that example, we compare similar people who live in the same area and went to the same college and ask whether Democrats and Republicans tend to make choices that pay less or more. We regress a measure of income associated with a major, occupation, or industry on indicators for whether the individual is a Republican or a Democrat (so the estimates are relative to those who are neither), controlling for CZ, interactions of birth year, race, and gender, and the college that the individual attended.[40] Figure 7 displays these coefficient estimates. We find that Democrats choose majors with an average income 0.026 log points (about $1,400) lower than the majors that Republicans choose, occupations with wages 0.041 log points ($3,800) lower, and industries with pay about 0.043 ($3,600) lower.

These relationships suggest that Democrats and Republicans sort into jobs in a way that correlates with income, even conditional on detailed observable characteristics and a measure of education. One possibility is that Democrats and Republicans choose different jobs because they have different preferences for the canonical attributes studied by the literature on job amenities: fringe benefits, such as employer-provided health insurance, or conditions of work, such as the option to telecommute. It is possible that Democrats and Republicans value these amenities differently, and this could explain the pay gap between the two parties. For example, Democrats may prefer jobs that pay less but offer the option to telecommute.

We find little evidence for this in previously collected data. Maestas et al. (2023) measure workers' willingness to pay for better working conditions using the American Life Panel, and while their survey did not measure political behavior, the panel participants were asked about their party affiliation in an unrelated earlier survey. This enables us to reanalyze their data and test whether Democrats and Republicans differ in their WTP for the job attributes they study. We show these results in Appendix Figure A10, finding that there are no major differences between Republicans and Democrats in their WTP for the job attributes in their experiment.[41]

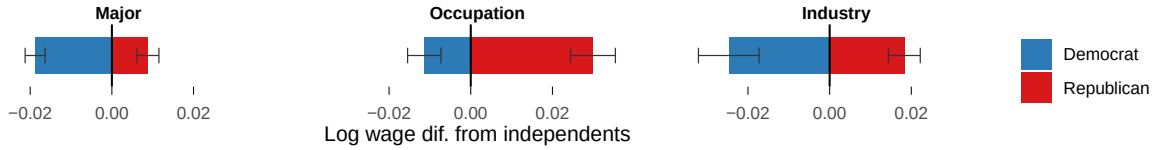
5.2 Valuing ideological amenities

If worker preferences are to explain political sorting, then the fact that partisan employees do not seem to value these canonical amenities differently suggests that they may pay attention to attributes of the job that have been studied less. Therefore, instead of amenities like telecommuting, we focus on "ideological amenities": characteristics that might appeal to Democrats and Republicans differently because of their *ideological* valence. This is closer in spirit to subjective

[40] We measure median wage income by field of degree and state from the 2012 to 2022 ACS samples, mean annual wage by 5-digit "broad" SOC code in each state from the 2022 BLS Occupational Employment and Wage Statistics, and mean annual pay by 4-digit NAICS code by state from the 2022 BLS Quarterly Census of Employment and Wages.

[41] While none of the characteristics they study are related explicitly to ideology or politics, they do measure workers' valuations for a job with "frequent opportunities to make a positive impact on your community or society" versus "occasional opportunities." We find no significant difference in Democrats' and Republicans' WTP for this attribute.

Figure 7: Partisanship and Pay by Major, Occupation, and Industry



Notes: The figure shows coefficient estimates from regressions of a group-level measure of log income on party indicators with fixed effects for the individual's CZ; interactions of year of birth, gender, and race; and the college that the individual attended. Estimates are relative to those who are neither Republicans nor Democrats. See the text for data sources. Standard errors are clustered at the level that income is measured, and horizontal bars are 95% confidence intervals.

amenities such as a desire for "meaningful" or "mission-driven" versus "immoral" work (e.g., Cassar and Meier, 2018; Burbano et al., 2022; Schneider et al., 2020); dignity at work (Dube et al., 2022); corporate social responsibility (e.g., Burbano, 2016; Hedblom et al., 2019); or environmental, social, and governance practices (e.g., Colonnelli et al., 2023). However, ideological amenities are not reducible to any of those alone. We study the bundle of attributes that appeal differently to Democratic and Republican employees to isolate the role of worker preferences in driving segregation.

It is notoriously difficult, however, to estimate workers' preferences from observational data. As the canonical Rosen (1986) model of equalizing differences highlights, the equilibrium relationship between wages and amenities reflects the preferences of the marginal worker, not the average Democrat or Republican.[42] Moreover, in our context, understanding the contribution of labor supply to segregation requires the *distribution* of workers' preferences, not just the average. In Appendix E, we extend the basic Rosen model to include two groups of workers with different preferences, which produces a key prediction: the degree of segregation between the two groups decreases with the variance of how much workers value the amenity. In other words, holding fixed the average Republican's preferences, the existence of Republicans who prefer a liberal workplace constrains the extent of segregation that ideological amenities can create, since these individuals will end up at workplaces where they mix with Democrats.

To estimate the distribution of worker preferences for ideological amenities, we adopt a stated-preference, discrete-choice experimental design that follows Mas and Pallais (2017), Wiswall and Zafar (2018), and Maestas et al. (2023). This lets us observe the same individual choosing between (hypothetical) jobs multiple times, eliminating bias due to unobserved individual characteristics and incomplete knowledge of the choice set, and abstracting from employer preferences and matching frictions. While stated-preference results cannot substitute for variation in real jobs, Mas and Pallais (2017) show that results based on surveys and experimental variation in wages are largely

[42] In addition, the observational relationship between partisanship and pay captures unobserved individual-level heterogeneity in skills, employer preferences, and matching frictions, not just worker preferences alone (e.g., Bonhomme and Jolivet, 2009; Lavetti, 2023). For example, Republicans might be more skilled, conditional on observable education, and therefore have jobs with higher wages *and* more Republican coworkers. This would bias the estimated relationship between wages and Republican affiliation upwards relative to the component due to worker preferences. Another issue is that we do not observe the full set of jobs that each individual could have chosen. Search frictions may imply that Republicans rarely have the option to choose a firm that emphasizes liberal values. Even tracking individuals over time fails to account for the endogeneity of job changes and leads to biased estimates.

consistent.[43]

5.3 Ideological signals in company descriptions

What are the ideological amenities that Democratic and Republican workers value differently? To answer this, we turn to the text of company pages in the LinkedIn data. These pages include short descriptions (on average, 100 words) that sometimes contain appeals to prospective employees. For example, the Nike page describes the company as "a purpose-driven organization energized by a shared commitment to move the world forward through the power of sport. We champion diversity and amplify individual passions ... We work together, embracing our differences and pushing boundaries, to achieve collective success..."[44]

Using our merged LinkedIn-voter file data, we associate each company's description with the shares of Democrats and Republicans who work there. Then, we test whether the company's partisanship predicts the use of signaling words, accounting for the company's industry.[45] Estimates for 20 selected words are shown in Figure 8. We find that Republican firms tend to highlight relatively "neutral" or nonpartisan concepts like service, professionalism, and integrity, relative to Democratic firms in the same industry. The differences are more striking, however, for the distinctively Democratic words such as community, diversity, and sustainability. In fact, Republican companies are signaled more clearly by the *absence* of words related to these liberal values.[46]

Although the words themselves may not be particularly surprising, the fact that partisan employers (in the same industry) describe themselves very differently to prospective employees leads us to ask whether partisan workers value these signals and use them to sort into politically compatible workplaces. Put differently, our question is as follows: does the fact that a company presents itself as a "purpose-driven" organization that "champions diversity" appeal more to Democrats versus Republicans, and would Democrats accept a smaller paycheck to work there versus a similar company that does not highlight the same values? If so, these preferences might explain political sorting across workplaces.

5.4 Framing experiment

Our first survey experiment operationalizes ideological amenities by measuring how much a job ad appeals to Democrats and Republicans when it is *framed* to highlight the company's liberal commitments. We now describe how we generate these job ads before explaining the survey design and presenting results.

5.4.1 Creating synthetic job ads

Our goal is to mimic real job ads like those on an online job board while ensuring that the choices we present are realistic for the survey respondent, and our approach is to generate synthetic or

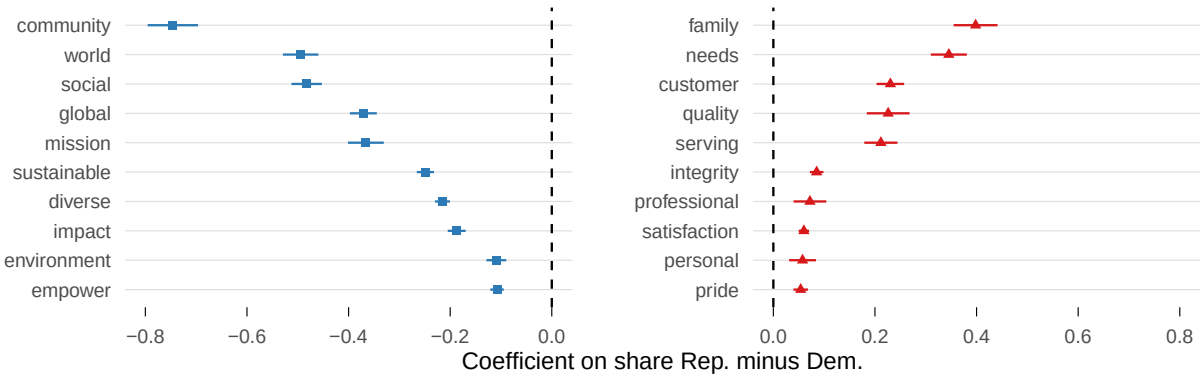
[43] Among stated-preference methods for estimating job-seekers' preferences, Feld et al. (2022) find that discrete choice experiments perform best.

[44] From linkedin.com/company/nike/about, accessed August 20, 2024.

[45] Specifically, we estimate $W_j = \beta P_j + I_j + \varepsilon_j$, where W_j is the number of times a given word is used in the company's description per 1000 words, P_j is the share of Republicans minus the share of Democrats at the company, and I_j are industry fixed effects.

[46] This asymmetry is consistent with Barari (2024), who finds that corporations use speech with liberal or Democratic appeals.

Figure 8: Ideological Signals in Company Descriptions



Notes: The figure shows estimates from separate regressions of the number of times a word is used in the company's description (per 1,000 words) on the share of Republicans minus the share of Democrats at the company, controlling for industry fixed effects. We show estimates for 20 selected words. Horizontal bars are 95% confidence intervals computed using heteroskedasticity-robust standard errors.

hypothetical job advertisements.[47] We use a large language model (OpenAI's ChatGPT 4o) to make these specific to job titles that the respondent has expressed interest in and to optionally emphasize liberal values in a naturalistic manner.[48]

Following from our text analysis, we focus on the presence of terms indicating liberal commitments, generating job ads that do or do not emphasize liberal values using two different prompts. The first includes the instruction to "Emphasize a liberal value like equity, inclusion, anti-racism, or sustainability," and the second includes the instruction to "Be neutral and don't mention anything political. Don't mention diversity, sustainability, or the environment." Generating three pairs of ads for each job title and industry yields 29,796 synthetic job ads; we show examples in Table 3. Following a review of real job postings, we also add one of several possible closing sentences that highlights liberal commitments, e.g., "Self-care is encouraged here, and we are committed to providing a positive work culture" or does not, e.g., "Our energetic environment creates a unique opportunity for professional growth." Further details about the way we construct the list of job titles and map them to industries, as well the text of the LLM prompts and the list of closing sentences, are in Appendix C.

To test whether the synthetic jobs signal partisanship the same way that the LinkedIn company descriptions do, we extend the exercise in Figure 8 by estimating the relationship between the text of company descriptions and the company's partisanship. Define D_j as the two-party share of Democrats at company j, I_j as the company's industry, and $\mathbf{T}_j$ as the vector of word frequencies used in j's description. We fit the following relationship using a support vector machine, a machine learning model selected to account for the high dimensionality and sparsity of $\mathbf{T}_j$:[49]

[47] Other studies have added non-standard amenities such as the company's mission to a conjoint experiment in the typical tabular format (e.g., Non et al., 2022). Our approach is inspired by Colonnelli et al. (2023), but we use an LLM to generate ads specific to the individual's desired job title.

[48] Recent advances in large language models have made highly personalized survey experiments more feasible. For example, Velez and Liu (2023) use an LLM to test the effects of personalized counter-attitudinal information. We also follow Miano (2023) in personalizing the treatment using the respondent's reported occupation.

[49] $\mathbf{T}_j$ contains the 3,000 most common words in the corpus weighted by inverse frequency (TF-IDF). We choose this

Table 3: Examples of Synthetic Job Descriptions

Liberal framing	**Neutral framing**
Title: Web Designer *Industry*: Marketing	*Title*: Web Designer *Industry*: Marketing
We are a progressive marketing agency dedicated to fostering creativity, inclusion, and social impact in every campaign we launch. Our services range from innovative social media strategies to comprehensive digital marketing solutions, aimed at driving engagement and building lasting relationships between brands and their audiences. Grounded in the principle of equity, we strive to level the playing field for businesses and diverse audiences alike, ensuring all voices are represented and amplified. With a talented, diverse team and a commitment to ethical practices, we believe in creating a positive and lasting impact in the world of marketing.	We are a dynamic midsize marketing agency providing next-level digital solutions to a diverse portfolio of clients across multiple industries. Our main activities include brand strategy development, innovative content creation, and precise digital marketing campaigns with a substantial focus on user experience and website optimization. We also specialize in SEO, social media management, and data-driven analytics, aligning all efforts with client objectives to generate measurable results. Our team of marketing experts is committed to delivering creative, high-impact marketing strategies that drive business growth and reach target audiences effectively.
We are seeking a Web Designer to craft stunning, user-friendly websites that resonate with diverse communities and help our clients achieve their marketing goals.	We are seeking a Web Designer to create visually compelling and user-friendly websites, ensuring seamless functionality and adhering to our clients' branding guidelines.

Notes: The table shows two examples of LLM-generated job ads. Both are for a "Web Designer" in the marketing industry at a midsize firm, but one emphasizes liberal values and the other does not. See text for details and Appendix C for the LLM instructions used to generate these ads.

$$D_j = \mathbf{T}_j' \boldsymbol{\gamma}_j + I_j + \varepsilon_j.$$

Then, we use this model to predict $\hat{D}_j$, the Democratic share that each of our synthesized job offers corresponds to. We find that the liberal ads have a predicted Democratic share that is 10.2 p.p. higher than the neutral ads, on average. This confirms that the liberal ads generated by the LLM indeed use words that are associated with more Democratic firms in the observational data. Later, we will use this model to connect the survey estimates to the cross-sectional variation in partisanship and to our second survey.

5.4.2 Survey design

To provide respondents realistic job choices, we first ask them for their current (or most recent) job title and annual salary. Based on their occupation, we then ask them to choose from up to 10 closely related job titles that they would like to see advertisements for (for example, someone who entered "Human Resources Specialist" could choose to see ads for "Talent Acquisition Manager").

simple approach for its transparency and interpretability, but the results are similar using word stems or embeddings generated using BERT. We train this model on the 299,562 company descriptions in our data that are associated with at least 10 people in our sample and are at least 100 characters long. In Appendix Table A13, we show the performance of alternative algorithms and the WTP estimates that they imply.

We inform participants that their responses will help us train an AI tool to make personalized job recommendations, so respondents are aware that the ads they see are hypothetical.[50]

The experiment consists of five pairs of job choices, each presented with four distinct salary offers, for a total of 20 choices per respondent. For each, we first pick a job title (from among those the respondent is interested in) and an industry corresponding to that job title.[51]

Then, for four of the five choices, we present one synthetic job ad that uses liberal values and one that uses neutral values. To ensure the respondent has read and understood these ads, we ask them to briefly describe the main differences between the two jobs. We then draw four salary differentials between 0 and 12% (varying the salary offers around the respondent's current salary), assign them randomly to the two jobs, and ask respondents which job they prefer. An example of the survey layout is shown in Appendix Figure C1.

The remaining choice differs from the others in two ways. First, both job descriptions are generated using either neutral or liberal values. Second, the last pair of salary offers is always the same as the first. Since results from stated-preference experiments may be affected by inattentive respondents, we present the same salary pair twice in quick succession to test whether respondents' choices are consistent. We use the share who respond inconsistently to estimate the inattention rate, which we incorporate directly in our estimation of WTP.[52]

5.4.3 Estimation

Our econometric framework follows Mas and Pallais (2017) and Maestas et al. (2023). We posit the following linear indirect utility function

$$V_{ij} = \alpha + Z_{ij}\beta + \gamma \log w_{ij} + \varepsilon_{ij}$$

where V_{ij} is the indirect utility for individual i choosing job j, $Z_{ij} = 1$ if the job uses the liberal framing, and w_{ij} is the wage offered to individual i for job j. We make the standard assumption that ε_{ij} follows an extreme value type I distribution. The probability that an individual who is fully attentive chooses Job A over Job B is

$$P_A = \Pr\left(\varepsilon_{i2} - \varepsilon_{i1} < (Z_{i,A} - Z_{i,B})\beta + \gamma(\log w_{i,A} - \log w_{i,B})\right).$$

Since $\varepsilon_{i2} - \varepsilon_{i1}$ follows a logistic distribution,

$$P_A = F\left((Z_{i,A} - Z_{i,B})\beta + \gamma(\log w_{i,A} - \log w_{i,B})\right)$$

where F is the logistic CDF. Later, we assess this distributional assumption by plotting the choice data nonparametrically.

[50] This also provides an incentive for respondents to report their preferences accurately. While we do not show respondents job recommendations based on their choices (e.g., Bustelo et al., 2023), Drake et al. (2023) argue that this approach has advantages relative to a setup using real jobs in which workers might "try to appease potential employers by altering statements about their preferences."

[51] Industry options are those that at least 3% of people in the occupation (6-digit SOC code) corresponding to the title have worked in, according to BLS data.

[52] We find, like Mas and Pallais (2017), that this correction mostly affects the tails of the estimated WTP distribution, not the median.

Given β and γ, individual i's willingness to pay for liberal framing makes them indifferent between Job A with $Z_{i,A} = 0$ and wage w and Job B with $Z_{i,B} = 1$ and wage $w - WTP_i$, where all other attributes of the jobs are the same. Thus

$$\gamma \log w + \varepsilon_{i1} = \beta + \gamma \log(w - WTP_i) + \varepsilon_{i2}$$

and

$$WTP_i = w \left[1 - \exp\left(-\frac{\beta + (\varepsilon_{i2} - \varepsilon_{i1})}{\gamma} \right) \right].$$

We compute quantiles of the WTP distribution. Under the parametric assumption above, the qth quantile of WTP for the liberal framing is

$$WTP^{(q)} = w \left[1 - \left(\frac{1-q}{q} \right)^{1/\gamma} \exp\left(-\frac{\beta}{\gamma} \right) \right].$$

We report the WTP expressed as the share of the wage the respondent would give up for the $Z = 1$ job, WTP$/w$.

Inattention: Inattentive individuals are equally likely to select each job. If 2α of individuals are inattentive, in expectation, half of them will choose each job by chance, so the probability that an individual in our survey chooses Job A is

$$\tilde{P}_A = (1 - 2\alpha) P_A + \alpha. \tag{4}$$

Our estimate of α is the share of people who respond inconsistently when presented with the same choice (with the same salary offers) twice in quick succession.[53] We compute this separately for each group: by party or vote choice in our main results, and by party and demographic characteristics when we analyze heterogeneity.

Given $\hat{\alpha}$, we estimate the above equation via maximum likelihood to recover $\hat{\beta}$ and $\hat{\gamma}$. To account for variability in the estimation of the inattention rate, we compute standard errors using the bootstrap, clustering at the individual level.

Nonparametric correction: Following Mas and Pallais (2017), to assess our parametric assumptions, we can also plot the share of individuals who chose the liberal framing at each wage difference corrected for inattentiveness.

Let $P_{\Delta \log w}$ be the probability that an attentive individual chooses the $Z = 1$ job when the wage difference between the $Z = 1$ and $Z = 0$ job is $\Delta \log w$. Then, for a given $\Delta \log w$, the share of survey respondents (who may be inattentive) who choose $Z = 1$ is

$$Y_{\Delta \log w} = (1 - 2\alpha) P_{\Delta \log w} + \alpha + \nu,$$

[53] This is different from Mas and Pallais (2017), who use a dominated choice to estimate inattention, and closer to the approach in Clayton et al. (2023).

where ν is sampling error. We use $\hat{\alpha}$ to transform this share so that it is an unbiased estimate of the share of respondents whose WTP for the liberal framing is more than $-\Delta \log w$:

$$\tilde{Y}_{\Delta w} = \frac{Y_{\Delta w} - \hat{\alpha}}{1 - 2\hat{\alpha}} = P_{\Delta \log w} + \nu.$$

5.4.4 Sample

We ran the survey through Prolific, an online survey platform, in August 2024, restricting the sample to individuals over 25 years old who are registered to vote and employed or seeking a job. We received 1,002 responses in total. After dropping 5 respondents who were not interested in any of the job titles we offered and 22 who reported their current salary as less than $5,000 or more than $300,000, we are left with 975 respondents. Summary statistics for the survey sample are shown in Appendix Table C1. The sample approximately aligns with the merged LinkedIn-voter file sample on important characteristics, particularly the distributions of educational attainment and party affiliation: 54% of the survey sample has a bachelor's degree and 25% has a graduate degree, compared to 57% and 26% of the LinkedIn sample (conditional on reporting education). Our survey sample contains 48% Democrats and 33% Republicans, compared to 40% and 29% of the LinkedIn sample. The average reported salary of respondents is $68,456. We also cover a broad range of occupations: respondents span 247 6-digit SOC codes.

5.4.5 Results

We first present a visual summary of the experimental data in Figure 9, which shows a binned scatterplot of the fraction of respondents who chose the liberal framing versus the salary differential (in log points). The blue series shows estimates for Democrats and the red for Republicans. We show the raw data and estimates corrected for inattention. We also overlay predictions from the model in Eq. 4 (estimated using maximum likelihood), which can be interpreted as CDFs of the WTP distribution.

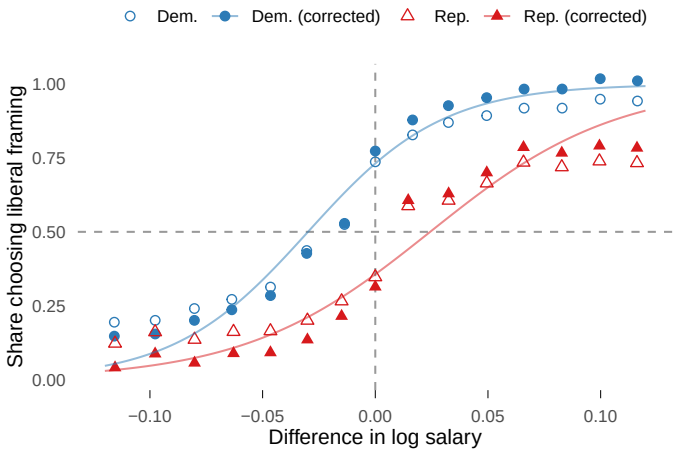
Both series slope upward – indicating that respondents prefer jobs that pay more, in general – but the blue series dominates the red. This means that Democrats are more likely than Republicans to choose the job with the liberal framing no matter the salary differential. We further see that the median WTP – the location where each series crosses 0.5 on the vertical axis, which is where the respondent is indifferent between the two jobs – is about -0.030 for Democrats and 0.024 for Republicans, indicating that the median Democrat would trade off 2.9% of their salary ($1,700 per year) for the liberal framing and the median Republican would trade off 2.4% ($1,600) to avoid it. The model also implies that 25% of Democrats would trade off 6.0% or more of their salary for the liberal framing, while 25% of Republicans would trade off 6.9% or more to avoid it.[54] This suggests that there is a group of workers for whom the liberal framing might be especially salient or important.

Indeed, these differences are larger if we examine a more detailed measure of party identification or Presidential vote choice, which we display in Figure 10. First, we show in Panel (a) that *strong* Democrats are willing to trade off 3.5% of their salary for the liberal framing, while strong Republicans are willing to trade off 3.6% to avoid it.[55] Independents are, on average, indifferent to

[54] We show quantiles of the WTP distribution by party in Appendix Table A11.

[55] After the initial party affiliation question, we ask those who identify as Democrats whether they "are a strong Democrat or lean Democrat," and likewise for Republicans.

Figure 9: Willingness to Pay for Liberal Framing



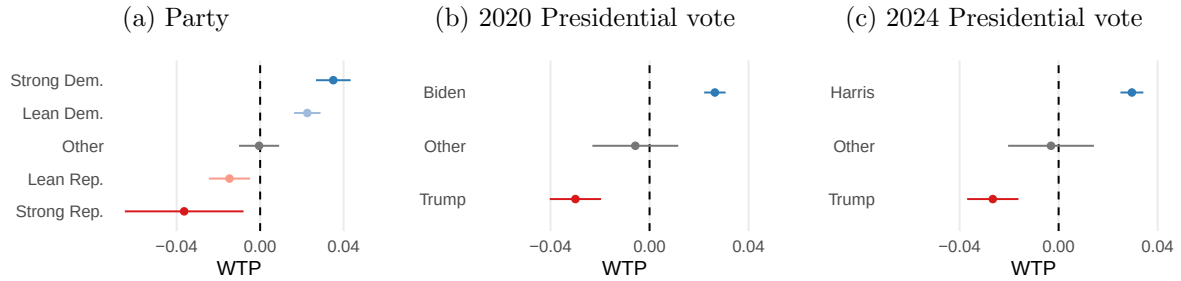
Notes: The open markers show the share of respondents who chose the job with the liberal framing versus the difference in log salary between the two jobs. The filled markers show these shares corrected for inattention. We show separate series for Democrats (the blue circles) and Republicans (the red triangles). The lines display predictions from the model in Eq. 4.

framing when the jobs do not have a salary difference. Second, these patterns persist when we split respondents by their reported 2020 Presidential vote choice in Panel (b) or 2024 vote intention in Panel (c). The difference in WTP for the liberal framing is 5.6 p.p. between those who voted for Joe Biden versus Donald Trump in 2020, and also 5.6 p.p. between those who intend to vote for Kamala Harris versus Trump in 2024.

Are these magnitudes reasonable? A potential concern with surveys that infer valuations of non-market goods is "hypothetical bias": individuals presented with hypothetical choices might report higher valuations than revealed by their choices in real markets (e.g., List, 2001). The magnitudes of our estimates, however, are roughly in line with other work, and in fact somewhat smaller. Colonnelli et al. (2023) report that job-seekers in Brazil value ESG signals as equivalent to 10% of average wages; Non et al. (2022) find that, among students in the Netherlands, the estimated value of a company mission advertising corporate social responsibility is about 10% of the salary offer; Kesternich et al. (2021) find in Germany that work meaning reduces the reservation wage by about 18% for those who say this is a very important aspect of their job; and Burbano (2016) finds that workers on an online labor market reduce their wage demands by up to 44% when informed about the employer's social responsibility. Our setting is different in that Republicans view the liberal framing as a *disamenity*, but the size of the gap between Republicans and Democrats is similar to previous work on nonstandard amenities, especially since we manipulate relatively small differences in wording, not verifiable signals such as an ESG designation.

We next examine heterogeneity in the WTP estimates by demographic groups within parties, which we show in Figure 11. For the most part, we find that the valuations are remarkably similar across these groups. In Panel (a), we split the sample by party and educational status, and in Panel (b) by the respondent's current annual (reported) salary. In Panel (c), we split the sample by the number of coworkers reported by the individual. We see some indication that Republicans at larger workplaces (with more than 100 coworkers) have a smaller WTP to avoid the liberal framing. In Panels (d) through (f), we split the sample by birth year, race, and gender, respectively. Here, we

Figure 10: Willingness to Pay for Liberal Framing by Party and Presidential Vote



Notes: The figures show the median willingness to pay (expressed as a fraction of the salary) for a job ad that uses the liberal framing, separately by party identification in Panel (a), reported 2020 Presidential vote in Panel (b), and intended 2024 Presidential vote in Panel (c). Estimates are corrected for inattention, and horizontal bars are 95% confidence intervals with standard errors computed using the bootstrap with 500 replications.

find that older and white Republicans are particularly averse to the liberal framing, while women tend to prefer to it more than men, especially among Republicans and independents.

While we detect only small amounts of *within*-party heterogeneity, these demographic characteristics are correlated with partisanship – e.g., women and non-white individuals are more likely to be Democrats, so it is true that women and non-white respondents have a higher WTP for the liberal framing. However, partisanship predicts an individual's likelihood of preferring the liberal framing better than any other demographic variable that we measure. In Appendix Table A12, we show results from a "horse-race" regression where we regress an indicator for whether the respondent chose the liberal framing on all of these demographics simultaneously, controlling for the wage difference between the two job offers. The coefficient estimates are large and significant for the female, African American, and Hispanic/Latino indicators, but they are substantially smaller than the coefficients on the partisanship indicators. This is consistent with the patterns in Figure 11 and confirms that workers differ in their valuation for the liberal framing primarily along partisan lines.

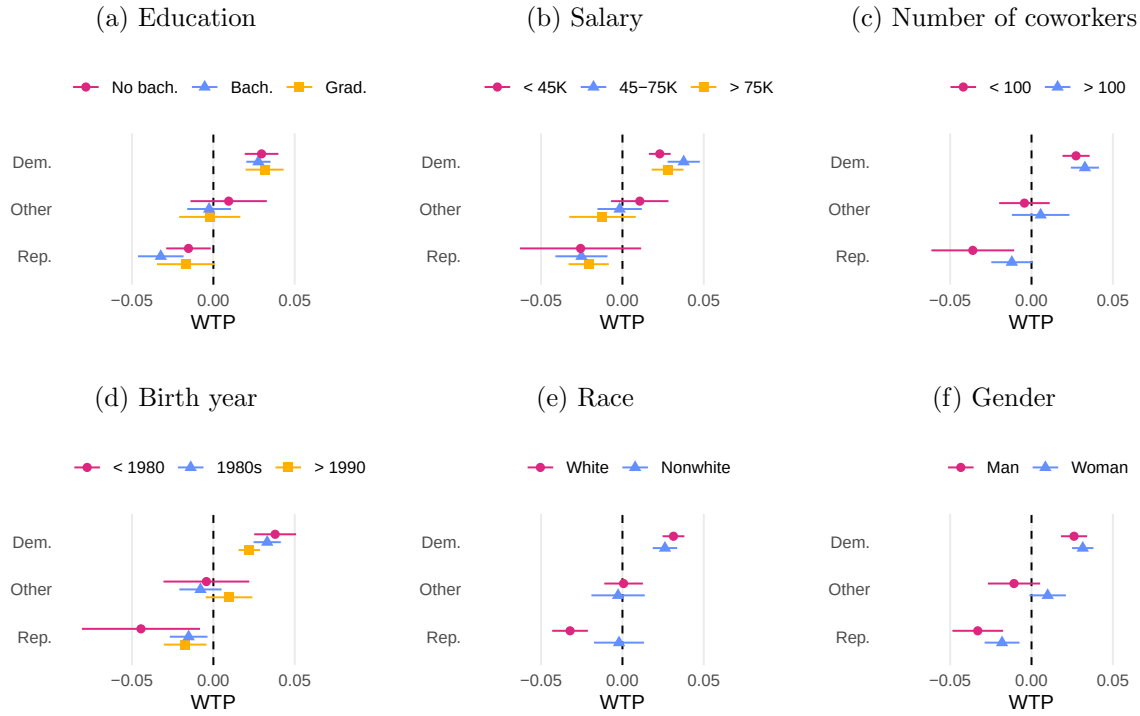
5.4.6 Interpretation

These results show that Democrats and Republicans value relatively minor changes in the language used in job ads, and do so in a way that corresponds to how partisan companies are described on their LinkedIn pages. Next, we make this connection explicit: we re-estimate Eq. 4, but we replace the indicator for liberal framing Z_{ij} with $\hat{D}_{ij}$, the predicted share of Democrats corresponding to the job ad, and calculate willingness to pay for this variable. We scale $\hat{D}_{ij}$ by the standard deviation of the Democratic share in the training sample, which is 24.2 p.p., so the estimates can be interpreted as the median respondent's willingness to pay for a job ad that corresponds to a company with a Democratic share 1 SD higher.[56]

Using our baseline prediction algorithm, we find that the median Democrat is willing to trade off

[56]This is not exactly equivalent to scaling the WTP estimates by the average difference in $\hat{D}_{ij}$ between liberal and neutral job ads because there is heterogeneity in $\hat{D}_{i,A} - \hat{D}_{i,B}$ across pairs and because the WTP estimand is a nonlinear function of the model parameters.

Figure 11: Heterogeneity in Willingness to Pay for Liberal Framing



Notes: The figures show the median willingness to pay (expressed as a fraction of the salary) for a job ad that uses the liberal framing, separately by party and by characteristics of the respondent. Estimates are corrected for inattention, and horizontal bars are 95% confidence intervals with standard errors computed using the bootstrap with 500 replications.

4.7% ($2,700) of their salary to go to a company whose description signals 1 SD more Democratic coworkers – e.g., from Farmer's Insurance (49% Democrats) to Microsoft (73%). Strong Democrats are willing to pay 5.3%. Republicans would pay 4.3% ($2,800) of their salary to avoid this; strong Republicans 6.3%. We show these estimates and results using alternative prediction algorithms in Appendix Table A13. The baseline algorithm yields conservative estimates because it predicts the largest difference between the liberal and neutral job ads, 10.2 p.p.

These estimates are similar in magnitude to workers' WTP for other amenities measured by Maestas et al. (2023) such as the ability to telecommute (4.5%) or a relaxed versus fast pace of work (4.4%), but they are smaller than WTP for amenities with clear monetary value, like 10 days of paid time off (17.8%). However, Democrats and Republicans both have similar, positive valuations for these canonical amenities – so they cannot generate partisan sorting across workplaces – while the two groups' preferences for ideological amenities go in opposite directions.

This scaling also enables us to compare the estimates from the framing experiment with those from the conjoint experiment, which we describe next.

5.5 Conjoint experiment

When workers look for jobs, they may rely on the language in job advertisements to assess their ideological compatibility with potential employers. The framing experiment, which mimics these

ads, estimates workers' preferences for ideological amenities in a more natural way than showing respondents a table of characteristics about the job and their coworkers.[57] Moreover, the effect of this bundle is the relevant quantity for understanding the overall contribution of these amenities to political segregation at the workplace. However, the framing experiment cannot tell us exactly which attributes workers are responding to – whether Democrats prefer the liberal framing because they explicitly think their coworkers will be Democrats, or will be more likely to be women, or because they imagine they will fit the company culture better, for example. To un-bundle this treatment, we use a conjoint design to disentangle workers' preferences for co-partisan coworkers from other coworker demographics and company characteristics.[58]

5.5.1 Survey design

The second survey starts the same way as the first: we ask respondents for their current job title and annual salary, then for up to 10 job titles that they would like to see advertisements for. The experiment consists of 12 pairs of jobs, each presented with one pair of randomly varying wage offers anchored around the respondent's current salary. For each pair, we pick one job title, which never varies between Job A and Job B, but in contrast to the previous experiment, we allow the industry to potentially vary. In addition to industry, we vary the following: coworker characteristics including the share of Democrats, the share of men, the share who are white, and the average age; company characteristics including the number of employees, whether and to which party the company makes political contributions; and two job benefits, such as medical insurance or tuition reimbursement. We vary only four of these characteristics for each pair to reduce the complexity of the comparisons that respondents must make. More details about the options we provide for each attribute are in Appendix D, and Appendix Figure D1 shows how the job choices are presented to respondents.

5.5.2 Sample

We collected responses in August 2024 through the Harvard Digital Lab for the Social Sciences, a volunteer panel operated by Harvard University.[59] After dropping 1 respondent who was not interested in any of the job titles we offered and 13 who reported their current salary as more than $300,000, we are left with 574 responses. Summary statistics on the analysis sample are in Appendix Table D1. On partisanship, we again roughly match the LinkedIn sample, with 35% Democrats and 30% Republicans (compared to 40% and 29% of the LinkedIn sample). However, this sample is more educated on average: 32% of the survey sample has a bachelor's degree and 39% has a graduate degree, compared to 57% and 26% of the LinkedIn sample. The survey sample also skews older and male, and considerably higher-income than our other survey sample and the LinkedIn population, with an average reported salary of $90,725. Importantly, however, we again cover workers in a wide range of occupations (203 SOC codes).

[57] Horiuchi et al. (2022) find, however, that conjoint designs can mitigate social desirability bias due to simultaneous randomization over sensitive and non-sensitive attributes. We see our two surveys as complements, each addressing a different set of concerns.

[58] Schuh (2024) conducts a similar survey experiment to analyze how employees' preferences about the gender composition of their workplace contribute to occupational gender segregation. This survey is also similar to Mummolo and Nall (2017), who use a conjoint design to study the aspects of neighborhoods that partisans care about in order to understand the drivers of partisan residential segregation. Lastly, Shafranek (2021) uses a similar design to assess the impact of partisanship on roommate selection, finding that respondents preferred to avoid living with a member of the other party about as much as they preferred to avoid someone described as "not at all clean and tidy."

[59] Strange et al. (2019) show that samples of volunteers can replicate other social science findings and produce responses comparable to paid subjects.

5.5.3 Results

Figure 12 shows the median willingness to pay for each attribute separately by party; Appendix Table A14 contains standard errors for the differences in WTP between Democrats and Republicans and between strong Democrats and strong Republicans. Starting with coworker characteristics, we estimate that Democrats are willing to trade off 3.3% of their salary for a company with a share of Democrats that is 1 SD (21.8 p.p.) higher, while Republicans are willing to trade off 4.0% of their salary to go in the other direction. This difference (7.3 p.p.) is more than three times larger than the differential WTP for 1 SD more female coworkers (2.1 p.p.) or nonwhite coworkers (2.1 p.p.), and much larger than for older coworkers, which is approximately zero.

Turning to company characteristics, we first see that partisans are willing to pay to work in industries with a higher share of copartisans. Note that we did not inform respondents about the partisanship of each industry, so these valuations are based on what the individual *perceives* about the industry options we show them. We estimate that Democrats are willing to trade off 0.9% of their salary for an industry with a 1 SD larger Democratic share (9.4 p.p., the two-party share as measured in our observational LinkedIn data), while Republicans are willing to trade off 1.3% to go in the opposite direction. The difference in these two estimates is statistically significant at the 90% level.

We also find that partisans are relatively indifferent to companies that donate to their own party, but are quite averse to companies that donate to the out-party. Democrats would trade off 10.0% of their salary to avoid a company that contributes to Republicans, while Republicans would trade off 7.6% of their salary to avoid a company that contributes to Democrats. Independents prefer companies that do not contribute to either party.

We do not find evidence of differential preferences for company size or – consistent with our reanalysis of Maestas et al. (2023) – for the benefits that we ask about. One exception, however, is that Democrats seem to value health insurance slightly more than others. We also reproduce the same patterns we showed for the framing survey: the differential WTP for 1 SD more Democratic coworkers between *strong* Democrats and strong Republicans is 9.1 p.p., and between Biden and Trump voters in 2020 is 6.4 p.p. (see Appendix Figure A11).

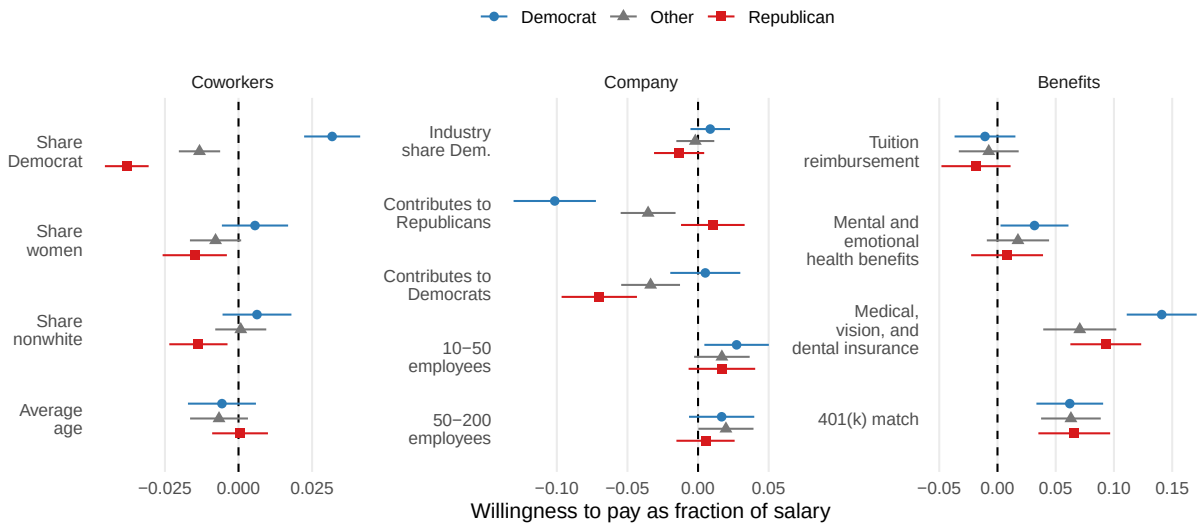
5.6 Comparing the two surveys

We found in the first experiment (under the scaling described in Section 5.4.6) that Democrats are willing to trade off 4.7% of their salary for 1 SD more Democratic coworkers, and Republicans would pay 4.3% of their salary to avoid this. The analogous figures from the conjoint experiment are somewhat smaller: 3.3% and 4.0%. Overall, the gap in WTP between Democrats and Republicans estimated from the conjoint experiment is about 80% as large as the gap estimated from the framing experiment.

As described previously, the two experiments target different estimands: the first measures workers' preferences for a bundle of attributes, while the second attempts to estimate components of that bundle separately. When a respondent reads a job ad with the liberal framing, they likely form expectations not just about the share of coworkers who will be Democrats, but also about the racial and gender composition of the workplace and whether the company is likely to contribute to Democrats or Republicans. Although we cannot simply add up the marginal preferences for these other characteristics, the conjoint experiment gives us some insight into what workers value.

Evidently, workers do care about whether their coworkers share their party affiliation – it seems

Figure 12: Willingness to Pay for Coworker and Job Characteristics



Notes: The figure shows the median willingness to pay (expressed as a fraction of the salary) for coworker and job characteristics, separately by party. Coworker demographic variables are standardized, so the estimate corresponds to the respondent's willingness to pay for a job where the share of Democrats, women, or nonwhite coworkers is 1 SD higher or the average age is 1 SD higher. The industry-level share of Democrats is also standardized, so the estimate corresponds to the respondent's willingness to pay for a job in an industry where the share of Democrats is 1 SD higher. The reference category for corporate contributions is "No political contributions," for company size is "200+ employees," and for benefits is "Life insurance." The covariance estimates for the regression coefficients are clustered by individual, and standard errors for the WTP estimates are computed using the delta method. Horizontal bars are 95% confidence intervals. Appendix Table A14 reports standard errors for the WTP differences between Republicans and Democrats.

unlikely that the preferences estimated from the first experiment are driven solely by preferences for characteristics that are merely correlated with partisanship, such as gender or race. The estimated preferences for corporate contributions, however, are substantial. For example, Democrats would trade off significantly more to avoid a company that donates to Republicans than to avoid a company where the share of Republican employees is 1 SD higher. The estimates of preferences for industry partisanship, though smaller, indicate that the *nature* of the company matters, too.[60] These findings suggest that workers' preferences for ideological amenities are driven partly by whom they expect to run into at the water cooler, but perhaps more so by whether they think the organization shares their ideological stance and whether the work they will be doing broadly accords with their values and interests.

5.7 Comparing to observed segregation

Finally, to understand the relationship between the observed level of segregation and the survey estimates, we return to the Rosen model of equalizing differences. We adjust the model to include three groups – Democrats, Republicans, and others – and plug in our estimates of each group's

[60] These are not, however, part of the estimated effect from the framing experiment, since in that survey we did not vary the industry between jobs within a pair.

WTP from the model estimated in Eq. 4.

Closing the model requires the distribution of firm costs for providing a job with the liberal framing. Since this paper focuses on labor supply, we do not have any estimates of these costs. While we were able to generate ads using an LLM at little expense, in the real world, firms might be held accountable to the values they espouse in a job ad, and their public stances might affect investors' or customers' perceptions. It is also reasonable that the cost of advertising liberal values is slightly positive for some firms (e.g., those in traditionally conservative industries) and slightly negative for others (in more progressive industries). Instead of assuming a particular distribution, we remain agnostic about the firm side, and instead check which levels of segregation are compatible with the estimated preferences by varying the market compensating differential w (the wage difference between the liberal and neutral job) between -10% and 10%. For each w, we calculate the share of workers of each party who would choose a job with the liberal framing, and compare these shares to the CZ-level shares of partisans to recover the overexposure ratio S_{cz}.

These results (shown in Appendix Figure A12) indicate that the model generates a level of segregation corresponding to the observational data ($S_{cz} = 1.100$) when the job with the neutral framing pays 2.3% more than the job with the liberal framing. The magnitude of this differential is similar to the observed differences in partisanship and pay in Section 5.1. Focusing on the range of w between the median Democrat's WTP (2.9%) and median Republican's (-2.4%), S_{cz} implied by the model ranges from 1.097 to 1.143, consistent with the observed data. Indeed, the *maximum* overexposure ratio that is compatible with the survey estimates and the shares of Democrats and Republicans in each CZ is 1.143.

This is just a back-of-the-envelope calculation, given that the model abstracts from employer discrimination and other forces that may impede sorting. But it illustrates the key insight of the compensating differentials framework applied to this context: even without any other frictions, the overlap in Democrats' and Republicans' preferences limits the extent to which they end up in different jobs. Thus, the labor market equilibrium naturally constrains sorting across workplaces. This dynamic is reminiscent of the finding in Mummolo and Nall (2017) that partisans do not sort into politically compatible neighborhoods because of the dearth of housing that fulfills other important criteria like affordability. Our mechanism is different, but the theme is similar: market forces can check the expression of ideology.

6 Conclusion

Political segregation is a pervasive feature of modern American society, but there is little systematic evidence on the different paths that Democrats and Republicans take through the labor market and why they do so. Some have suggested that the workplace might exhibit less polarization than other environments and therefore soothe partisan tensions. Our results show that this is unlikely: we document sizeable differences in the choices that Democrats and Republicans make, demonstrate that they lead to substantial segregation across workplaces, and show that workers are willing to give up wages to sort into ideologically compatible jobs.

While our study focuses on workers' preferences – the labor supply side – future research should investigate employers' preferences. Audit studies suggest that employers discriminate against minority party applicants (Gift and Gift, 2015), and evidence from Brazil shows that firm owners favor workers from their own party (Colonnelli et al., 2024). But in the U.S., would employers pay a premium to hire a co-partisan over an equally qualified worker from the opposing party? Or does

competitive pressure prevent ideological discrimination?

Understanding whether employer behavior exacerbates or mitigates political segregation can help predict the future of workplace polarization. Our results suggest that if workers were unconstrained by labor market forces that push towards more mixing, jobs might be even more segregated than they currently are. Examining other forces that shape political sorting in the labor market is essential not just to understand the trade-offs faced by politically motivated workers and employers, but also to gauge whether the workplace can foster more cross-party exposure compared to other areas of Democrats' and Republicans' increasingly divided lives.

References

Abramowitz, Alan I. and Kyle L. Saunders, "Ideological Realignment in the U.S. Electorate," *The Journal of Politics*, 1998, *60* (3), 634–652.

Adrjan, Pawel, Svenja Gudell, Emily Nix, Allison Shrivastava, Jason Sockin, and Evan Starr, "We've Got You Covered: Employer and Employee Responses to Dobbs v. Jackson," Technical Report, IZA Discussion Papers 2023.

Akerlof, George A. and Rachel E. Kranton, "Identity and the Economics of Organizations," *Journal of Economic Perspectives*, 2005, *19* (1), 9–32.

Alonso-Villar, Olga, Coral Del Rio, and Carlos Gradín, "The Extent of Occupational Segregation in the United States: Differences by Race, Ethnicity, and Gender," *Industrial Relations: A Journal of Economy and Society*, 2012, *51* (2), 179–212.

Åslund, Olof and Oskar Nordström Skans, "How to Measure Segregation Conditional on the Distribution of Covariates," *Journal of Population Economics*, 2009, *22*, 971–981.

Ba, Bocar, Jacob Kaplan, Dean Knox, Mayya Komisarchik, Gregory Lanzalotto, Rachel Mariman, Jonathan Mummolo, Roman Rivera, and Michelle Torres, "Political Diversity in U.S. Police Agencies," Technical Report 2023.

Barari, Soubhik, "Political Speech from Corporate America: Sparse, Mostly for Democrats, and Somewhat Representative," *Journal of Quantitative Description: Digital Media*, 2024, *4*.

Bartels, Larry M., "Partisanship and Voting Behavior, 1952-1996," *American Journal of Political Science*, 2000, pp. 35–50.

Bayard, Kimberly, Judith Hellerstein, David Neumark, and Kenneth Troske, "New Evidence on Sex Segregation and Sex Differences in Wages from Matched Employee-Employer Data," *Journal of Labor Economics*, 2003, *21* (4), 887–922.

Bermiss, Y. Sekou and Rory McDonald, "Ideological Misfit? Political Affiliation and Employee Departure in the Private-Equity Industry," *Academy of Management Journal*, 2018, *61* (6), 2182–2209.

Bonhomme, Stéphane and Grégory Jolivet, "The pervasive absence of compensating differentials," *Journal of Applied Econometrics*, 2009, *24* (5), 763–795.

Bonica, Adam, "Mapping the Ideological Marketplace," *American Journal of Political Science*, 2014, *58* (2), 367–386.

_ , "Are Donation-Based Measures of Ideology Valid Predictors of Individual-Level Policy Preferences?," *The Journal of Politics*, 2019, *81* (1), 327–333.

_ , "Database on Ideology, Money in Politics, and Elections: Public version 3.1 [computer file]," *Stanford, CA: Stanford University Libraries*, 2023.

_ **, Howard Rosenthal, Kristy Blackwood, and David J. Rothman**, "Ideological Sorting of Physicians in Both Geography and the Workplace," *Journal of Health Politics, Policy and Law*, 2020, *45* (6), 1023–1057.

Brown, Jacob R., "Partisan Conversion Through Neighborhood Influence: How Voters Adopt the Partisanship of their Neighbors," *The Journal of Politics*, 2024.

_ **and Ryan D. Enos**, "The Measurement of Partisan Sorting for 180 Million Voters," *Nature Human Behaviour*, 2021, *5* (8), 998–1008.

_ **, Enrico Cantoni, Ryan D. Enos, Vincent Pons, and Emilie Sartre**, "The Increase in Partisan Segregation in the United States," 2022.

_ **,** _ **, Sahil Chinoy, Martin Koenen, and Vincent Pons**, "The Effect of Childhood Environment on Political Behavior: Evidence from Young US Movers, 1992–2021," 2023.

Burbano, Vanessa C., "Social Responsibility Messages and Worker Wage Requirements: Field Experimental Evidence from Online Labor Marketplaces," *Organization Science*, 2016, *27* (4), 1010–1028.

_ **, Nicolas Padilla, and Stephan Meier**, "Gender Differences in Preferences for Meaning at Work," *American Economic Journal: Economic Policy*, 2024, *16* (3), 61–94.

_ **, Olle Folke, Stephan Meier, and Johanna Rickne**, *The Gender Gap in Meaningful Work: Explanations and Implications*, Centre for Economic Policy Research, 2022.

Bustelo, Monserrat, Ana Maria Diaz, Jeanne Lafortune, Claudia Piras, Luz Magdalena Salas, and José Tessada, "What Is the Price of Freedom? Estimating Women's Willingness to Pay for Job Schedule Flexibility," *Economic Development and Cultural Change*, 2023, *71* (4), 1179–1211.

Cantoni, Enrico and Vincent Pons, "Does Context Outweigh Individual Characteristics in Driving Voting Behavior? Evidence from Relocations within the U.S.," *American Economic Review*, 2022, *112* (4), 1226–1272.

Card, David, Alexandre Mas, and Jesse Rothstein, "Tipping and the Dynamics of Segregation," *The Quarterly Journal of Economics*, 2008, *123* (1), 177–218.

Carrington, William J. and Kenneth R. Troske, "On Measuring Segregation in Samples with Small Units," *Journal of Business & Economic Statistics*, 1997, *15* (4), 402–409.

Cassar, Lea and Stephan Meier, "Nonmonetary Incentives and the Implications of Work as a Source of Meaning," *Journal of Economic Perspectives*, 2018, *32* (3), 215–238.

Chetty, Raj and Nathaniel Hendren, "The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects," *The Quarterly Journal of Economics*, 2018, *133* (3), 1107–1162.

_ **, John N. Friedman, Nathaniel Hendren, Maggie R. Jones, and Sonya R. Porter**, "The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility," Technical Report, National Bureau of Economic Research 2018.

Choi, Jung Ho, Joseph Pacelli, Kristina M. Rennekamp, and Sorabh Tomar, "Do Jobseekers Value Diversity Information? Evidence from a Field Experiment and Human Capital Disclosures," *Journal of Accounting Research*, 2023, *61* (3), 695–735.

Clayton, Katherine, Yusaku Horiuchi, Aaron R. Kaufman, Gary King, and Mayya Komisarchik, "Correcting Measurement Error Bias in Conjoint Survey Experiments," 2023.

Colonnelli, Emanuele, Timothy McQuade, Gabriel Ramos, Thomas Rauter, and Olivia Xiong, "Polarizing Corporations: Does Talent Flow to "Good" Firms?," Technical Report, National Bureau of Economic Research 2023.

_ , **Valdemar Pinho Neto, and Edoardo Teso**, "Politics at Work," 2024.

Conway, Jacob and Levi Boxell, "Consuming Values," 2024.

Cutler, David M., Edward L. Glaeser, and Jacob L. Vigdor, "The Rise and Decline of the American Ghetto," *Journal of Political Economy*, 1999, *107* (3), 455–506.

Davis, Donald R., Jonathan I. Dingel, Joan Monras, and Eduardo Morales, "How Segregated Is Urban Consumption?," *Journal of Political Economy*, 2019, *127* (4), 1684–1738.

d'Haultfoeuille, Xavier and Roland Rathelot, "Measuring segregation on small units: A partial identification analysis," *Quantitative Economics*, 2017, *8* (1), 39–73.

Drake, Marshall, Neil Thakral, and Linh T. Tô, "Wage Differentials and the Price of Workplace Flexibility," 2023.

Dube, Arindrajit, Suresh Naidu, and Adam D. Reich, "Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers," Technical Report, National Bureau of Economic Research 2022.

Engelberg, Joseph, Jorge Guzman, Runjing Lu, and William Mullins, "Partisan Entrepreneurship," Technical Report, National Bureau of Economic Research 2022.

Estlund, Cynthia, *Working Together: How Workplace Bonds Strengthen a Diverse Democracy*, Oxford University Press, 2003.

Feld, Brian, AbdelRahman Nagy, and Adam Osman, "What do jobseekers want? Comparing methods to estimate reservation wages and the value of job attributes," *Journal of Development Economics*, 2022, *159*, 102978.

Firoozi, Daniel, "The Effect of Selective Colleges on Student Partisanship," 2023.

Fos, Vyacheslav, Elisabeth Kempf, and Margarita Tsoutsoura, "The Political Polarization of Corporate America," 2022.

Frake, Justin, Reuben Hurst, and Max Kagan, "Partisan segregation in the U.S. workplace is large and rising," 2024.

Gentzkow, Matthew and Jesse M. Shapiro, "Ideological Segregation Online and Offline," *The Quarterly Journal of Economics*, 2011, *126* (4), 1799–1839.

_ , _ , **and Matt Taddy**, "Measuring Group Differences in High-dimensional Choices: Method and Application to Congressional Speech," *Econometrica*, 2019, *87* (4), 1307–1340.

Gift, Karen and Thomas Gift, "Does Politics Influence Hiring? Evidence From a Randomized Experiment," *Political Behavior*, 2015, *37*, 653–675.

Glaeser, Edward L. and Bryce A. Ward, "Myths and Realities of American Political Geography," *Journal of Economic Perspectives*, 2006, *20* (2), 119–144.

Goncalves, Felipe and Cody Tuttle, "Police Work and Political Identity," 2024.

Green, Donald P., Bradley Palmquist, and Eric Schickler, *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters*, Yale University Press, 2004.

Hedblom, Daniel, Brent R. Hickman, and John A. List, "Toward an Understanding of Corporate Social Responsibility: Theory and Field Experimental Evidence," 2019.

Hellerstein, Judith K. and David Neumark, "Workplace Segregation in the United States: Race, Ethnicity, and Skill," *The Review of Economics and Statistics*, 2008, *90* (3), 459–477.

Hertel-Fernandez, Alexander, "Power and Politics in the U.S. Workplace," *Economic Policy Institute*, 2020.

Hoang, Thao, Phong T.H. Ngo, and Le Zhang, "Polarized Corporate Boards," 2022.

Horiuchi, Yusaku, Zachary Markovich, and Teppei Yamamoto, "Does Conjoint Analysis Mitigate Social Desirability Bias?," *Political Analysis*, 2022, *30* (4), 535–549.

Huber, Gregory A. and Neil Malhotra, "Political Homophily in Social Relationships: Evidence from Online Dating Behavior," *The Journal of Politics*, 2017, *79* (1), 269–283.

Hurst, Reuben and Saerom Ronnie Lee, "Do Diversity Claims Cause Labor Market Sorting by Political Partisanship? Evidence from Experiments," 2024.

Indeed, "Navigating Politics in the Workplace During a Divisive Election Year," 2024. Accessed 2024-08-09.

Iyengar, Shanto, Tobias Konitzer, and Kent Tedin, "The Home as a Political Fortress: Family Agreement in an Era of Polarization," *The Journal of Politics*, 2018, *80* (4), 1326–1338.

Kempf, Elisabeth and Margarita Tsoutsoura, "Partisan Professionals: Evidence from Credit Rating Analysts," *The Journal of Finance*, 2021, *76* (6), 2805–2856.

Kesternich, Iris, Heiner Schumacher, Bettina Siflinger, and Stefan Schwarz, "Money or meaning? Labor supply responses to work meaning of employed and unemployed individuals," *European Economic Review*, 2021, *137*, 103786.

Kim, Woojin, "Political polarization in medicine," 2024.

Lavetti, Kurt, "Compensating Wage Differentials in Labor Markets: Empirical Challenges and Applications," *Journal of Economic Perspectives*, 2023, *37* (3), 189–212.

LaViers, Lisa and Jason Sandvik, "The Impact of Workplace Gender Diversity Disclosures on Job Selection Decisions," 2023.

Levendusky, Matthew, *The Partisan Sort: How Liberals Became Democrats and Conservatives Became Republicans*, University of Chicago Press, 2009.

Levy, Ro'ee, "Social Media, News Consumption, and Polarization: Evidence from a Field Experiment," *American Economic Review*, 2021, *111* (3), 831–870.

List, John A., "Do Explicit Warnings Eliminate the Hypothetical Bias in Elicitation Procedures? Evidence from Field Auctions for Sportscards," *American Economic Review*, 2001, *91* (5), 1498–1507.

Maestas, Nicole, Kathleen J. Mullen, David Powell, Till Von Wachter, and Jeffrey B. Wenger, "The Value of Working Conditions in the United States and Implications for the Structure of Wages," *American Economic Review*, 2023, *113* (7), 2007–2047.

Mangold, Frank, David Schoch, and Sebastian Stier, "Ideological self-selection in online news exposure: Evidence from Europe and the US," *Science Advances*, 2024, *10* (37).

Martin, Gregory J. and Steven W. Webster, "Does Residential Sorting Explain Geographic Polarization?," *Political Science Research and Methods*, 2020, *8* (2), 215–231.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, 2017, *107* (12), 3722–3759.

Mason, Lilliana, *Uncivil Agreement: How Politics Became Our Identity*, University of Chicago Press, 2018.

Matzat, Johannes and Aiko Schmeißer, "Do Unions Shape Political Ideologies at Work?," 2023.

McConnell, Christopher, Yotam Margalit, Neil Malhotra, and Matthew Levendusky, "The Economic Consequences of Partisanship in a Polarized Era," *American Journal of Political Science*, 2018, *62* (1), 5–18.

Miano, Armando, "Search Costs, Outside Options, and On-the-Job Search," 2023.

Mummolo, Jonathan and Clayton Nall, "Why Partisans Do Not Sort: The Constraints on Political Segregation," *The Journal of Politics*, 2017, *79* (1), 45–59.

Mutz, Diana C. and Jeffery J. Mondak, "The Workplace as a Context for Cross-Cutting Political Discourse," *The Journal of Politics*, 2006, *68* (1), 140–155.

Non, Arjan, Ingrid Rohde, Andries de Grip, and Thomas Dohmen, "Mission of the company, prosocial attitudes and job preferences: A discrete choice experiment," *Labour Economics*, 2022, *74*, 102087.

Nyhan, Brendan, Jaime Settle, Emily Thorson, Magdalena Wojcieszak, Pablo Barberá, Annie Y. Chen, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Andrew M. Guess, Edward Kennedy, Young Mie Kim, David Lazer, Neil Malhotra, Devra Moehler, Jennifer Pan, Daniel Robert Thomas, Rebekah Tromble, Carlos Velasco Rivera, Arjun Wilkins, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker, "Like-minded sources on Facebook are prevalent but not polarizing," *Nature*, 2023, *620* (7972), 137–144.

Pan, Jessica, "Gender Segregation in Occupations: The Role of Tipping and Social Interactions," *Journal of Labor Economics*, 2015, *33* (2), 365–408.

Panagopoulos, Costas, Donald P. Green, Jonathan Krasno, Michael Schwam-Baird, and Kyle Endres, "Partisan Consumerism: Experimental Tests of Consumer Reactions to Corporate Political Activity," *The Journal of Politics*, 2020, *82* (3), 996–1007.

Perez-Truglia, Ricardo, "Political Conformity: Event-Study Evidence from the United States," *Review of Economics and Statistics*, 2018, *100* (1), 14–28.

Ransom, Michael R., "Sampling Distributions of Segregation Indexes," *Sociological Methods & Research*, 2000, *28* (4), 454–475.

Rice, Anthony B. and Christoph Schiller, "When Values Align: Corporate Philanthropy and Employee Turnover," 2024.

Rosen, Sherwin, "The Theory of Equalizing Differences," *Handbook of Labor Economics*, 1986, *1*, 641–692.

_ , "Markets and Diversity," *American Economic Review*, 2002, *92* (1), 1–15.

Schelling, Thomas C., "Dynamic Models of Segregation," *Journal of Mathematical Sociology*, 1971, *1* (2), 143–186.

Schneider, Florian, Fanny Brun, and Roberto A. Weber, "Sorting and Wage Premiums in Immoral Work," *University of Zurich, Department of Economics Working Paper*, 2020.

Schuh, Rachel, "Miss-Allocation: The Value of Workplace Gender Composition and Occupational Segregation," *Federal Reserve Bank of New York Staff Reports*, 2024.

Shafranek, Richard M., "Political Considerations in Nonpolitical Decisions: A Conjoint Analysis of Roommate Choice," *Political Behavior*, 2021, *43* (1), 271–300.

Steel, Reilly, "The Political Transformation of Corporate America, 2001–2022," 2024.

Strange, Austin M., Ryan D. Enos, Mark Hill, and Amy Lakeman, "Online Volunteer Laboratories for Human Subjects Research," *PloS One*, 2019, *14* (8), e0221676.

Stuckatz, Jan, "Political Alignment Between Firms and Employees in the United States: Evidence From a New Dataset," *Political Science Research and Methods*, 2022, *10* (1), 215–225.

Velez, Yamil R. and Patrick Liu, "Confronting Core Issues: A Critical Test of Attitude Polarization," 2023.

Wang, Kitty and Shijie Lu, "Corporate Political Positioning and Sales: Evidence from a Natural Experiment," 2022.

Williams, Heidi L., "Equalizing wage differentials," 2015.

Wiswall, Matthew and Basit Zafar, "Preference for the Workplace, Investment in Human Capital, and Gender," *The Quarterly Journal of Economics*, 2018, *133* (1), 457–507.

Zingher, Joshua N., "Diploma divide: Educational attainment and the realignment of the American electorate," *Political Research Quarterly*, 2022, *75* (2), 263–277.

A Additional tables and figures

Table A1: Sample Comparison to Cooperative Election Study

	Sample	CES
Birth year	1974.6	1976.6
Male	0.491	0.508
Democrat	0.399	0.391
Republican	0.286	0.273
White	0.799	0.696
Black	0.076	0.112
Hispanic	0.081	0.095
Asian	0.037	0.058
Missing educ.	0.410	0.000
High school	0.039	0.267
Associate's degree	0.062	0.263
Bachelor's degree	0.336	0.295
Graduate degree	0.153	0.175
Num. obs.	34,533,903	12,702

Notes: The 2022 Cooperative Election Study sample is restricted to those who were registered, in places that record party registration, and who did not report that they were retired, homemakers, disabled, or students. We use self-reported party identification and weight the CES using weights for the validated voter sub-sample, so that the target population is registered adults. In our sample, "Democrat" and "Republican" refers to registering with the Democratic or Republican party in the latest observed registration.

Table A2: Coverage by NAICS Sector

Sector	NAICS code	Num. in sample	Share of sample	Share of ACS	Sh. ACS (rewtd.)	Diff. (p.p.)
Professional, Scientific, and Technical Services	54	2,114,030	0.103	0.076	0.127	-2.36
Health Care and Social Assistance	62	2,099,497	0.102	0.145	0.166	-6.40
Manufacturing	31-33	2,087,734	0.102	0.083	0.072	2.98
Finance and Insurance	52	1,977,181	0.096	0.049	0.073	2.30
Educational Services	61	1,821,706	0.089	0.095	0.145	-5.64
Information	51	1,797,491	0.088	0.019	0.027	6.11
Retail Trade	44-45	1,693,074	0.082	0.108	0.073	0.99
Public Administration	92	1,585,545	0.077	0.054	0.065	1.20
Administrative and Support Services	56	901,022	0.044	0.043	0.029	1.45
Other Services (except Public Administration)	81	768,460	0.037	0.048	0.037	-0.01
Accommodation and Food Services	72	748,992	0.036	0.073	0.035	0.11
Real Estate and Rental and Leasing	53	613,632	0.030	0.019	0.021	0.89
Construction	23	542,574	0.026	0.068	0.036	-0.91
Wholesale Trade	42	506,771	0.025	0.024	0.022	0.24
Arts, Entertainment, and Recreation	71	467,075	0.023	0.024	0.023	-0.05
Transportation and Warehousing	48-49	455,375	0.022	0.047	0.029	-0.71
Utilities	22	155,497	0.008	0.008	0.008	-0.01
Mining, Quarrying, and Oil and Gas Extraction	21	123,168	0.006	0.004	0.003	0.35
Agriculture, Forestry, Fishing and Hunting	11	57,863	0.003	0.011	0.007	-0.39
Management of Companies and Enterprises	55	9,808	0.000	0.001	0.002	-0.13

Notes: Sectors are 2-digit NAICS codes. "Share of ACS" refers to the proportion of individuals in the 2017-2022 ACS who are in the labor force and live in a place that records party affiliation who work in the given sector. We also show industry shares re-weighting the ACS sample to match our analysis sample according to the proportions of individuals with a high school, associate's, bachelor's, and graduate degree.

Table A3: Coverage by SOC Occupation

Occupation title	SOC code	Num. in sample	Share sample	Share ACS	Sh. ACS (rewtd.)	Diff. (p.p.)
Management	11	6,778,091	0.282	0.106	0.151	13.14
Office and Administrative Support	43	2,617,227	0.109	0.111	0.090	1.86
Business and Financial Operations	13	2,015,396	0.084	0.056	0.096	-1.23
Sales	41	1,699,996	0.071	0.099	0.086	-1.54
Computer and Mathematical	15	1,497,399	0.062	0.032	0.054	0.83
Healthcare Practitioners and Technical	29	1,438,259	0.060	0.062	0.097	-3.74
Educational Instruction and Library	25	1,432,930	0.060	0.063	0.107	-4.73
Arts, Design, Entertainment, Sports, and Media	27	962,507	0.040	0.021	0.034	0.60
Architecture and Engineering	17	667,388	0.028	0.018	0.031	-0.30
Community and Social Service	21	640,897	0.027	0.018	0.031	-0.43
Transportation and Material Moving	53	618,639	0.026	0.075	0.032	-0.66
Production	51	526,976	0.022	0.048	0.021	0.06
Food Preparation and Serving	35	511,195	0.021	0.056	0.023	-0.17
Healthcare Support	31	421,501	0.018	0.034	0.019	-0.17
Installation, Maintenance, and Repair	49	397,570	0.017	0.031	0.014	0.29
Legal	23	353,705	0.015	0.013	0.021	-0.66
Life, Physical, and Social Science	19	352,169	0.015	0.011	0.020	-0.53
Protective Service	33	348,793	0.015	0.023	0.020	-0.50
Personal Care and Service	39	312,013	0.013	0.028	0.018	-0.49
Construction and Extraction	47	268,350	0.011	0.051	0.018	-0.69
Building and Grounds Cleaning and Maintenance	37	123,367	0.005	0.038	0.014	-0.84
Farming, Fishing, and Forestry	45	21,205	0.001	0.005	0.002	-0.12

Notes: Occupations are SOC "major" 2-digit occupation codes. "Share of ACS" refers to the proportion of individuals in the 2017–2022 ACS who are in the labor force and live in a place that records party affiliation with the given occupation. We also show occupation shares re-weighting the ACS sample to our analysis sample according to the proportions of individuals with a high school, associate's, bachelor's, and graduate degree.

Table A4: Summary Statistics by Group

	Democrat	Other	Republican	Has coworker	DIME
Birth year	1974.9	1977.7	1970.7	1976.8	1965.8
Male	0.423	0.527	0.545	0.483	0.521
White	0.707	0.802	0.923	0.792	0.891
Black	0.142	0.054	0.011	0.077	0.040
Hispanic	0.101	0.088	0.045	0.085	0.036
Asian	0.044	0.048	0.017	0.040	0.029
Democrat	1.000	0.000	0.000	0.408	0.594
Republican	0.000	0.000	1.000	0.274	0.219
High school	0.056	0.068	0.078	0.055	0.028
Associate's degree	0.093	0.111	0.119	0.098	0.052
Bachelor's degree	0.560	0.573	0.582	0.568	0.517
Graduate degree	0.292	0.248	0.221	0.279	0.403
Missing field of degree	0.607	0.609	0.655	0.544	0.599
Missing industry	0.336	0.362	0.320	0.328	0.198
Missing occupation	0.284	0.270	0.265	0.134	0.309
Has DIME score	0.164	0.066	0.084	0.103	1.000
Num. connections	136.8	133.6	123.3	154.2	189.9
Num. jobs	3.2	3.2	2.8	3.8	3.5
Block group HH income	89,739	91,396	89,841	91,263	104,895
Lists employer in 2022	0.627	0.631	0.656	0.755	0.640
At least one coworker	0.600	0.593	0.565	1.000	0.550
Share Dem. coworkers	0.632	0.230	0.192	0.382	0.478
Share Rep. coworkers	0.140	0.151	0.600	0.260	0.234
Num. coworkers	374.7	324.4	243.2	446.4	357.4
Num. obs. (millions)	13.787	10.884	9.862	20.293	3.810

Notes: This table shows means of covariates for five subsamples of the analysis data. "Democrat" and "Republican" refer to the subsamples who registered with the Democratic or Republican party in their latest observed registration, while "Other" refers to those who registered with another party, no party, or as an independent. "DIME" refers to the subsample that we link to the DIME database and who have nonmissing DIME score. Number of connections is topcoded at 500.

Table A5: Partisan Sorting Accounting for Covariates

	σ_g^X										
	No controls	CZ	+ Demo.	+ Coll.	+ Major	+ Occ.	+ Ind.	Num. groups	Num. individuals	Mean size	Median size
College	0.215	0.137	0.124	–	–	–	–	3,970	16,987,677	4,279	502
Major	0.148	0.127	0.117	0.091	–	–	–	186	13,069,305	70,265	26,550
Occupation	0.122	0.111	0.085	0.077	0.067	–	–	776	25,053,834	32,286	4,793
Industry	0.151	0.128	0.107	0.095	0.085	0.068	–	145	23,072,098	159,118	78,729
Employer	0.270	0.214	0.190	0.177	0.171	0.162	0.155	530,099	20,343,829	38	9
County	0.253	0.146	0.116	0.104	0.102	0.101	0.098	3,099	34,533,815	11,144	490
ZIP code	0.296	0.213	0.169	0.155	0.153	0.150	0.146	27,490	34,523,842	1,256	127
Block group	0.315	0.239	0.190	0.177	0.175	0.172	0.167	178,079	34,436,545	193	139

Notes: This table shows the standard deviation of the partisan gap, weighted by group size, across various groups. We include only groups (e.g., companies) with more than 5 people. We also show the standard deviation after residualizing partisanship by a series of fixed effects, which are included cumulatively. When we residualize, we include those missing a particular covariate (e.g., college) in a separate category. College refers to the IPEDS identifier of the associate or bachelor's degree school; major refers to the detailed field of degree as recorded in the ACS; occupation refers to a 6-digit SOC code; industry refers to the industry as reported on the individual's profile; and employer refers to the name of the employer that the individual lists on their profile.

Table A6: College Selectivity and the Partisan Gap

	Admission rate		Avg. SAT score	
	(1)	(2)	(3)	(4)
Partisan gap	0.425***		-162.9***	
	(0.053)		(38.5)	
Partisan gap (resid.)		0.595***		-454.9***
		(0.084)		(57.3)
Observations	1,806	1,806	1,394	1,394
R^2	0.200	0.133	0.067	0.172
Dep. var. mean	0.663	0.663	1,065.3	1,065.3
Dep. var. std. dev.	0.206	0.206	132.9	132.9

Notes: This table shows regressions of college-level characteristics – the admission rate and the average SAT score of admitted students – on the partisan gap, or the share of Republicans minus the share of Democrats. Regressions are weighted by the number of individuals in our sample from the college. We also show estimates after residualizing individual partisanship using CZ fixed effects and interactions of year of birth, gender, and race. Average SAT score and admission rate are from the College Scorecard and refer to the 2013-2014 cohort. Standard errors are heteroskedasticity-robust. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table A7: Summary Statistics for One-Time Movers

	All N = 1,109,019		Moved before 2016 N = 427,376		Moved after 2016 N = 541,453		Uncommitted in 2012 N = 584,574		Partisan in 2012 N = 524,445	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Male	0.511	0.500	0.509	0.500	0.513	0.500	0.519	0.500	0.503	0.500
Birth year	1977.9	11.7	1976.9	11.8	1978.7	11.6	1980.1	10.8	1975.5	12.2
White	0.763	0.425	0.772	0.419	0.755	0.430	0.746	0.435	0.781	0.414
Registered in 2012	0.720	0.449	0.739	0.439	0.705	0.456	0.469	0.499	1.000	0.000
Voted in 2012	0.616	0.486	0.634	0.482	0.601	0.490	0.447	0.497	0.804	0.397
Dem. in 2012	0.275	0.447	0.286	0.452	0.266	0.442	0.000	0.000	0.581	0.493
Rep. in 2012	0.198	0.398	0.204	0.403	0.194	0.395	0.000	0.000	0.419	0.493
Registered in 2016	0.788	0.409	0.799	0.401	0.779	0.415	0.694	0.461	0.892	0.310
Voted in 2016	0.716	0.451	0.732	0.443	0.702	0.458	0.629	0.483	0.810	0.392
Dem. in 2016	0.351	0.477	0.358	0.479	0.344	0.475	0.172	0.377	0.549	0.498
Rep. in 2016	0.236	0.425	0.242	0.428	0.231	0.422	0.094	0.292	0.394	0.489
Registered in 2020	0.892	0.311	0.888	0.315	0.894	0.308	0.879	0.326	0.905	0.293
Voted in 2020	0.750	0.433	0.751	0.433	0.749	0.434	0.722	0.448	0.781	0.414
Dem. in 2020	0.410	0.492	0.412	0.492	0.409	0.492	0.303	0.459	0.530	0.499
Rep. in 2020	0.253	0.435	0.257	0.437	0.250	0.433	0.157	0.364	0.360	0.480
High school	0.043	0.203	0.049	0.216	0.038	0.192	0.046	0.209	0.040	0.196
Associate's degree	0.096	0.295	0.101	0.302	0.093	0.290	0.101	0.302	0.091	0.288
Bachelor's degree	0.549	0.498	0.548	0.498	0.549	0.498	0.553	0.497	0.544	0.498
Graduate degree	0.312	0.463	0.301	0.459	0.320	0.467	0.300	0.458	0.325	0.468
Num. connections	228.9	187.9	221.7	187.5	233.6	188.4	222.7	187.9	234.6	188.0
Origin workplace Dem. share	0.351	0.135	0.352	0.136	0.349	0.133	0.342	0.132	0.361	0.137
Dest. workplace Dem. share	0.359	0.137	0.359	0.136	0.359	0.137	0.352	0.134	0.367	0.140
Origin workplace Rep. share	0.237	0.130	0.237	0.132	0.238	0.130	0.227	0.126	0.249	0.135
Dest. workplace Rep. share	0.242	0.134	0.241	0.133	0.244	0.135	0.233	0.130	0.253	0.138
Num. in origin workplace	843.1	1819.0	775.7	1742.5	897.5	1877.9	871.3	1866.5	813.1	1767.0
Num. in dest. workplace	868.9	2000.1	938.6	2125.9	802.7	1873.5	905.4	2064.5	814.9	1902.2
Linked to contributions	0.144	0.351	0.150	0.357	0.139	0.346	0.114	0.318	0.178	0.382

Notes: This table shows summary statistics for one-time movers between workplaces with at least 10 individuals (among the group of non-one-time-movers) observed working there over the 2012 to 2021 period. We show statistics for the entire group; then for those who moved before 2016 and those who moved after 2016; then for those who were uncommitted (independent or unregistered) in 2012 and those who were registered partisans in 2012.

Table A8: Effect of the Workplace on Partisanship in 2016 – Effects on Registration with Each Party

	Democrat (1)	Republican (2)
Moved before × dest. exposure	0.044***	0.055***
	(0.009)	(0.008)
Dest. workplace exposure	0.086***	0.061***
	(0.007)	(0.006)
Origin workplace exposure	0.106***	0.098***
	(0.006)	(0.005)
Individual controls	✓	✓
Origin × dest. controls	✓	✓
Observations	545,326	545,326
R^2	0.165	0.146
Dependent variable mean	0.173	0.093
Dependent variable std. dev.	0.378	0.291

Notes: This table shows regression results corresponding to column (1) of Table 2, separately for registration with the Democratic and Republican Party. Standard errors are clustered by individual. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table A9: Effect of the Workplace on Partisanship in Each Election, 2014–2020

	2014	2016	2018	2020	General	Presidential
	(1)	(2)	(3)	(4)	(5)	(6)
Moved before × dest. exposure	0.033***	0.050***	0.054***	0.050***	0.048***	0.050***
	(0.007)	(0.006)	(0.007)	(0.012)	(0.005)	(0.006)
Dest. workplace exposure	0.045***	0.074***	0.089***	0.108***		
	(0.003)	(0.005)	(0.007)	(0.012)		
Origin workplace exposure	0.056***	0.102***	0.121***	0.145***		
	(0.003)	(0.004)	(0.005)	(0.005)		
Individual controls	✓	✓	✓	✓	✓	✓
Origin × dest. controls	✓	✓	✓	✓	✓	✓
Dest. workplace exposure × Outcome year					✓	✓
Origin workplace exposure × Outcome year					✓	✓
Observations	1,085,606	1,090,652	1,110,324	1,141,098	4,427,680	2,231,750
R^2	0.163	0.169	0.183	0.206	0.208	0.204
Dependent variable mean	0.068	0.133	0.175	0.230	0.152	0.183
Dependent variable std. dev.	0.251	0.340	0.380	0.421	0.360	0.386

Notes: This table shows regression results corresponding to column (1) of Table 2, except in columns (1) through (4), the outcome variable is the individual's partisanship in each general election year between 2014 and 2020, and the independent variable of interest is whether the individual moved between workplaces before that year versus afterwards. Column (5) pools all election years and column (6) pools the two presidential election years (2016 and 2020). These estimates correspond to Eq. 3. Standard errors are clustered by individual. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table A10: Effect of the Workplace on Partisanship in 2016 – Split-Sample IV Results

	SSIV		OLS	
	(1)	(2)	(3)	(4)
Dest. workplace exposure	0.187***	0.211***	0.074***	0.037***
	(0.018)	(0.019)	(0.005)	(0.003)
Moved before × dest. exposure	0.071***	0.069***	0.050***	0.023***
	(0.011)	(0.012)	(0.006)	(0.004)
Origin workplace exposure	0.251***	0.232***	0.102***	0.048***
	(0.017)	(0.018)	(0.004)	(0.002)
Individual controls	✓	✓	✓	✓
Origin × dest. controls	✓	✓	✓	✓
Workplace size restriction	✓		✓	
Observations	1,163,068	1,495,142	1,164,038	1,677,700
R^2	0.218	0.194	0.224	0.203
Dependent variable mean	0.134	0.134	0.134	0.134
Dependent variable std. dev.	0.341	0.343	0.340	0.341

Notes: This table shows regression results corresponding to column (1) of Table 2, except that in columns (3) and (4), we instrument for origin and destination workplace partisanship using a split-sample instrument. See the text for details. Columns (1) and (3) include only workplaces where we observe more than 10 individuals, while columns (2) and (4) include all workplaces. Standard errors are clustered by individual. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table A11: Willingness to Pay for Liberal Framing

	Democrat		Other	Republican	
	All	Strong		All	Strong
	(1)	(2)	(3)	(4)	(5)
10th percentile	-0.035	-0.028	-0.108	-0.115	-0.156
	(0.006)	(0.007)	(0.025)	(0.014)	(0.099)
25th	-0.003	0.004	-0.053	-0.069	-0.094
	(0.003)	(0.004)	(0.015)	(0.008)	(0.048)
50th	0.029	0.035	-0.000	-0.024	-0.036
	(0.003)	(0.004)	(0.009)	(0.005)	(0.015)
75th	0.060	0.065	0.049	0.019	0.019
	(0.005)	(0.007)	(0.011)	(0.005)	(0.014)
90th	0.090	0.095	0.096	0.059	0.071
	(0.007)	(0.011)	(0.018)	(0.009)	(0.029)
Inattention rate	0.134	0.175	0.054	0.178	0.168
	(0.024)	(0.038)	(0.052)	(0.031)	(0.044)
Observations	8,816	4,788	3,515	6,175	2,945
Num. individuals	464	252	185	325	155

Notes: This table shows estimated quantiles of WTP for the liberal framing from the model in Eq. 4, estimated via maximum likelihood, separately for (1) all Democrats, (2) strong Democrats, (3) those who are neither Democrats nor Republicans, (4) all Republicans, and (5) strong Republicans. These estimates are corrected for inattention, and we also show the estimated inattention rate for each group. Bootstrapped standard errors based on 500 samples are in parentheses.

Table A12: Framing Experiment: Heterogeneity

	(1)
Log wage diff.	3.76***
	(0.079)
Democrat	0.105***
	(0.019)
Republican	-0.079***
	(0.021)
Female	0.036***
	(0.014)
Asian/Asian Am.	-0.011
	(0.023)
Black/Af. Am.	0.052***
	(0.018)
Hispanic/Latino	-0.052*
	(0.028)
Middle Eastern	0.013
	(0.066)
Other race	-0.011
	(0.045)
Born 1960s	-0.041
	(0.064)
Born 1970s	-0.012
	(0.056)
Born 1980s	-7.06×10^{-5}
	(0.054)
Born 1990s	0.003
	(0.054)
Born 2000s	-0.182***
	(0.057)
Bachelor's deg.	-0.025
	(0.018)
Graduate deg.	-0.010
	(0.021)
Log current salary	-0.002
	(0.011)
Observations	15,536

Notes: This table shows coefficient estimates from a regression of whether the respondent chose the job with the liberal framing on their demographics, controlling for the wage difference between the jobs. We exclude observations where both jobs use the same framing. The omitted category for party is "Independent," for race is "White," for year of birth is "Born before 1960," and for education is "No bachelor's degree." Standard errors are clustered by individual. ***, **, and * indicate significance at 1%, 5%, and 10% levels.

Table A13: Willingness to Pay for Share of Democrats Predicted by Company Description

				Democrat		Other	Republican	
Algorithm	Type	MAE	Framing Dif.	All	Strong		All	Strong
SVM	Words	0.109	0.102	0.047	0.053	0.000	-0.043	-0.063
SVM	Stems	0.107	0.102	0.047	0.054	-0.005	-0.046	-0.067
SVM	Embeddings	0.098	0.102	0.049	0.058	-0.005	-0.043	-0.065
LightGBM	Words	0.100	0.061	0.054	0.065	-0.005	-0.066	-0.103
LightGBM	Stems	0.098	0.061	0.058	0.073	-0.011	-0.064	-0.101
LightGBM	Embeddings	0.095	0.060	0.059	0.067	-0.005	-0.047	-0.079

Notes: This table shows WTP estimates for a company whose description signals a share of Democrats 1 SD higher using various algorithms to predict the partisanship of the synthetic job ads. We present results using a linear support vector machine as well as LightGBM, using as predictors the 3,000 most frequent words, the 3,000 most frequent word stems, and embedding vectors generated by BERT. We include industry as a categorical predictor in all models. We train these models on the 299,562 company descriptions in our data that are associated with at least 10 people in our sample and are at least 100 characters long. We present the mean absolute error of each algorithm on a held-out 10% test sample, as well as the average difference between the share of Democrats predicted for our synthetic liberal versus neutral ads. Then, we show the median WTP implied by each of these algorithms for an ad that signals a 1 SD higher share of Democrats (24.2 p.p.), separately by party affiliation.

Table A14: Willingness to Pay for Coworker and Job Characteristics

	Democrat		Other	Republican		Difference	
	All (1)	Strong (2)	(3)	All (4)	Strong (5)	(1) – (4)	(2) – (5)
Share Democrat	0.033	0.041	-0.014	-0.040	-0.050	0.073	0.091
	(0.005)	(0.006)	(0.004)	(0.007)	(0.012)	(0.009)	(0.013)
Share nonwhite	0.006	0.004	0.001	-0.015	-0.024	0.020	0.028
	(0.006)	(0.006)	(0.005)	(0.006)	(0.009)	(0.008)	(0.011)
Share women	0.006	0.008	-0.009	-0.016	-0.018	0.021	0.025
	(0.006)	(0.007)	(0.005)	(0.006)	(0.008)	(0.009)	(0.011)
Average age	-0.006	0.001	-0.007	-0.000	0.003	-0.006	-0.001
	(0.006)	(0.006)	(0.005)	(0.005)	(0.007)	(0.008)	(0.009)
10-50 employees	0.028	0.020	0.018	0.018	0.031	0.010	-0.011
	(0.012)	(0.014)	(0.011)	(0.014)	(0.019)	(0.018)	(0.024)
50-200 employees	0.017	0.020	0.022	0.007	0.022	0.011	-0.003
	(0.012)	(0.013)	(0.011)	(0.012)	(0.015)	(0.017)	(0.020)
Contributes to Democrats	0.007	0.023	-0.035	-0.076	-0.060	0.083	0.083
	(0.013)	(0.013)	(0.012)	(0.017)	(0.022)	(0.021)	(0.025)
Contributes to Republicans	-0.100	-0.097	-0.038	0.011	0.036	-0.111	-0.132
	(0.018)	(0.019)	(0.011)	(0.013)	(0.017)	(0.022)	(0.025)
Industry share Dem.	0.009	0.014	-0.002	-0.013	-0.009	0.022	0.023
	(0.007)	(0.009)	(0.007)	(0.010)	(0.016)	(0.013)	(0.018)
401k company match	0.061	0.055	0.069	0.074	0.070	-0.013	-0.015
	(0.015)	(0.017)	(0.013)	(0.017)	(0.021)	(0.023)	(0.027)
Medical, vision, and dental insurance	0.150	0.153	0.077	0.104	0.107	0.046	0.046
	(0.019)	(0.023)	(0.016)	(0.019)	(0.026)	(0.027)	(0.035)
Mental and emotional health benefits	0.032	0.025	0.019	0.010	-0.001	0.022	0.026
	(0.015)	(0.017)	(0.014)	(0.018)	(0.023)	(0.023)	(0.029)
Tuition reimbursement	-0.013	-0.004	-0.008	-0.017	-0.018	0.004	0.015
	(0.014)	(0.015)	(0.014)	(0.018)	(0.023)	(0.023)	(0.028)
Observations	1738	1309	1815	1452	902		
Num. individuals	158	119	165	132	82		

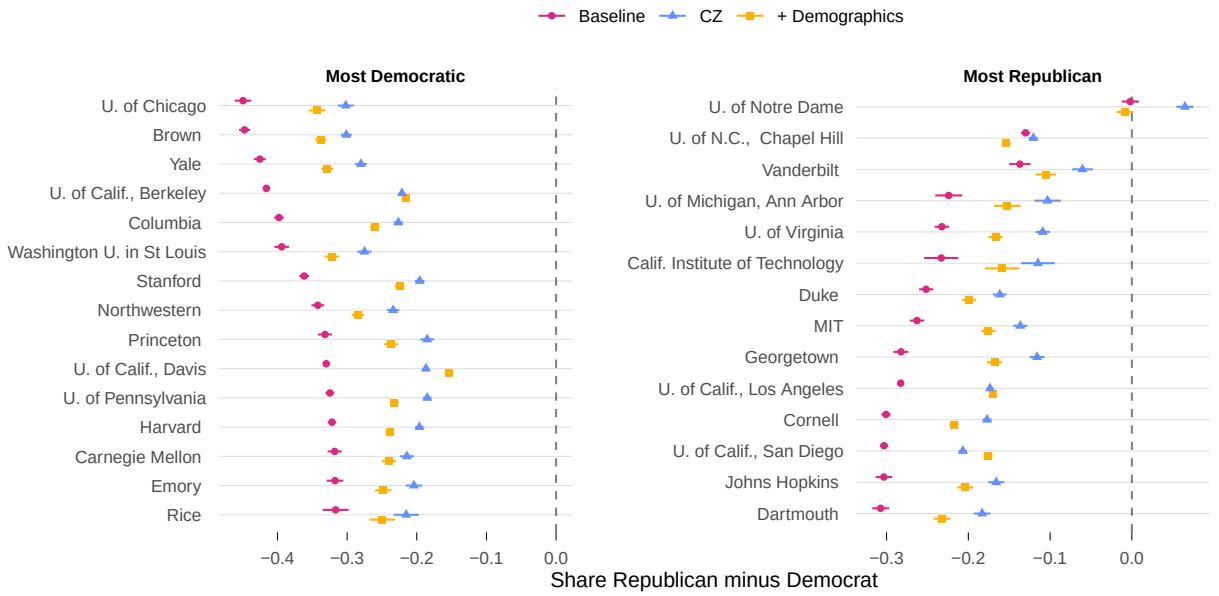
Notes: This table shows estimates of the median willingness to pay (expressed as a fraction of the salary) for coworker and job characteristics, separately by party, estimated from the conjoint experiment in Section 5.5. The final two columns show the difference in WTP between Democrats and Republicans and between strong Democrats and strong Republicans, respectively. Coworker demographic variables are standardized, so the estimate corresponds to the respondent's WTP for a job where the share of Democrats, women, or nonwhite coworkers is 1 SD higher or the average age is 1 SD higher. The industry Democrat share is also standardized, so the estimate corresponds to the respondent's WTP for a job in an industry where the share of Democrats is 1 SD higher. The reference category for corporate contributions is "No political contributions," for company size is "200+ employees," and for benefits is "Life insurance." The covariance estimates for the regression coefficients are clustered by individual, and standard errors for the WTP estimates are computed using the delta method. Horizontal bars are 95% confidence intervals.

Figure A1: Most Republican and Democratic Colleges



Notes: The figure shows, among the 100 largest colleges in the sample, the 30 most Republican and 30 most Democratic, as measured by the partisan gap: the share of Republicans minus the share of Democrats. The baseline estimates (without controls) are shown in magenta, and are relative to the partisan shares for all others in the sample. The blue series includes fixed effects for the individual's CZ and the yellow series adds fixed effects for interactions of year of birth, gender, and race. Horizontal bars are 95% confidence intervals.

Figure A2: Partisanship of Top Colleges



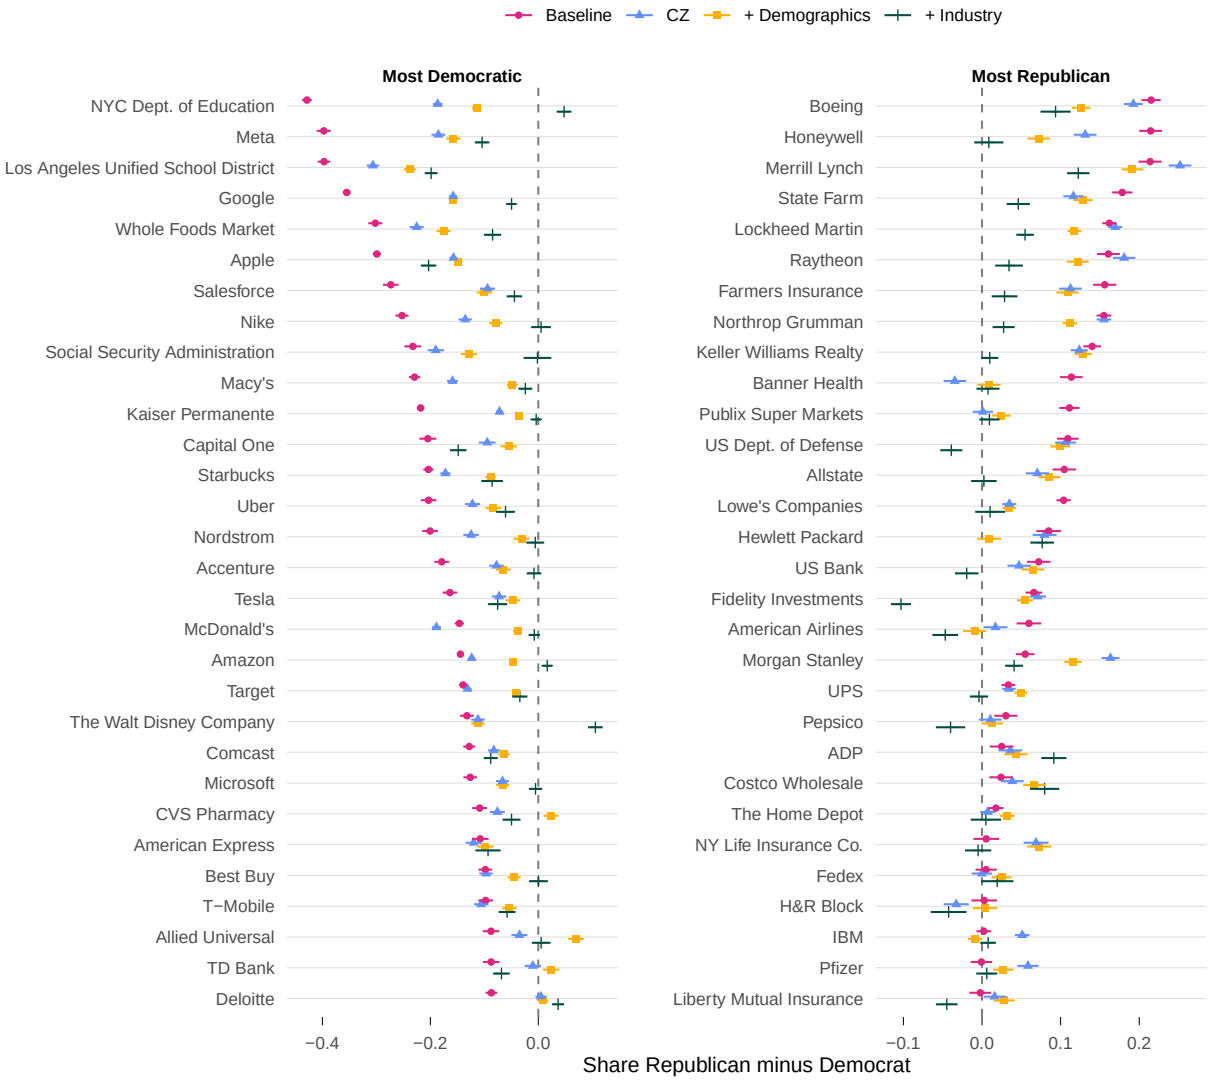
Notes: The figure shows the estimates for Republican and Democratic partisanship among the 30 highest-ranked schools in the 2024 U.S. News & World Report National University rankings. The baseline estimates (without controls) are shown in magenta, and are relative to the partisan shares for all others in the sample. The blue series includes fixed effects for the individual's CZ and the yellow series adds fixed effects for interactions of year of birth, gender, and race. Horizontal bars are 95% confidence intervals.

Figure A3: Most Republican and Democratic Majors, Occupations, and Industries



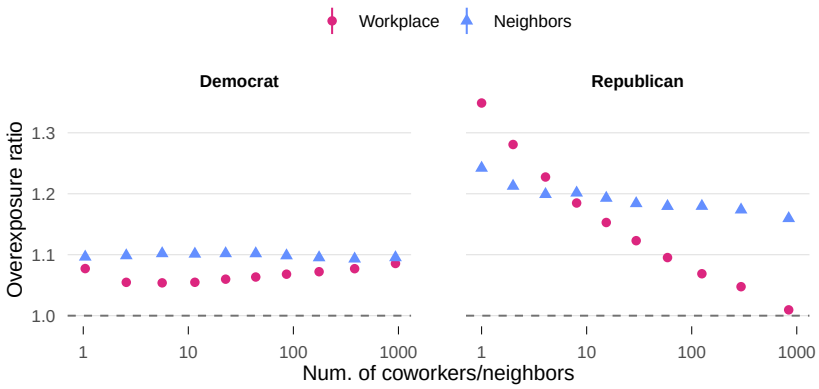
Notes: The figures show, among the 50 largest college majors, industries, and occupations in the sample, the 10 most Republican and 10 most Democratic, as measured by the partisan gap: the share of Republicans minus the share of Democrats. Major refers to the detailed field of degree as recorded in the ACS. Industry refers to the industry as reported on the individual profile, *not* the industry associated with their employer. Occupations are 5-digit "broad" SOC codes. The baseline estimates (without controls) are shown in magenta, and are relative to all individuals in our sample who are not included in one of the 50 largest majors (in the first row), occupations (in the second), or industries (in the third). The blue series includes fixed effects for the individual's CZ and the yellow series adds fixed effects for interactions of year of birth, gender, and race. Horizontal bars are 95% confidence intervals.

Figure A4: Most Republican and Democratic Employers



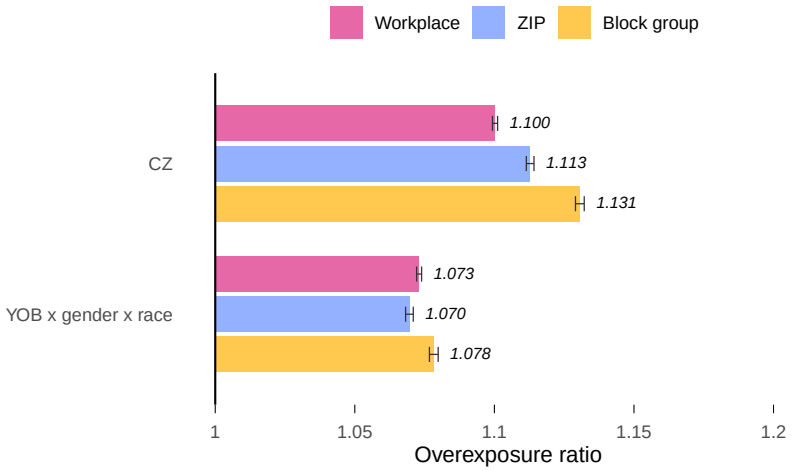
Notes: The figure shows, among the 100 largest employers in the sample, the 30 most Republican and 30 most Democratic, as measured by the partisan gap: the share of Republicans minus the share of Democrats. We exclude universities and the military. The baseline estimates (without controls) are shown in magenta, and are relative to the partisan shares at all other employers. The blue series includes fixed effects for the individual's CZ, the yellow series adds fixed effects for interactions of year of birth, gender, and race, and the green series adds fixed effects for industry. Horizontal bars are 95% confidence intervals.

Figure A5: Partisan Segregation by Number of Coworkers and Neighbors



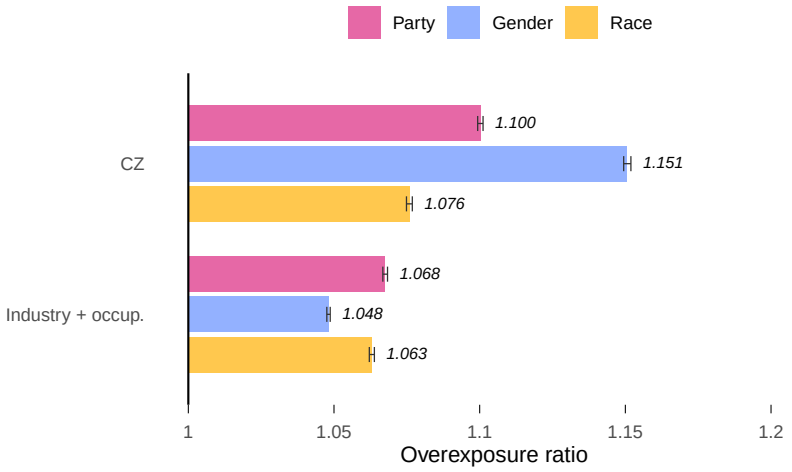
Notes: The figure shows the relationship between partisan segregation and the size of the comparison group – coworkers and neighbors – separately for Democrats and Republicans. For each Democrat and Republican, we first compute their share of co-partisan coworkers and the expected share of co-partisan coworkers based on their CZ. Then, we compute the share of co-partisans among the same number of geographically closest neighbors (in our matched sample) and the expected share of co-partisan neighbors, again based on CZ. Finally, we compute 10 equally-sized bins of the individual's number of coworkers, separately for Democrats and Republicans, and take the means of the overexposure ratio within each bin, which we show on the vertical axis for coworkers (in magenta) and neighbors (in blue). The horizontal axis, which is logarithmic, shows the number of coworkers or neighbors. We drop individuals with more than 1,500 coworkers. Neighbors must be in the same CZ as the focal individual.

Figure A6: Comparing Partisan Segregation at the Workplace, ZIP Code, and Block Group



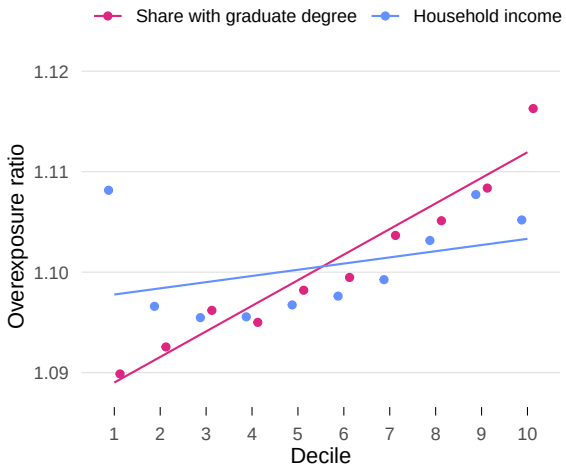
Notes: The figure compares the partisan overexposure ratio at the workplace, ZIP code, and Census block group. The first row compares the own-party share of coworkers or neighbors to the expected share based on CZ. In the second row, we predict party with fixed effects for CZ and the interaction of year of birth, gender, and race. Standard errors are computed using the bootstrap with 50 replications, and horizontal bars are 95% confidence intervals.

Figure A7: Comparing Partisan Segregation to Racial and Gender Segregation at the Workplace



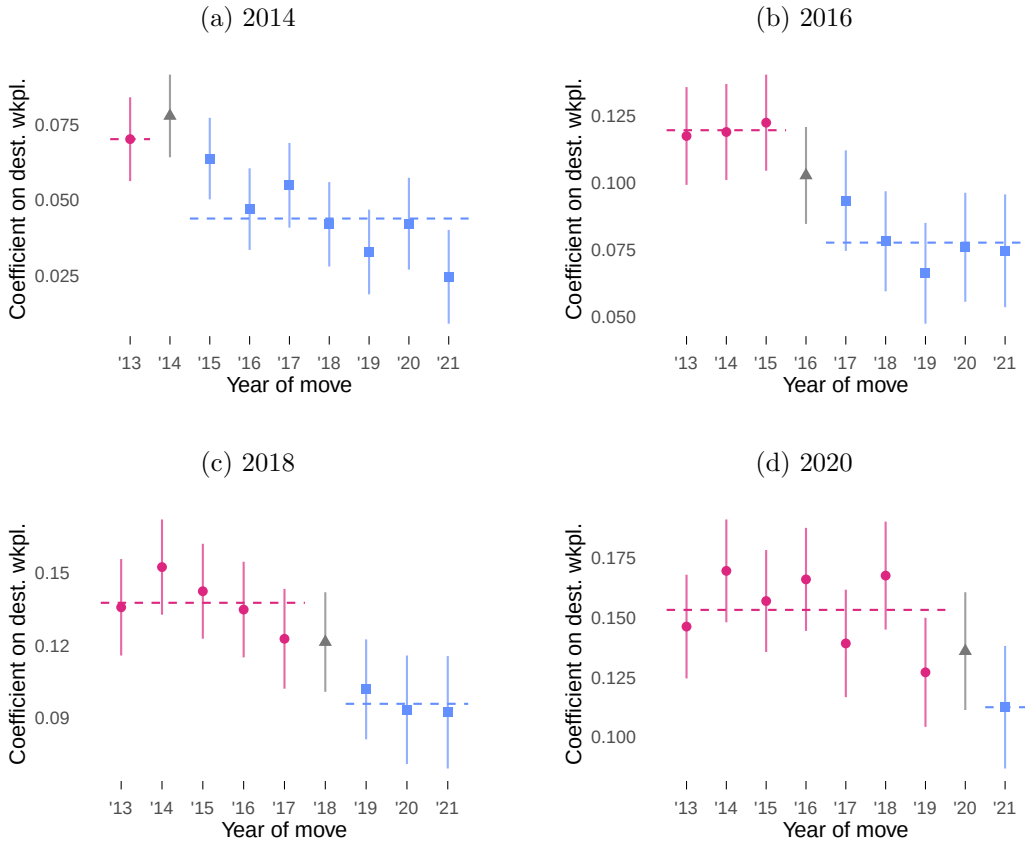
Notes: The figure compares the partisan overexposure ratio to the gender and racial overexposure ratios at the workplace. The first row compares the own-party, own-race, or own-gender shares of coworkers to the expected shares based on CZ. In the second row, we predict party, race, and gender with CZ, industry, and occupation fixed effects. Missing values are a separate category (i.e., we do not drop observations with missing values). We use the industry reported on the individual's profile, and occupation is the 5-digit SOC code associated with their job title. Standard errors are computed using the bootstrap with 50 replications, and horizontal bars are 95% confidence intervals.

Figure A8: Workplace Partisan Segregation by Tract-Level Education and Income



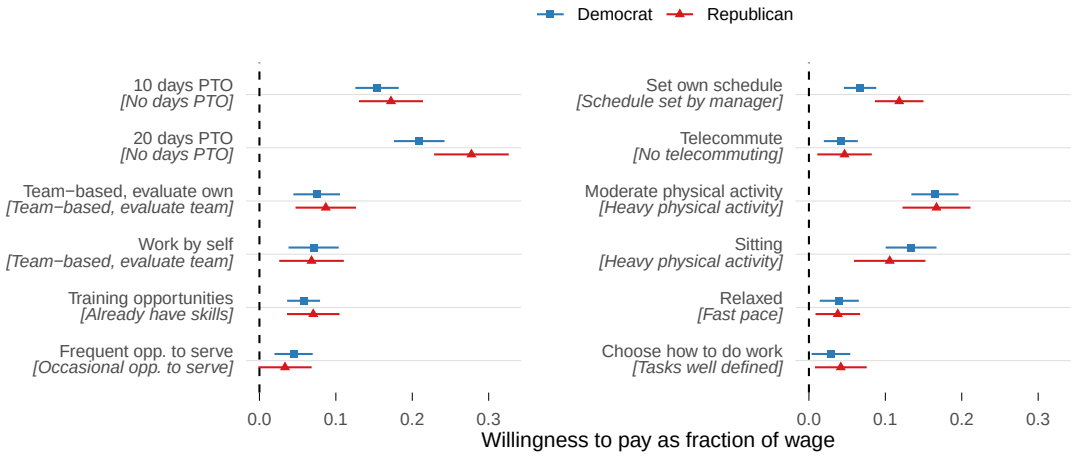
Notes: The figure shows the own-party overexposure ratio, computed separately for individuals in 10 deciles of tract-level education (share with a graduate degree) or income (median household income). To compute overexposure, we compare actual own-party exposure to exposure predicted by coworkers' CZs. Tract-level data is from the 2016-2021 ACS.

Figure A9: Effect of the Workplace on Partisanship, By Year of Move and Election Year



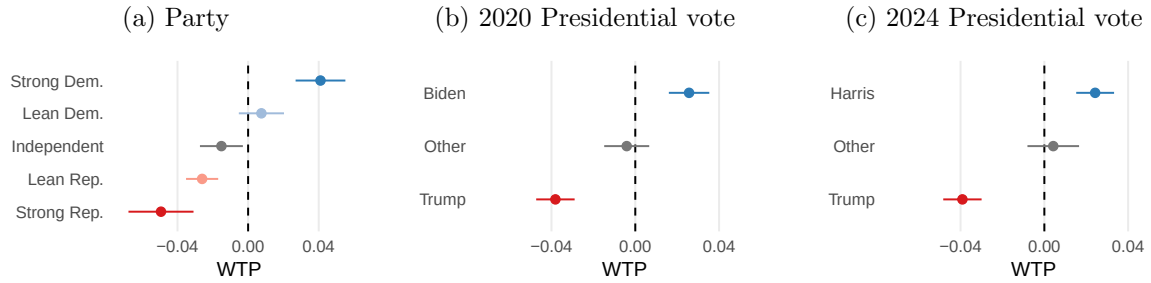
Notes: The figure shows estimates corresponding to Eq. 1 of the relationship between the destination workplace partisanship and the worker's own partisanship measured in 2014, 2016, 2018, or 2020 (in Panels (a) to (d), respectively). Each coefficient corresponds to individuals who moved in a particular year. The magenta coefficients are estimated for those who moved before partisanship is measured and the blue coefficients are for those who moved after. Vertical bars are 95% confidence intervals computed using standard errors clustered by individual.

Figure A10: Willingness to Pay for "Canonical" Amenities by Party



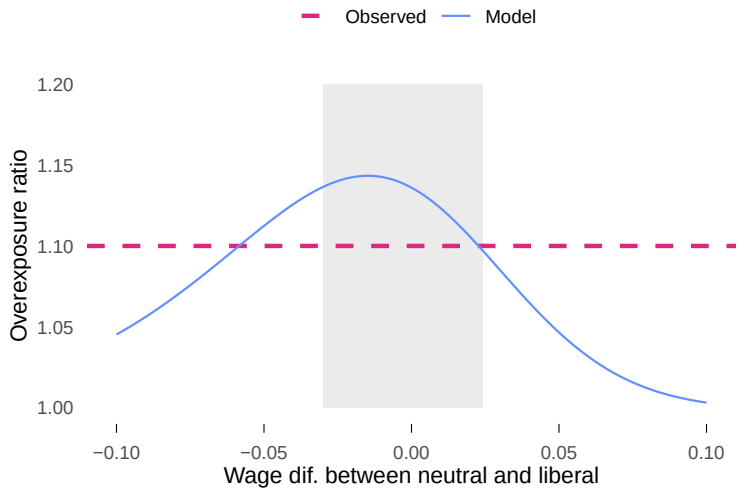
Notes: The figure shows results of our reanalysis of stated-preference experiments in the second wave of the American Working Conditions Survey, fielded December 2015 to February 2016 on the RAND ALP and originally reported in Maestas et al. (2023). We combine this data with information about party affiliation from a separate survey about midterm voting intentions (survey 406 on the RAND ALP) fielded between September and October 2014 to some of the same respondents. The sample consists of 323 Republicans and 505 Democrats. On the horizontal axis, we show respondents' willingness to pay for each of the 9 job attributes studied by Maestas et al. (2023), expressed as a fraction of the respondent's current wage, separately for Democrats and Republicans. Below each attribute, we report in italics the "baseline" level of that characteristic. The regression specification is the same as in that paper. Horizontal bars are 95% confidence intervals computed using standard errors clustered by individual.

Figure A11: Willingness to Pay for Democratic Coworkers by Party and Presidential Vote



Notes: The figures show estimates from the conjoint experiment described in Section 5.5 of the average willingness to pay (expressed as a fraction of the salary) for a company with a share of Democratic coworkers 1 SD higher. We show these separately by party identification in Panel (a), reported 2020 Presidential vote in Panel (b), and intended 2024 Presidential vote in Panel (c). The covariance estimates for the regression coefficients are clustered by individual, and standard errors for the WTP estimates are computed using the delta method. Horizontal bars are 95% confidence intervals.

Figure A12: Calibration of Rosen Model Using Survey WTP Estimates



Notes: The figure shows the overexposure ratio versus the market compensating differential (the log salary difference between the neutral and the liberal job). We calculate the overexposure ratio implied by the Rosen model with a binary amenity, plugging in the survey WTP estimates separately for Democrats, Republicans, and others, and comparing the shares of partisans in each type of job to the CZ-level partisan shares. The blue series shows the model estimates, the magenta dashed line shows S_{cz} measured in the observational data, and the shaded area is the range between the median Democrat and the median Republican's WTP from the framing experiment.

B Data and matching

This appendix describes how we process the LinkedIn data and match it to voter file and campaign finance records.

B.1 Processing LinkedIn data

College: We process the colleges reported on LinkedIn by matching them to an IPEDS identifier. First, we drop school names that correspond to high schools. Then, we use approximate string matching (Jaro-Winkler distance less than 0.03) to find the best match between each school name and the 2021 IPEDS list of institution names and aliases. We then manually check and edit this mapping.

School type: We first exclude educational experiences where the description indicates that the individual did not graduate (contains "attended," "took classes," "did not complete," etc.). For each education item, we first check whether the description contains "high school" or similar. We then check whether the experience corresponds to an associate's, bachelor's, or graduate degree by looking for a degree title (e.g., "B.A." or "bachelor of arts") in the description. If these are not present but we matched the school to an IPEDS identifier, we check whether school offers bachelor's degrees (`HLOFFER` ≥ 5) or associate's degrees (`HLOFFER` ≥ 3) to infer the school type. We use these classifications to infer the individual's highest level of education.

Major: For each educational experience corresponding to an associate or bachelor's degree, we assign a major using approximate string matching (Jaro-Winkler distance less than 0.15) to the list of degree fields (`DEGFIELDD`) from the ACS and CIP codes from the NCES. We then manually check and edit this mapping.

Location: We process the locations reported on LinkedIn by first attempting to place them within a state using string matching. We then further match them to a Census place, county, or CBSA, using approximate string matching (Jaro-Winkler distance less than 0.1).

Occupation: We assign each reported job title to an SOC code using approximate string matching (Jaro-Winkler distance less than 0.15) to the O*NET list of alternate titles. We then manually check and edit this mapping.

B.2 Creating employment histories

Our goal is to assign each individual to a job for each year from 2012 to 2022. To do this, we first exclude jobs listed with an employer such as "self-employed," "freelancer", "independent contractor," etc. We remove consecutive duplicates of the same title, employer, and location, keeping the start date of the first and the end date of the last position.

Each position listed on LinkedIn can include a start and end year, but these are sometimes missing. We fill them as follows. If a position does not have an end year and it is the *most recent* job for a given person, we assign the end year to the year the profile was last updated. If it is not the most recent job for a given person, we assign the end year to the start year of the next job minus one. Similarly, if the position does not have a start year and it is the *first* job listed for a given person, we assign the start year to the year the profile was created or the position's end year, whichever is earlier. If it is not the first job listed for a given person, we assign the start year to the previous job's end year plus one.

Finally, if there are multiple positions overlapping for the same year, we assign the individual to the one that took up the greater fraction of the year.

B.3 Matching LinkedIn and voter file data

The LinkedIn data contains name and location, but most individuals do not report their birth year, which is useful in disambiguating matches to the voter file. So, we first attempt to predict birth year for each individual who does not report it, prioritizing:

1. High school start year minus 14 (if ≥ 1950)
2. High school graduation year minus 18 (if ≥ 1952)
3. Bachelor's degree start year minus 19 (if ≥ 1954)
4. Bachelor's degree graduation year minus 23 (if ≥ 1956)
5. First job start year minus 23 (if ≥ 1960)

Then, we attempt to match the LinkedIn to the voter file data in successive rounds with different sets of identifiers, blocking on state (so we never match observations across states, although one LinkedIn row can be in a CBSA that spans multiple states). In each round, we require that the LinkedIn rows be complete (have non-missing values for all the matching variables) and unique (there is only one LinkedIn row with that combination of matching variables within the state). We require that the voter file rows are complete, but not unique, which means that one LinkedIn row can match multiple voter file rows but each voter file row can match only one LinkedIn row in each round. When we include birth year in the set of matching variables, we enforce that the difference in birth years is less than or equal to 10.[B1] We also do not allow middle initials, when non-missing, to disagree.

We use 72 matching rounds, each with a different set of identifiers. To summarize, we vary:

- *Name*: full name both including and excluding middle initial, first four letters of first and last name, first name cleaned using a list of nicknames and diminutives
- *Birth year*: reported by the individual, predicted from high school or bachelor's degree year, predicted using any method
- *Location*: Census place, nearest Census place, county, CBSA

We achieve an overall match rate of 45.7%. We then drop the 5.2% of LinkedIn rows who are matched to two or more individuals in the voter file for a unique match rate of 40.5%. Table B2 shows the match rate by state.

B.4 Matching DIME and TargetSmart data

The DIME data we use includes individuals' contributions from the 2012 through 2020 election cycles, governor, and presidential races. The data has the contributor's full name and address, which we use to merge it to the TargetSmart data.

[B1] When we match on first and last name alone, we also enforce the 10 year birth year threshold, but only if the LinkedIn birth year is sourced from high school of bachelor's degree start or end year, and we allow it to be missing.

Like above, we attempt to match the two in successive rounds with different sets of identifiers, blocking on state. In each match round, we require that rows in both DIME and TargetSmart be complete and unique. We also enforce that middle initials, when non-missing, cannot disagree.

We use 15 matching rounds, each with a different set of identifiers. To summarize, we vary:

- *Name*: full name both including and excluding middle initial, first four letters of first and last name, first name cleaned using a list of nicknames and diminutives
- *Location*: Full street address, street name and ZIP code, ZIP code, census tract

Overall, we match 62.8% of DIME rows to the voter file. Table B3 contains the match rate for each method and shows that about 80% of matches use the exact address from DIME and the voter file.

Table B1: Match Rates by Method for LinkedIn-TargetSmart Match

Merging variables	Num. of obs.	Prop.
first, last, YOB, exact place	9,628,387	0.146
first, last, exact place	8,308,492	0.126
first, last, YOB, county	6,194,872	0.094
first, last, county	5,423,143	0.082
f4, l4, address number, street name, ZIP	4,981,437	0.076
f4, l4, county	4,290,788	0.065
first, last, YOB, CBSA	2,638,841	0.040
first, last, YOB	2,536,964	0.039
first, last, YOB (strict), exact place	2,327,415	0.035
first, last	1,908,054	0.029

Notes: "f4" and "l4" refer to the first four letters of the first and last name, respectively. "Strict" birth year refers to birth year reported by the individual or sourced from high school or bachelor's degree year.

Table B2: Match Rates by State for LinkedIn-TargetSmart Match

	Num.			Share		
State	0 matches	1 match	2+ matches	0 matches	1 match	2+ matches
All	**88,197,973**	**65,854,387**	**8,383,481**	**0.543**	**0.405**	**0.052**
Alabama	612,645	809,629	112,488	0.399	0.528	0.073
Alaska	110,306	128,754	12,044	0.439	0.513	0.048
Arizona	1,267,384	1,397,291	182,777	0.445	0.491	0.064
Arkansas	374,334	488,255	59,556	0.406	0.529	0.065
California	8,337,840	8,113,917	1,213,428	0.472	0.459	0.069
Colorado	1,236,699	1,502,031	152,811	0.428	0.519	0.053
Connecticut	493,835	646,165	74,395	0.407	0.532	0.061
Delaware	295,350	318,025	61,714	0.437	0.471	0.091
District of Columbia	1,432,991	425,780	103,319	0.730	0.217	0.053
Florida	3,842,567	4,361,117	546,590	0.439	0.498	0.062
Georgia	2,057,953	2,279,693	300,502	0.444	0.492	0.065
Hawaii	223,689	196,025	20,877	0.508	0.445	0.047
Idaho	248,003	321,945	22,012	0.419	0.544	0.037
Illinois	2,415,314	2,814,149	373,329	0.431	0.502	0.067
Indiana	995,141	1,350,717	154,390	0.398	0.540	0.062
Iowa	403,165	591,789	43,023	0.388	0.570	0.041
Kansas	905,463	625,384	75,084	0.564	0.389	0.047
Kentucky	480,228	660,454	91,980	0.390	0.536	0.075
Louisiana	683,610	821,579	103,380	0.425	0.511	0.064
Maine	185,488	256,896	26,102	0.396	0.548	0.056
Maryland	984,420	1,176,234	133,815	0.429	0.513	0.058
Massachusetts	1,769,398	1,777,296	241,417	0.467	0.469	0.064
Michigan	1,450,962	2,101,538	232,314	0.383	0.555	0.061
Minnesota	1,030,031	1,336,300	108,925	0.416	0.540	0.044
Mississippi	270,124	380,082	45,202	0.388	0.547	0.065
Missouri	949,620	1,152,045	123,883	0.427	0.518	0.056
Montana	132,201	185,538	10,085	0.403	0.566	0.031
Nebraska	272,938	354,818	23,761	0.419	0.545	0.036
Nevada	530,067	565,562	63,939	0.457	0.488	0.055
New Hampshire	175,346	226,789	18,551	0.417	0.539	0.044
New Jersey	1,690,408	1,745,876	321,963	0.450	0.465	0.086
New Mexico	311,455	306,115	41,115	0.473	0.465	0.062
New York	5,221,418	3,896,801	533,226	0.541	0.404	0.055
North Carolina	1,716,228	2,221,722	295,328	0.405	0.525	0.070
North Dakota	103,782	112,118	7,726	0.464	0.501	0.035
Ohio	1,723,169	2,434,040	262,732	0.390	0.551	0.059
Oklahoma	721,004	661,320	80,802	0.493	0.452	0.055
Oregon	732,194	954,170	89,022	0.412	0.537	0.050
Pennsylvania	2,116,243	2,570,103	303,765	0.424	0.515	0.061
Rhode Island	211,701	224,637	23,968	0.460	0.488	0.052
South Carolina	695,558	932,974	97,772	0.403	0.540	0.057
South Dakota	97,681	126,235	6,101	0.425	0.549	0.027
Tennessee	1,045,805	1,239,960	169,138	0.426	0.505	0.069
Texas	5,121,510	5,476,573	855,475	0.447	0.478	0.075
Utah	586,635	668,631	60,467	0.446	0.508	0.046
Vermont	92,005	131,372	7,050	0.399	0.570	0.031
Virginia	1,260,768	1,561,587	186,509	0.419	0.519	0.062
Washington	1,499,941	1,735,517	160,161	0.442	0.511	0.047
West Virginia	145,752	200,158	22,282	0.396	0.544	0.061
Wisconsin	821,307	1,198,779	120,781	0.384	0.560	0.056
Wyoming	88,534	89,902	6,405	0.479	0.486	0.035

Table B3: Match Rates by Method for DIME-TargetSmart Merge

	Num.	Prop. of total
Total number of individuals in DIME	34,101,629	1.000
Merged using any method	21,420,056	0.628
Merged using f4, l4, address number, street name, ZIP	16,974,259	0.498
Merged using f4, l4, street name, ZIP	305,191	0.009
Merged using cleaned first name, l4, street name, ZIP	436,925	0.013
Merged using first name, middle initial, last name, tract	403,260	0.012
Merged using first name, middle initial, last name, ZIP	261,268	0.008
Merged using first name, last name, tract	1,622,393	0.048
Merged using first name, last name, ZIP	1,123,708	0.033
Merged using cleaned first name, middle initial, last name, tract	4,588	0.000
Merged using cleaned first name, middle initial, last name, ZIP	4,359	0.000
Merged using cleaned first name, last name, tract	74,908	0.002
Merged using cleaned first name, last name, ZIP	68,454	0.002
Merged using f4, middle initial, last name, tract	2,919	0.000
Merged using f4, middle initial, last name, ZIP	3,654	0.000
Merged using f4, last name, tract	60,155	0.002
Merged using f4, last name, ZIP	74,015	0.002

Notes: "f4" and "l4" refer to the first four letters of the first and last name, respectively. "Cleaned first name" uses a crosswalk of common nicknames, e.g., "John" to "Johnathan."

C Framing experiment

This appendix contains further details about the framing experiment.

C.1 Sample

We ran the survey from August 26 to 28, 2024 on Prolific, an online survey platform. We limited eligibility to those who were over 25 years old, registered to vote, English-speaking, and currently employed or looking for a job. We sampled to approximately match the LinkedIn sample on education and party affiliation. In total, we recruited 1,002 respondents. The median respondent took 15.3 minutes to complete the survey.

After dropping 5 respondents who were not interested in any of the job titles we offered, and 22 who reported their current salary as less than $5,000 or more than $300,000, we are left with 975 responses. Summary statistics for the analysis sample are in Table C1.

C.2 Survey flow

The survey proceeds as follows:

1. Ask for current or most recent job title and annual pay.[C1]
2. Based on their current occupation, ask which job titles (up to 10) the respondent is interested in. This mapping uses the O*NET list of "related occupations."
3. Present five pairs of jobs. For all choices except the placebo (choice #4):
   - First, pick a job title (from among those the respondent expressed interest in) and an industry corresponding to that job title.[C2] Draw a company size from {small, medium, large} randomly.
   - Generate one company and job description using liberal values, and one using neutral values. (See next section for details.)
   - Draw center of salary offers: $c = \text{Uniform}[0.95, 1.15]$. The salary range for *both* Job A and Job B is $[0.9cp, 1.1cp]$, rounded to nearest $1,000, where p is the respondent's current annual pay.
   - Show the respondent the two job descriptions and the salary range, and ask them to describe the main differences between the two jobs.
   - Draw four pairs of salary offers randomly. For each, draw without replacement the salary differential d from {0, 0.02, 0.04, 0.06, 0.08, 0.10 0.12}.[C3] The low salary offer is $(c - 0.5d)p$ and the high salary offer is $(c + 0.5d)p$, both rounded to the nearest $100.
   - Ask respondents to choose between "Strongly prefer Job A," "Prefer Job A," "Prefer Job B," and "Strongly prefer Job B."

[C1] If the job title corresponds to more than one occupation (SOC code), we ask the respondent to clarify. The mapping of job titles to occupations is a manually edited version of the O*NET list of "alternate titles."

[C2] Industry options are those that at least 3% of people in the occupation corresponding to the title worked in during 2022 according to the BLS Occupational Employment and Wage Statistics.

[C3] We draw $d = 0$ with half the probability as the other values.

Table C1: Summary Statistics for the Framing Experiment

	Mean	Std. dev.
Birth year	1984.3	10.4
Male	0.425	0.495
White	0.665	0.472
Democrat	0.476	0.500
Republican	0.333	0.472
Strong Democrat	0.258	0.438
Lean Democrat	0.217	0.413
Lean Republican	0.174	0.380
Strong Republican	0.159	0.366
High school	0.040	0.196
Associate's degree	0.164	0.371
Bachelor's degree	0.541	0.499
Graduate degree	0.253	0.435
Full- or part-time	0.848	0.359
Self-employed	0.067	0.250
Unemployed	0.072	0.258
Current salary	68,456	41,324
Num. of obs.	975	

Notes: The table shows summary statistics for the analysis sample.

The placebo choice differs in two ways: first, both job descriptions are generated using either neutral or liberal values (with equal probability). Second, the last pair of salary offers is always the same as the first, to check for consistency. This means that each respondent sees 19 unique sets of salary offers across 5 pairs of jobs.

Figure C1 shows how the job choices are presented.

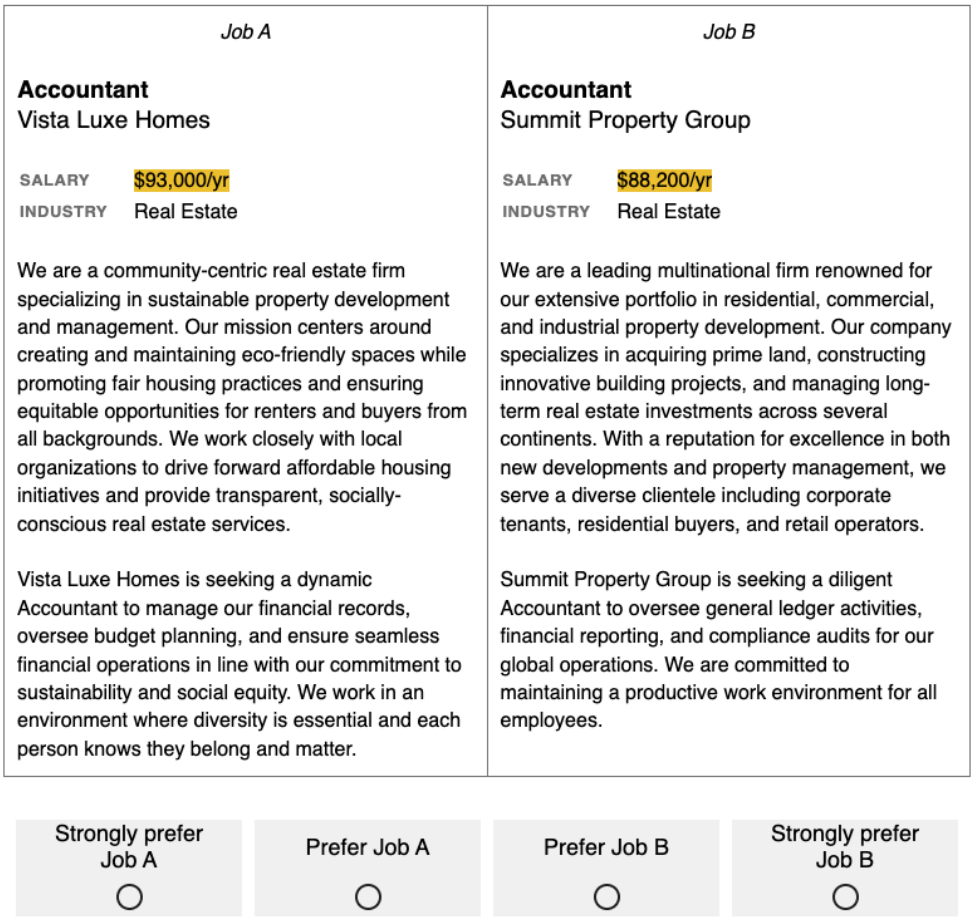
4. Ask about demographics; the respondent's current job and coworkers; party affiliation, voter registration, and presidential vote choice in 2020 and 2024.

C.3 Generating job descriptions

We use OpenAI's ChatGPT 4o model to generate a company and job description that incorporates the job title, company's industry, and size, and optionally emphasizes liberal values. The two sets of LLM instructions are:

- *Liberal values*: You are an HR specialist who writes 2-paragraph job advertisements for liberal companies. Given a job title and an industry, first write a 1-paragraph description of a company. Emphasize one liberal value like equity, inclusion, anti-racism, or sustainability. The paragraph is 3 sentences long and starts with "We are". Describe the company's main activities or products. Be specific and don't include the company's name. The second paragraph is 1 sentence long and describes the job's tasks. Begin the sentence with 'We are seeking a [title] to'
- *Neutral values*: You are an HR specialist who writes 2-paragraph job advertisements. Given a job title and an industry, first write a 1-paragraph description of a company. The paragraph

Figure C1: Framing Experiment Layout



Notes: The figure shows an example of a job choice presented in the framing experiment.

is 3 sentences long and starts with "We are". Describe the company's main activities or products. Be specific and don't include the company's name. Be neutral and don't mention anything political. Don't mention diversity, sustainability, or the environment. The second paragraph is 1 sentence long and describes the job's tasks. Begin the sentence with 'We are seeking a [title] to'

Then, we replace "We" in the second paragraph with a company name. We randomly draw these company names from a list specific to the industry, again generated using ChatGPT 4o. We also randomly assign an adjective to appear after the word "seeking" from one of the following: *dedicated, passionate, diligent, dynamic, committed, driven, motivated, ambitious, enthusiastic*, and *exceptional*.

Finally, we draw one of the following sentences at random (without replacement) and append it to the last paragraph. For liberal ads:

- We are committed to creating a diverse and inclusive work environment.
- Through our talent, we believe that diversity make us a more vibrant place to work.

- We work in an environment where diversity is essential and each person knows they belong and matter.
- Self-care is encouraged here, and we are committed to providing a positive work culture.
- We are on an anti-racist journey, committed to learning, living, and breathing inclusion for all.
- We strive to recruit a workforce that represents our diverse and culturally rich community.

For neutral ads:

- We prioritize creating an environment that encourages excellence.
- We are dedicated to building a culture of innovation and initiative.
- Our culture is built on steadfast commitment to responsibility and excellence.
- Our energetic environment creates a unique opportunity for professional growth.
- We offer a unique opportunity to contribute to the success of a dynamic organization.
- Our culture empowers our employees to take initiative, be proactive, and contribute to our collective success.

D Conjoint experiment

This appendix contains further details about the conjoint experiment.

D.1 Sample

We collected responses from August 22 to September 1, 2024 through the Harvard Digital Lab for the Social Sciences, a volunteer panel operated by Harvard University. In total, we recruited 539 respondents. The median respondent took 15.4 minutes to complete the survey and spent 36.8 seconds on each of the 12 pairs of jobs.

After dropping 1 respondent who was not interested in any of the job titles we offered, 12 who reported their current salary as more than $300,000, and 109 who responded inconsistently when presented with the same choice twice, we are left with 417 responses. Summary statistics for the analysis sample are in Appendix Table D1.

D.2 Survey flow

The survey flow is similar to the survey described in Appendix C, and proceeds as follows:

1. Ask for current or most recent job title or occupation and annual pay.
2. Based on their current occupation, ask which job titles (up to 10) the respondent is interested in.
3. Present 12 pairs of jobs, varying the attributes of Job A and Job B. (See next section for details.) Ask respondents to choose between "Strongly prefer Job A," "Prefer Job A," "Prefer Job B," and "Strongly prefer Job B."

   The tenth choice is a repeat of the first one (with the labels Job A and Job B flipped). This means each respondent sees 11 unique pairs of alternatives.

   Figure D1 shows how the job choices are presented. We highlight the attributes that vary between Job A and Job B to focus the respondent's attention.
4. Ask about demographics; the respondent's current job and coworkers; party affiliation, voter registration, and presidential vote choice in 2020 and 2024.

D.3 Hypothetical job profiles

To create each pair of job alternatives A and B, we first randomly choose one job title from among the titles that the respondent expresses interest in. The job title never varies between Job A and Job B.

To reduce the number of comparisons the respondent must make, for each choice pair, we pick **exactly four** of the following attributes to vary between Job A and Job B. However, the probability that each attribute varies is not identical. Salary varies with 10 times the probability as the rest of the attributes, and the share of partisan coworkers varies with five times the probability.

1. *Share Democrats among coworkers*: Randomly sample from Uniform[10, 90], rounded to the nearest 10.
2. *Salary*:

Table D1: Summary Statistics for the Conjoint Experiment

	Mean	Std. dev.
Birth year	1964.0	14.9
Male	0.671	0.470
White	0.901	0.299
Democrat	0.352	0.478
Republican	0.298	0.458
Strong Democrat	0.265	0.442
Lean Democrat	0.087	0.282
Lean Republican	0.110	0.313
Strong Republican	0.188	0.391
High school	0.160	0.367
Associate's degree	0.091	0.287
Bachelor's degree	0.319	0.467
Graduate degree	0.389	0.488
Full- or part-time	0.531	0.499
Self-employed	0.085	0.280
Unemployed	0.383	0.487
Current salary	90,725	55,770
Num. of obs.	574	

Notes: The table shows summary statistics for the analysis sample.

- *When chosen to vary*: Draw two salaries from Uniform[0.9p, 1.3p] (where p is the respondent's current annual pay), rounded to the nearest $1,000. Redraw until the salaries differ by more than 2% and less than 15%. Assign the lower of the two salaries to the job with more Democratic coworkers with probability $\frac{3}{4}$, and to the job with fewer Democratic workers with probability $\frac{1}{4}$. If the share of Democrats does not vary across jobs, then randomly choose which job has the lower salary.
- *When not chosen to vary*: Randomly sample from Uniform[0.9p, 1.3p] (where p is the respondent's current annual pay), rounded to the nearest $1,000.

3. *Industry*: Randomly sample from possible industries.[D1]
4. *Number of employees*: Randomly sample from {10–50, 50–200, 200+}.
5. *Share men among coworkers*: Randomly sample from Uniform[20, 80], rounded to the nearest 10.
6. *Share white among coworkers*: Randomly sample from Uniform[30, 90], rounded to the nearest 10.
7. *Average age of coworkers*: Randomly sample from Uniform[30, 50], rounded to the nearest 5.
8. *Corporate political contributions*: Choose among {"No political contributions", "Contributes to Democrats", "Contributes to Republicans"} with the following probabilities:
   - *When chosen to vary:*

[D1] Industry options are those that at least 3% of people in the occupation (6-digit SOC code) corresponding to the title have worked in, according to BLS data.

Figure D1: Conjoint Experiment Layout

	Job A	Job B
Job title	Economist	Economist
Industry	Business Consulting	Insurance
Job description	We are seeking an Economist to conduct thorough economic research, analyze data trends, and develop actionable business strategies that drive client success.	We are seeking an Economist to conduct market research, analyze economic data, and provide forecasts that will inform our pricing strategies and risk assessments.
Size	100-500 employees	100-500 employees
Corporate political contributions	No political contributions	No political contributions
Salary	$148,000/year	$130,000/year
Benefits	- Cellphone service reimbursement - Life insurance	- Cellphone service reimbursement - Life insurance

Coworker information

	Job A	Job B
Gender	70% men 30% women	70% men 30% women
Political party	50% Democrats 50% Republicans	70% Democrats 30% Republicans
Average age	35 years old	30 years old
Race	50% white 50% nonwhite	50% white 50% nonwhite

Notes: The figure shows an example of a job choice presented in the conjoint experiment.

  - For the company with more Democrats: $\{\frac{1}{4}, \frac{3}{4}, 0\}$
  - For the company with fewer Democrats: $\{\frac{1}{4}, 0, \frac{3}{4}\}$
  - If the companies have the same share of Democrats: $\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\}$
  - Redraw until the values for Job A and Job B are different
- *When not chosen to vary:*
  - "No political contributions" with probability $\frac{1}{2}$
  - "Contributes to Democrats" with probability $\frac{1}{4}$
  - "Contributes to Republicans" with probability $\frac{1}{4}$

9. *Benefits*: Choose two of the following (without replacement)[D2]
   - Medical, vision, and dental insurance
   - Mental and emotional health benefits
   - Life insurance

[D2] When benefits are randomly selected to vary, only one of the two benefits is different across Job A and Job B.

- 401(k) company match
- Tuition reimbursement

Finally, the one-sentence job description is a deterministic function of the job title, industry, and number of employees, generated by an LLM. The purpose of this sentence is to ensure that the respondent has a concrete sense of the job's general roles and responsibilities. It does not factor into our estimation procedure.

E Model

This appendix sets up a model of two groups of workers with heterogeneous preferences for a workplace amenity and shows how it leads to segregation. The setup directly follows the Rosen (1986) model of equalizing differences, and our exposition follows Williams (2015).

E.1 Setup

We consider a static, partial equilibrium, perfect information model comprising a unit measure of workers **I** and a unit measure of firms **F**. Workers choose a firm at which they will be employed while firms choose whether to be "liberal" or "neutral." We assume that each firm employs a single worker, and study segregation in this framework by measuring segregation across *sectors* rather than segregation across *firms*.

Workers have heterogeneous preferences for two types of jobs: liberal ($D = 0$) and neutral ($D = 1$). The wages in each job are w_0 and w_1. Workers are equally productive and differ only in their preference for D.

Labor supply: Workers have quasilinear preferences over consumption C and whether the job is liberal D: $U_i(C, D) = C - Z_i D$. Thus, Z_i is the compensating differential for $D = 1$ compared to $D = 0$ – that is, the consumption necessary to make the worker indifferent between the two types of jobs. If $Z_i > 0$, then the worker prefers the liberal job.

Neutral jobs offer $(w_1, 0)$, while liberal jobs offer $(w_0, 1)$. The market equalizing wage differential is $\Delta w = w_1 - w_0$. Workers choose $D = 1$ if $\Delta w > Z_i$ and choose $D = 0$ if $\Delta w < Z_i$. Note that this model does not specify an equilibrium *level* of wages, only the differential.

Let n_d be the fraction of Democrats in the population, and let $Z_i \sim F_d(Z)$ for Democrats and $Z_i \sim F_r(Z)$ for Republicans, so that $f(Z) = n_d \times f_d(Z) + (1 - n_d) \times f_r(Z)$ is the density of Z_i in the population.

Let L_1^s be the fraction of workers applying to $D = 1$ jobs (that is, the fraction of workers with $Z_i < \Delta w$), and L_0^s be the fraction of workers applying to $D = 0$ jobs (that is, the fraction of workers with $Z_i > \Delta w$). Then we have

$$L_1^s = \int_0^{\Delta w} f(Z)dZ = n_d F_d(\Delta w) + (1 - n_d)F_r(\Delta w)$$
$$L_0^s = \int_{\Delta w}^{\infty} f(Z)dZ = n_d(1 - F_d(\Delta w)) + (1 - n_d)(1 - F_r(\Delta w))$$

Labor demand: Firms hire one worker and sell a good y to the market; we normalize the price of y to 1. Firms choose whether to provide the $D = 1$ or $D = 0$ job and have heterogeneous costs B_f of doing so. The production technology is:

$$y = a_1 \text{ if } D = 1$$
$$y = a_0 \text{ if } D = 0$$

Define $B_f = a_1 - a_0$, which is the cost of providing the liberal job in terms of foregone output. If $B_f > 0$, then it is costly to provide the liberal job. Let $B_f \sim G(B)$, so that $g(B)$ is the density of B_f in the population of firms.

If the firm provides the liberal job ($D = 0$), it can hire a worker at a wage Δw lower. So, firms choose $D = 1$ if $B_f > \Delta w$ and choose $D = 0$ if $B_f < \Delta w$.

Let L_1^d be the fraction of firms demanding workers in jobs with $D = 1$ (that is, the fraction of firms with $B_f > \Delta w$), and L_0^d be the fraction of firms demanding workers in jobs with $D = 0$ (that is, the fraction of firms with $B_f < \Delta w$). Then we have:

$$L_1^d = \int_{\Delta w}^{\infty} g(B)dB = 1 - G(\Delta w)$$
$$L_0^d = \int_{0}^{\Delta w} g(B)dB = G(\Delta w)$$

Equilibrium: Δw adjusts so that $L_1^s = L_1^d$ and $L_0^s = L_0^d$. The equilibrium wage differential Δw^* thus solves

$$n_d F_d(\Delta w^*) + (1 - n_d)F_r(\Delta w^*) = 1 - G(\Delta w^*)$$

Segregation: In equilibrium, workers and firms are systematically matched: workers in $D = 1$ jobs have the lowest WTP for the liberal job, and firms offering $D = 1$ jobs have the highest costs of providing it. This produces negative assortative matching – workers with larger Z_i values (Democrats) are systematically found in firms with smaller B_f values (e.g., mental healthcare), and workers with small Z_i values (Republicans) are systematically found in firms with larger B_f values (e.g., oil and gas). This generates segregation between the two groups, which we can measure with the overexposure ratio. We define S_r as the fraction of co-partisans in the same sector (liberal or neutral) as a Republican worker, normalized by the fraction of Republicans in the population; S_d is defined similarly, and the weighted average of the two is the overall overexposure ratio S:

$$S = \frac{n_{0r}^2}{n_r n_0} + \frac{n_{1r}^2}{n_r n_1} + \frac{n_{0d}^2}{n_d n_0} + \frac{n_{1d}^2}{n_d n_1}$$

where n_d is the number of Democrats, n_r is the number of Republicans, and we normalize $n_d + n_r = 1$. The number of Democrats in $D = 0$ jobs is n_{0d}, and in $D = 1$ jobs is n_{1d}, which means $n_{0d} + n_{1d} = n_d$. Definitions are symmetric for Republicans.

Parameterization: Assume worker preferences and firm costs are normally distributed: $F_d(\Delta w) \sim N(\mu_d, \sigma^2)$, $F_r(\Delta w) \sim N(\mu_r, \sigma^2)$, and $G \sim N(\mu_f, \sigma_f^2)$. Also, assume Democrats value the liberal job more than Republicans on average and that the average firm cost of providing the liberal job falls in the middle of these two: $\mu_d > \mu_f > \mu_r$.

E.2 Comparative statics

Proposition: As the standard deviation of worker preferences increases, the overexposure ratio decreases, or $\frac{\partial S}{\partial \sigma} < 0$. We prove this in five steps.

Step 1: We show that $\frac{\partial S}{\partial \sigma} \propto \frac{\partial S_r}{\partial \sigma}$ by rewriting the overexposure ratio as

$$
\begin{aligned}
S &= n_r S_r + \frac{(n_0 - n_{0r})^2}{(1-n_r)n_0} + \frac{(n_1 - n_{1r})^2}{(1-n_r)n_1} \\
&= \frac{n_r}{1-n_r} S_r + \frac{1-2n_r}{1-n_r}.
\end{aligned}
$$

Step 2: The share of workers in $D = 0$ jobs who are Republican is smaller than the share of workers in $D = 1$ jobs who are Republican, or $\frac{n_{0r}}{n_0} < \frac{n_{1r}}{n_1}$, which follows from our assumptions on F_r and F_d:

$$
\begin{aligned}
F_r(\Delta w) &> F_d(\Delta w) \\
F_r(\Delta w) &> n_r F_r(\Delta w) + n_d F_d(\Delta w) \\
\frac{n_{1r}}{n_r} &> n_1
\end{aligned}
$$

and likewise $\frac{n_{0r}}{n_r} < n_0$, from which the result follows.

Step 3: The equilibrium wage differential (Δw^*) lies between the mean WTP of Democrats and Republicans, or $\mu_r < \Delta w^* < \mu_d$. The wage differential is defined implicitly by

$$
n_d F_d(\Delta w^*) + n_r F_r(\Delta w^*) = 1 - G(\Delta w^*).
$$

Thus

$$
\begin{aligned}
F_r(\Delta w^*) &> 1 - G(\Delta w^*) \\
\Phi\left(\frac{\Delta w^* - \mu_r}{\sigma}\right) &> 1 - \Phi\left(\frac{\Delta w^* - \mu_f}{\sigma_f}\right).
\end{aligned}
$$

Suppose $\Delta w^* < \mu_r$. Then $\Phi\left(\frac{\Delta w^* - \mu_f}{\sigma_f}\right) < \Phi\left(\frac{\Delta w^* - \mu_r}{\sigma}\right) < 0.5$. But then we have $x > 1 - y$ for $y < x < 0.5$, a contradiction. Similarly, we have

$$
\begin{aligned}
F_d(\Delta w^*) &< 1 - G(\Delta w^*) \\
\Phi\left(\frac{\Delta w^* - \mu_d}{\sigma}\right) &< 1 - \Phi\left(\frac{\Delta w^* - \mu_f}{\sigma_f}\right).
\end{aligned}
$$

Suppose $\Delta w^* > \mu_d$. Then $\Phi\left(\frac{\Delta w^* - \mu_f}{\sigma_f}\right) > \Phi\left(\frac{\Delta w^* - \mu_d}{\sigma}\right) > 0.5$. But then we have $x < 1 - y$ for $y > x > 0.5$, a contradiction.

Step 4: As the standard deviation of worker preferences increases, the number of Democrats in $D = 0$ jobs decreases, while the number of Republicans in $D = 0$ jobs increases, or $\frac{\partial n_{0d}}{\partial \sigma} < 0 < \frac{\partial n_{0r}}{\partial \sigma}$. This follows from the previous step:

$$n_{0d} = n_d(1 - F_d(\Delta w)) = n_d - n_d\Phi\left(\frac{\Delta w - \mu_d}{\sigma}\right)$$
$$\frac{\partial n_{0d}}{\partial \sigma} = -\frac{n_d}{\sigma}\phi\left(\frac{\Delta w - \mu_d}{\sigma}\right)\left(\frac{\partial \Delta w}{\partial \sigma} - \frac{\Delta w - \mu_d}{\sigma}\right)$$

which is negative if and only if

$$\frac{\partial \Delta w}{\partial \sigma} > \frac{\Delta w - \mu_d}{\sigma}.$$

Now, the equilibrium wage is defined by

$$n_d F_d(\Delta w) + n_r F_r(\Delta w) = 1 - G(\Delta w)$$
$$n_d\Phi\left(\frac{\Delta w - \mu_d}{\sigma}\right) + (1 - n_d)\Phi\left(\frac{\Delta w - \mu_r}{\sigma}\right) = 1 - \Phi\left(\frac{\Delta w - \mu_f}{\sigma_f}\right)$$

taking derivatives,

$$\frac{n_d}{\sigma}\phi\left(\frac{\Delta w - \mu_d}{\sigma}\right)\left(\frac{\partial \Delta w}{\partial \sigma} - \frac{\Delta w - \mu_d}{\sigma}\right) + \frac{n_r}{\sigma}\phi\left(\frac{\Delta w - \mu_r}{\sigma}\right)\left(\frac{\partial \Delta w}{\partial \sigma} - \frac{\Delta w - \mu_r}{\sigma}\right)$$
$$= -\frac{1}{\sigma_f}\phi\left(\frac{\Delta w - \mu_f}{\sigma_f}\right)\frac{\partial \Delta w}{\partial \sigma}.$$

Let $q_d = \frac{n_d}{\sigma}\phi\left(\frac{\Delta w - \mu_d}{\sigma}\right)$, $q_r = \frac{n_r}{\sigma}\phi\left(\frac{\Delta w - \mu_r}{\sigma}\right)$, and $q_f = \frac{1}{\sigma_f}\phi\left(\frac{\Delta w - \mu_f}{\sigma_f}\right)$. Then

$$\frac{\partial \Delta w}{\partial \sigma} = \frac{q_d\frac{\Delta w - \mu_d}{\sigma} + q_r\frac{\Delta w - \mu_r}{\sigma}}{q_d + q_r + q_f} > \frac{\Delta w - \mu_d}{\sigma}$$

since $\mu_r < \Delta w < \mu_d$ from the previous step, and the proof for $\frac{\partial n_{0r}}{\partial \sigma} > 0$ is symmetric.

Step 5: Taking the derivative of the overexposure ratio for Republicans with respect to σ:

$$S_r = \frac{1}{n_r^2}\left(\frac{n_{0r}^2}{n_0} + \frac{n_{1r}^2}{n_1}\right)$$
$$\frac{\partial S_r}{\partial \sigma} = \frac{1}{n_r^2}\left(\frac{2n_{0r}\frac{\partial n_{0r}}{\partial \sigma}n_0 - n_{0r}^2\frac{\partial n_0}{\partial \sigma}}{n_0^2} + \frac{2n_{1r}\frac{\partial n_{1r}}{\partial \sigma}n_1 - n_{1r}^2\frac{\partial n_1}{\partial \sigma}}{n_1^2}\right).$$

Since $\frac{\partial n_{0r}}{\partial \sigma} = -\frac{\partial n_{1r}}{\partial \sigma}$ and $\frac{\partial n_0}{\partial \sigma} = -\frac{\partial n_1}{\partial \sigma}$, this equals

$$\frac{\partial S_r}{\partial \sigma} = \frac{1}{n_r^2}\left(\frac{n_{0r}}{n_0} - \frac{n_{1r}}{n_1}\right)\left[\frac{\partial n_{0r}}{\partial \sigma}\left(2 - \frac{n_{0r}}{n_0} - \frac{n_{1r}}{n_1}\right) - \frac{\partial n_{0d}}{\partial \sigma}\left(\frac{n_{0r}}{n_0} + \frac{n_{1r}}{n_1}\right)\right],$$

which is negative by the previous steps, thus the claim follows from the first step.