Walking the Tightrope: Understanding Stakeholder Political Preferences and Their Influence on Corporate Sociopolitical Activism

by

Max I. Kagan

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Political Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor David Broockman, Chair
Professor Paul Pierson
Professor David Vogel

Summer 2024

Walking the Tightrope: Understanding Stakeholder Political Preferences and Their Influence on Corporate Sociopolitical Activism

Abstract

Walking the Tightrope: Understanding Stakeholder Political Preferences and Their Influence on Corporate Sociopolitical Activism

by

Max I. Kagan

Doctor of Philosophy in Political Science

University of California, Berkeley

Professor David Broockman, Chair

This dissertation focuses on three questions within non-market strategy—that is to say, at the intersection of business and politics.

Recently, companies have increasingly taken public stances on controversial political and social issues. While business has long sought to influence political policies in its own self-interest, this typically focused on economic issues with a direct connection to the financial bottom line—e.g., taxes, trade, and regulation—rather than contentious, morally-infused issues such as abortion, gun violence, LGBTQ rights, and racial (in)equality. In thinking about when and why companies may engage in this activism, I consider how stakeholders' influence may influence and constrain corporate decisions to engage on these issues. In this dissertation, the first two chapters focus consider employees, while the third focuses on customers.

The first chapter of this dissertation, **"Office Parties: Partisan Sorting in the United States Labor Market"** (co-authored with Justin Frake and Reuben Hurst)[1] explores the extent of partisan sorting among employees in the United States. We create an original longitudinal dataset by merging voter registration data with with 17 million online employee profiles covering 14.5 million unique workers from 2012–2022. This represents the largest-ever dataset of this kind, with significantly larger (and less socioeconomically biased) coverage than other measures of employees' partisanship. Using these original data, we present four main findings. First, we find significant evidence that individuals are politically sorted across workplaces. Whereas prior work in political science has held up the workplace as an ideal locus of cross-partisan context, we find that the average Democratic worker's co-workers are about 15 percentage points more Democratic than the average Republicans co-workers, and vice versa. The extent of observed sorting is largely—but not entirely—associated with partisan sorting among metropolitan areas, industries,

[1] As of the filing of this dissertation, "Office Parties" is a public working paper and is not yet published. I thank both Justin Frake and Reuben Hurst for their permission in allowing me to adapt this co-authored work for the purposes of this dissertation.

and occupations. Accounting for factors correlated with partisanship—geography, occupation, and industry—reduces this estimate to about 2 percentage points, which is similar in magnitude to our estimates of workplace sorting by gender and race. Second, we find that sorting appears more pronounced among workers who may have more labor market power (e.g., white workers, more senior workers, and workers in jobs which require more training) as well as workers who appear more politically engaged. Third, political sorting has increased among new joiners to firms since 2017. Fourth, because Democrats comprise a greater share of the workforce, sorting means that Republicans experience a significantly larger share of out-partisan co-workers.

The second chapter of this dissertation, "**Do Conservative Social Policies Harm Access to Employee Talent? Evidence from the North Carolina Bathroom Bill**," takes up a common rhetorical claim companies make to explain their decision to speak out on contentious socially issue. One common explanation for "corporate sociopolitical activism" is that government policies can make it harder to attract talented employees, and so companies expose these policies which have the potential to harm their businesses. While these employee-centric justifications for CSA are widespread, there is limited evidence measuring whether these types of policies actually do have a negative impact on firms' ability to attract employees. This paper tests whether government policies do harm businesses by reducing their access to employee talent. I look at employee responses to North Carolina's 2016 HB2 law (the "bathroom bill"), which prompted national attention and over 100 CEOs to argue that it would harm North Carolina's ability to attract workers. Using administrative data which cover 95% of U.S. employees, and the recently-developed Generalized Synthetic Control (GSC) technique (Xu 2017), I find no evidence that HB2 resulted in a decline in workers' in-migration to North Carolina or an increase in workers' out-migration from the state, meaning that the bill did not appear to reduce North Carolina employers' access to talent. This null finding is precisely estimated and holds both for overall in-migration and out-migration and there is no evidence of heterogeneity among young workers, college-educated workers, or workers in specific industries. I also rule out that this null result was driven by a compensating differential in wages for workers relocating to or from North Carolina or that it was the result of offsetting partisan shifts in migration patterns.

Finally, the third chapter, **"This Bud's For You? The Effect of Partisan Political Cues on Politically Polarized Brands,"** investigates whether partisan cues can potentially unwind or attenuate already-polarized customer preferences. Real-world events and political cues can induce political polarization in consumption patterns. When politically-motivated customers find out that seemingly neutral or apolitical brands support a political party or take a political stance, they may respond by boycotting or "buycotting" the brand in question. To date, no research has explored what happens when customers receive a political cue about a brand that is not viewed as neutral or apolitical, but one which already has a polarized reputation. This survey experiment takes advantage of the recent conservative boycott of Bud Light following the brand's partnership with a transgender social media influencer. Using the context of this (mostly) Republican boycott of Bud Light after the brand partnered with a transgender influencer, I find that partisan political cues likely constitute a one-way ratchet on polarized brand reputations. Republicans remain negatively disposed towards the brand even when told that Republican leaders (including President Trump) support the brand. I do that partisan cues can erode support for brands: When Democrats were told that fellow Democrats were boycotting Bud Light, they become less likely to select the brand in a choice-based conjoint task, and far more likely to report boycotting the brand. I also find that this polarization in expressed attitudes may not always be backed up with actual changes in economic behavior, as a significant portion of respondents who claim to be boycotting the brand give inconsistent survey responses which suggest they may not actually be boycotting.

Contents

List of Figures

List of Tables

Acknowledgments

Even without a global pandemic, completing a PhD is always a long journey. I would not have reached this destination without the help and support of many people.

My research interests have taken many twists and turns over the past five years. During my initial years at Berkeley, Ryan Brutger, Marika Landau-Wells, and Steve Vogel provided sage advice on success in academia. As my research interests became clearer, Zhao Li generously provided invaluable career guidance that continues to shape my research trajectory to this day. In addition to being excellent co-authors, Reuben Hurst and Justin Frake have helped me to better understand research in management and strategy. I have also enjoyed and been inspired by conversations with a number of researchers working at the intersection of business and politics, including Ben Egerod, Dan Nielsen, Neil Malhotra, and Jan Stuckatz.

Paul Pierson and David Vogel served as committee members and enthusiastic supporters on this dissertation project. David Broockman went above and beyond his formal role as chair of this dissertation. His enthusiasm for the project and willingness to help with a seemingly endless series of logistical tasks made this research possible over the last few years. For thoughtful suggestions on parts of this dissertation, I thank Adam Bonica, Jacob Brown, Eitan Hersh, David Primo, Stephanie Zonszein, and the members of the Political Behavior Lab at UC Berkeley.

I also gratefully acknowledge the support of a variety of organizations which provided financial support for my research, including the Berkeley Center for the Study of American Democracy (CSAD), the Berkeley Economics and Society Initiative (BESI), the Berkeley Graduate Division, UC Berkeley's Charles and Louise Travers Department of Political Science, the Institute for Humane Studies (IHS), the Law, Economics, and Policy (LEAP) Center at Berkeley Law, and the National Science Foundation's Graduate Research Fellowship Program. The research in Chapter 1 benefited from financial and in-kind support from the University of Maryland Smith School of Business and Advanced Research Computing at the University of Michigan. While this research would not have been possible without the generous support of these organizations, the findings and conclusions of this research are entirely my own and do not necessarily represent their views.

On a more personal note, I would like to thank Christelle Richard, Eric Chen, and James Dean for welcoming me into their lives in Berkeley.

Finally, my utmost thanks and love go to Emily Chen.

Chapter 1

Office Parties: Partisan Sorting in the United States Labor Market

1.1 Introduction

Political partisans in the United States are increasingly sorted along geographic and social dimensions in ways that limit cross-partisan contact. Democrats and Republicans tend to reside in predominately "blue" or "red" regions, states, cities, and neighborhoods (Bishop 2009; Brown and Enos 2021; Brown et al. 2023; Mummolo and Nall 2017). Democrats are more likely to attend college (Downey and Liu 2023; Firoozi 2023; Zingher 2022) while Republicans are more likely to attend religious services (Mason 2015). Worshipers from both parties tend to share the pews with co-partisans (Malina and Hersh 2021; Margolis 2018). With regard to romance and family, people are more likely to date and marry those who share their partisanship, and have also become increasingly likely to pass their partisanship on to their children (Huber and Malhotra 2017; Klofstad, McDermott and Hatemi 2013; Iyengar, Konitzer and Tedin 2018). This increased sorting has coincided with rising measures of "affective polarization"—i.e., the tendency of political partisans to dislike and distrust members of the opposing party (Iyengar and Westwood 2015; Iyengar et al. 2019)—raising concerns that declining inter-group contact among rival political partisans may be reinforcing this rise in partisan animosity (Allport 1954; Lipset 1981; Ahler and Sood 2018).

Despite progress in measuring various dimensions of partisan sorting, partisan sorting at work remains a critical, but understudied possibility. Americans spend far more of their waking hours at work than in any other activity (Bureau of Labor Statistics 2023).[1] Crucially, this time at work may constitute a rare opportunity for the sort of cross-partisan intergroup contact that can reduce prejudice (Allport 1979; Mutz and Mondak 2006; Corno, La Ferrara and Burns 2022). Employment also frequently serves as an important source of social identity (Akerlof and Kranton 2000), meaning that workplaces with partisan diversity may also constitute a source

[1]In the 2023 American Time Use Survey from the BLS, the average American reporting spending 3.5 hours per day on "working and work-related activities". This figure includes a large portion of Americans outside the labor force; among the subset who do work, the figure is 8.01 hours.

of "cross-cutting cleavages" which strengthen democratic societal norms and promote tolerance (Wojcieszak and Warner 2020). From this perspective, partisan sorting across workplaces likely reduces a critical opportunity for cross-partisan contact, potentially exacerbating affective polarization and a decline in societal norms of political deliberation. Notwithstanding these potentially significant implications, however, research on this topic in the U.S. context has been limited to experimental work on hiring (Iyengar and Westwood 2015; Gift and Gift 2015; McConnell et al. 2018), or focuses exclusively on firms' top management (Hoang, Ngo and Zhang 2022; Fos, Kempf and Tsoutsoura 2022). There lacks, to our knowledge, any large scale estimate of the extent of or trends in partisan sorting among rank-and-file workers in the U.S.

We address this gap by merging administrative data from a national voter file with a separate dataset of online employee profiles. Our main dataset, a snapshot of the US workforce in 2022, contains information on the work history and partisanship for 14.5 million unique workers—approximately 9% of the labor force. While we cannot achieve fully representative or comprehensive coverage of the U.S. labor force—most notably, our use of online employee profiles results in an over-representation of white-collar workers—we believe our dataset nevertheless represent the most complete and least biased dataset currently available.

We document four descriptive results. Our first and main result is that partisans are sorted by workplace. Members of a given party tend to work in workplaces with a greater share of co-partisans. Specifically, the average Democrat's workplace has about 15 percentage points more Democrats' than the average Republican and the average Republican's workplace has about 15 percentage points more Republicans than the average Democrat's. Because partisanship is correlated with a range of other factors, including geography, industry and occupation, we conduct a series of fixed effect regressions accounting for these factors. Inclusion of geographic (MSA), industry, and occupational fixed effects reduces this estimate to two percentage points, suggesting that partisan sorting is not entirely a function of these observable correlates and may be due in part to demand- or supply-side preferences of employees and/or employers. As a benchmarking exercise, we use the same methodology to estimate the degree of sorting by gender and race and find that the estimated magnitude of political sorting in the workforce is roughly similar to sorting by gender and race. Our main estimate of sorting is robust to a variety of alternative specifications and sample exclusion criteria.

Our second result is that this sorting varies significantly by worker characteristics. Most notably, we find that partisan sorting is substantially greater among the most politically active workers in our sample. Workers who voted in the 2020 presidential primary and/or who have made political donations have a much higher portion of co-workers who share their partisanship than those who do not. We also observe higher sorting among white workers, younger workers, more senior employees, and in occupations that require greater education, experience, and training. Although by no means conclusive, these patterns are consistent with the idea that workers may value working alongside politically like-minded co-workers, and that workers with greater market power are better able to sort themselves into these types of workplaces.

Our third result is that sorting by workplace is on the rise. Specifically, new joiners to firms—including those that are newly entering the labor market or switching firms—increasingly sort into employers in which a higher share of coworkers share their partisanship. Specifically, our three-way fixed effects estimates of the degree of sorting increased 100% (from 2 percentage point to 4 percentage points) between 2012 and 2022. This increase in sorting among new joiners began to increase c. 2017 and has been accelerating more rapidly in recent years.

Our fourth and final result relates to asymmetry across parties in the degree of out-partisan exposure that occurs at work. The mean Democrat labors in a workplace where 51% of his coworkers are also Democrats, and only 23% are Republicans (a ratio of 2.2). By contrast, the mean Republican works at a workplace where 36% of her coworkers are Republican and 36% are Democrats (a ratio of 1.1). This considerable gap is driven in part by the fact that Democrats comprise a greater share of workers, a characteristics of our data and nationally representative samples, which is in turn magnified by the sorting pattern we document in our first result. This provides an important caveat to the conclusion that the workplace is an important locus for cross-partisan contact (Mutz and Mondak 2006), suggesting this is much more the case for Republican versus Democrat workers. This asymmetry may hold important political consequences and could help explain perceptions that corporate America has been tacking to the left on cultural and political issues in recent years (Hersh and Shah 2023).

We close by outlining avenues for future research that can further elucidate the origins and consequences of labor market sorting by partisanship.

1.2 Theoretical Motivation

Partisan Sorting in the United States

Political polarization has risen sharply in the United States. While scholars debate the extent to which Americans' policy preferences have polarized ideologically (Mason 2015, 2018) as well as the extent to which rising polarization actually affects political outcomes (Broockman, Kalla and Westwood 2020), there is growing evidence that Democrats and Republicans increasingly dislike and distrust one another (Iyengar et al. 2019; Iyengar and Westwood 2015). This rise in "affective polarization" coincides with increasing partisan sorting across a range of non-political domains. Not only are Americans increasingly sorted into "red" and "blue" states, but they increasingly reside in politically homogeneous regions, cities, and neighborhoods (Bishop 2009; Brown et al. 2023; Brown and Enos 2021; Mummolo and Nall 2017).

In addition to geographic segregation, partisans make other life choices that further reduce cross-partisan exposure. Rising polarization by education means Americans with a college education are increasingly Democratic (Downey and Liu 2023; Firoozi 2023; Zingher 2022), although the direction of causality remains unclear (see Apfeld et al. 2023). Members of both parties prefer to live with roommates who share their political views (Shafranek 2021). Democrats are less likely to take part in organized religion (Mason 2015), and members of both parties who attend religious services are likely to belong to politically homogeneous congregations (Malina and Hersh 2021; Margolis 2018). Sorting extends to romance and family formation: partisans are more likely date and marry fellow partisans and to pass on their partisan identity to their children (Huber and Malhotra 2017; Iyengar, Konitzer and Tedin 2018; Klofstad, McDermott and Hatemi 2013). Recent evidence suggests that partisan political sorting also extends to the economic realm (Iyengar et al. 2019; McConnell et al. 2018). Partisans' consumption habits, for example, may be bifurcated by deliberate boycotts or "buycotts" which target specific products based upon these companies' political stances (Liaukonytė, Tuchman and Zhu 2022; Hou and Poliquin 2023). More generally, consumers (McConnell et al. 2018; Panagopoulos et al. 2020) and investors (Bolton et al. 2020) appear to favor firms aligned with their political identities.

Increased partisan sorting may have negative implications for democracy. Political theorists have long argued that cross-cutting interactions between members of society who disagree are vital for democracy (Lipset 1981; Nordlinger 1972). Given the prominence of partisanship in shaping political behaviors, these interactions are especially important with regard to partisanship. This is in part because intergroup contact prevents partisans from developing negative, biased perceptions of outpartisans that underpin affective polarization. Affective polarization is likely self-reinforcing inasmuch as it causes partisans to isolate one from another, which, in turn, drives further polarization. It is therefore critical to identify contexts for polarization-decreasing, cross-partisan contact, especially in a context like the U.S where these interactions seem increasingly rare.

Partisan Sorting by Workplace

The workplace constitutes a promising, but largely understudied, context for cross-partisan contact. First, the workplace may be a rare arena where Americans regularly interact with citizens who belong to an opposing political party (Mutz and Mondak 2006; Corno, La Ferrara and Burns 2022). Unlike nearly any other aspect of social life, individuals often have limited choice over the people with whom they interact at work (i.e., their co-workers), reducing the opportunity for self-selection. Second, while much scholarly attention has focused on partisan sorting by geography, neighbors often interact only superficially, and are less likely to interact at all if they do not share a partisan identity (Parker et al. 2018). By contrast, employees in most worksites typically have no choice but to actively interact and communicate with their co-workers. Finally, Americans simply spend far more of their waking hours at work than in any other activity. The average American spends 3.5 hours per day on work (8 hours among those who are employed), versus only about 15 minutes on "organizational, civic, and religious activities" such as church and volunteering (Bureau of Labor Statistics 2023). The workplace, therefore, may not only constitute an arena for meaningful cross-partisan contact, but also one in which contact occurs frequently and over long periods of time.

Prior literature implies that, due to these characteristics, cross-partisan workplace interactions are likely to reduce cross-party animus (Allport 1979). In fact, the workplace may be one of the few domains that meets Allport's (1979) criteria for prejudice-reducing intergroup contact inasmuch as it frequently is characterized by (a) relatively equal status among participants, (b) a common group goal that requires (c) cooperation across members of different groups, and (d) the support of relevant authorities. Almost by definition, co-workers often cooperate in teams to achieve a common objective that has been set by a manager. To the extent these co-workers hold opposing partisan identities, these workplace interactions may plausibly reduce affective polarization. Relatedly, given that one's work and employer frequently constitute an important social identity (Akerlof and Kranton 2000), workplaces with partisan diversity plausibly cultivate the sort of cross-partisan identities, or "cross-cutting cleavages," that can potentially ameliorate sources of prejudice such as affective polarization (Mutz and Mondak 2006; Gubler and Selway 2012; Wojcieszak and Warner 2020). From this perspective, partisan de-diversification resulting from sorting across workplaces may contribute to affective polarization.

Despite the apparent importance of workplace sorting by partisanship, research on this topic in the United States has focused exclusively on top managers or provided only small-scale experimental evidence that rank-and-file employees might sort by partisanship. Analyses of top managers have found that boards of directors have become increasingly politically homogeneous (Hoang, Ngo and Zhang 2022; Fos, Kempf and Tsoutsoura 2022). Experiments involving employee hiring suggest that rank-and-file employees prefer to work for companies that share their political ideology and/or partisanship (Carpenter and Gong 2015; Burbano 2021; McConnell et al. 2018) and that employers also prefer to hire politically like-minded candidates (Gift and Gift 2015; Colonnelli, Pinho Neto and Teso 2022). In fact, Iyengar and Westwood (2015) posit that, in the absence of legal or normative barriers to employers discriminating on the basis of partisanship, workplace sorting by partisanship may be more pronounced than sorting based on race or gender. Recent work suggests that, at least in Brazil, this may be the case (Colonnelli, Pinho Neto and Teso 2022). We are aware, however, of no large-scale attempt to quantify the extent of partisan sorting in the U.S. labor market as a whole.

1.3 Data and Merging

We estimate partisan sorting in the workplace by linking state voter registration files with a database of online worker profiles. This approach significantly improves upon previous studies of workplace partisanship in the United States in terms of the number and representativeness of workers it captures. Most existing research on firm-level political ideology relies upon data from legally-mandated public disclosure of individual political donations by the Federal Election Commission (FEC).[2] Because these public campaign finance disclosures contain names, addresses, job titles, and employer information, researchers have used them to measure firm-level political ideology (e.g., Gupta and Briscoe 2020; Li 2018; Stuckatz 2022). Many prior studies leverage the extensive Database on Ideology, Money in Politics, and Elections (DIME) data, developed by Bonica (2014), which not only captures this donation data, but also precisely characterizes donors' ideology based on the ideology of the candidates to whom they donate.[3]

[2]Historically, this disclosure was limited to donors who gave more than $200 in a single election cycle, although changes in the way donations are collected (the growing use of donation platforms ActBlue and WinRed) has meant that nearly all donors have been captured in more recent election cycles, including small-dollar donors.

[3]Political partisanship and ideology are theoretically distinct concepts, although they are highly correlated in the modern U.S. political landscape. One major strength of the DIME data is that it can make fine-grained ideological distinctions between more moderate and more extreme members of the same party. As we discuss, our measure of partisanship greatly expands coverage but at the cost of a coarser measurement.

The major trade-off for this measurement precision is narrow and skewed coverage. The vast majority of Americans do not donate to political campaigns, let alone in amounts large enough that historically required disclosure. As a result, donations-based measures cover a narrow and demographically unrepresentative slice of the US population—one that is not only far more politically engaged, but also far whiter, wealthier, and older than the U.S. population as a whole (Grumbach and Sahn 2020; Bonica and Grumbach 2022). Thus while DIME and other campaign finance-based datasets may provide good coverage for highly-educated, wealthy, and politically-engaged professions such as lawyers (Bonica, Chilton and Sen 2016) and doctors (Hersh and Goldenberg 2016), they are likely less representative when when looking at employees or companies in general.

In addition to questions about representativeness, linking FEC data with companies can also be technically challenging. FEC data is based upon reports filed by political campaigns or other political organizations, which must make a good faith effort to collect accurate information from their donors. Unlike with the voter file, however—where registered voters must provide their exact legal name in order to be allowed to vote—there is nothing preventing donors from providing variants or nicknames. Company names are also not standardized, nor can corporate parents easily be linked with subsidiaries (Stuckatz 2022). Donors may choose to withhold information, leading donation filings to read only that information has been "requested." Furthermore, donors may also give technically accurate but misleading information about their employment, such as listing a position in an industry trade association rather than their primary employer—a tactic that that is especially common among firms facing reputational challenges (Shanor, McDonnell and Werner 2022).

Given both the limited coverage and technical challenges involved with linking FEC disclosures with companies, previous efforts to link individual workers' partisanship with their employer have been orders of magnitude smaller than the data we present here. For instance, Stuckatz (2022) deployed extensive code to standardize and parse Federal campaign contributions from 2003 to 2016 and was able to identify only 85,109 individuals across 874 political action committees (PACs) associated with publicly traded companies. In another analysis, Barber and Blake (2023) estimated that 60% of U.S. companies have no donors that appear in DIME.

For some studies—such as those which focus on the ideology of the "upper echelons" within a firm or those which are primarily concerned with PAC contributions—these campaign finance measures may be an appropriate. Given our focus on the general employee population, we believe our approach provides more comprehensive coverage compared to existing donation-based measures. We discuss the representativeness of our sample in more detail in Section 1.3.

L2 Voter File

Our first data source is the national voter file. In the United States, individuals' vote choices are secret, but whether or not they register or turn out to vote is a matter of public record. Voters' registration and turnout history, combined with demographic and contact information, are widely used by political campaigns, commercial data vendors and marketers, and, increasingly, by academic researchers (Hersh 2015). In many states, these data can be accessed by contacting individual state and/or local officials, but researchers more typically rely on data vendors which aggregate, clean, and standardize the raw records received from election officials. Voter file data have been widely used in political science, economics, and other fields, often by merging these data with other datasets to study topics such as voter turnout (Barber and Holbein 2022; Bonica et al. 2021), as well the partisan makeup of various professions and groups, including board members (Fos, Kempf and Tsoutsoura 2022), college students (Firoozi 2023), government bureaucrats (Spenkuch, Teso and Xu 2023), patent examiners (Raffiee, Teodoridis and Fehder 2023), physicians (Hersh and Goldenberg 2016), the police (Ba et al. 2022), religious leaders (Malina and Hersh 2021), and spouses (Hersh and Ghitza 2018). Our voter file comes from the nonpartisan vendor L2 Data and contains information on approximately 185 million registered voters in the United States. Estimates from the 2022 Cooperative Election Study suggest that approximately 73% of the overall U.S. adult population are registered to vote; among those currently employed, it is 76%.

Below, we describe the variables we use from the L2 voter file:

Partisanship: Unlike many other countries, (e.g., Brazil; see Colonnelli, Pinho Neto and Teso 2022), political parties in the United States do not maintain nationwide member rolls. Instead, L2's information on partisan identification come from a variety of sources, which vary by state (Barber and Holbein 2022). In 30 states (plus D.C.), party ID is recorded as part of the voter registration process. In remaining states, L2 imputes partisanship based on other available data, such as voters' primary election participation, past data releases of party voter rolls, or modeling based on ethnicity, geography, and other relevant data.[4] Prior studies which make use of commercial voter file data and have conducted robustness checks excluding states with imputed partisanship have found no difference in results. A report by Pew Research which compared modeled data on commercial voter files with self-reported survey responses also found that modeled partisanship is correct in a majority of cases (Igielnik et al. 2018). Throughout the paper, we focus on registered Democrats and Republicans and treat all other registered voters (i.e., independents, non-partisans, and third-party registrants) as an "other" category.[5]

Race and Ethnicity: We also measure race and ethnicity (i.e., "ethnorace", see Grumbach and Sahn 2020) based on measures provided by L2. Six Southern states (Alabama, Florida, Georgia, Louisiana, North Carolina, and South Carolina) collect data on registrants' race. In other states, L2 imputes race using Bayesian probabilistic modeling techniques (Imai, Olivella and Rosenman 2022; Rosenman, Olivella and Imai 2023) which incorporate the racial makeup of first and last names, as well as decennial census data on the racial makeup of small-scale geographic units (i.e., Census blocks and/or tracts). The racial and ethnic categories used by L2 are European, Hispanic and Portuguese, East and South Asian, African-American, and other. For convenience, we refer

[4]Georgia, Illinois, Indiana, Michigan, Ohio, South Carolina, Texas, Virginia, and Washington have partisan primaries in which voters must choose which party ballot to receive. L2 records partisan primary participation and assigns party based upon the most recent primary in which individuals voted. In Michigan and Washington, partisan affiliation is not currently released by state election authorities, but L2 has access to historical data releases that are no longer publicly available. Finally, in Alabama, Hawaii, Minnesota, Missouri, Montana, North Dakota, and Vermont, partisan affiliation is modeled based upon ethnicity, geography, and other data points.

[5]Because it is difficult to characterize the political ideology of those who are registered independents, officially non-partisan, or otherwise unaffiliated with a party, we avoid drawing inferences about this group.

to these groups as (non-Hispanic) white, Hispanic, Asian, Black, and other; our analyses focus on comparing white workers with workers of color. While assigning racial categories is both theoretically and analytically complex, L2's imputed race data have been widely-used across the social sciences, including by scholars of race who research the disparate racial impacts of election laws on voter turnout (Bonica et al. 2021; Barber and Holbein 2022; Billings et al. 2022). Igielnik et al. (2018) find that modeled race in voter files is typically well-measured overall when compared with survey self-responses, with different voter files vendors accurately classifying between 74–85% of respondents.

Gender: We measure gender based upon information provided in state voter registration files.

MSA: To improve our ability to match between the voter file and employment data, we use home address ZIP code from the voter file to match individuals with the corresponding metropolitan statistical area (MSA).

Political Participation: To measure registered voters' degree of political participation, we include information on voters' turnout in general and primary elections. We also include a binary indicator for whether the individual was a political donor, based upon whether they gave at least one donation to a political campaign, as captured under FEC donation requirements, in the past four election cycles.

Worker Profiles from Revelio Labs

Prior work in other national contexts has leveraged administrative tax or pension data to study workplace partisanship (e.g, Colonnelli, Pinho Neto and Teso 2022). In the United States, federal administrative data are not publicly available, and made available to researchers only under certain narrow circumstances. Statutory guidelines for the use of personally-identifiable tax or social security microdata for research purposes likely preclude their use for studies which focus on political partisanship. For example, researcher data access often relies upon sampling or the use of synthetic and/or coarsened anonymized data, neither of which would be suitable for our purposes.[6]

Because these data are not available, we rely upon data on job positions from Revelio Labs. Revelio is a workforce intelligence company which uses proprietary technology to compile individual employment records based upon online professional profiles (e.g., LinkedIn), as well as job postings.[7] The Revelio data we use contain information on individual job positions, including the

[6]There have been notable cases where researchers have gained access to IRS data to study less politically-charged topics such as economic mobility (e.g., Chetty et al. 2014), but these have only been for subsets (i.e., individual birth cohorts) rather than the entire U.S. workforce, as well as on research topics which are comparatively less politically sensitive relative to partisan politics.

[7]Other competing data vendors include Lightcast (formerly known as Emsi Burning Glass) and LinkUp. Revelio data have been been used in a variety of studies in management, accounting, and other disciplines to study topics including ESG reporting and workplace diversity (Ahn et al. 2023; Baker et al. 2022; Cai, Chen, Rajgopal and Yang 2022; Cai, Dey, Grennan, Pacelli and Qiu 2022; Fadhel et al. 2021), the drivers and consequences of employee turnover (Arif, Yoon and Zhang 2022; Leung et al. 2023; Li et al. 2022), corporate culture (Pacelli, Shi and Zou 2023), and the influence of individual work history and life experiences on business outcomes (Agarwal et al. 2023; Gao, Wu and

date on which they started and ended. Our particular Revelio dataset was captured at the end of August 2022, meaning we did not capture positions that appeared online after this date. Accordingly, we limit our analysis of the contemporary state of the labor force to jobs that were active as of the start of 2022, which is final year in our data with significant coverage. After excluding records with missing data fields (MSA or employer), our Revelio data available for matching cover approximately 90 million positions, held by 76 million unique workers (approximately 46% of the U.S. labor force).

Below, we detail each of the variables we use from this dataset.

Employer: Revelio uses proprietary machine learning techniques to clean and standardize employer names, including grouping variant representations of the same company (e.g., "BofA" and "Bank of America"). Revelio also maps companies to their ultimate corporate parent (e.g., "Facebook" and "Instagram" share the ultimate corporate parent "Meta").

Gender: Revelio models gender probabilistically based upon first name.

Industry: Revelio classifies employers using sector, subsector, and industry NAICS codes.

MSA: Location is mapped based upon the metropolitan statistical area (MSA) associated with an individual work position.

Occupation: Revelio classifies jobs into standardized O*NET-SOC codes using a proprietary machine learning algorithm which incorporates job titles and the responsibilities and descriptions associated with job titles found in online job postings.

Seniority: Revelio models seniority into seven categories ranging from entry level (e.g., intern, trainee) to senior executive level (i.e., C-suite). This modeling is based upon current job title, individual job history (i.e., past employment), and age. Age is not directly reported in Revelio but estimated based upon years of employment history.[8]

Merging Strategy

Merging datasets is not a new challenge in social science research (Enamorado 2021). This particular application presents unique challenges, however. First, the Revelio data are sourced from online employment profiles, rather than from government or administrative sources. As a result, we lack the individual identifiers—dates of birth (DOBs), social security numbers (SSNs), or home addresses—that are commonly used to link datasets. Furthermore, individuals may use an entirely separate name professionally (e.g., a nickname, middle name, or maiden name) that does not correspond to the legal name they use when registering to vote. While social scientists have developed methods for linking administrative datasets, these are typically designed to address partially missing data, data entry errors (e.g., typos), or inconsistencies (e.g., inconsistent usage of abbreviations), rather than cases where bridging data fields are almost entirely absent. A second related challenge is the sheer size of the data. The L2 voter file contains approximately 185 million records and the Revelio data contain approximately 90 million distinct positions we

Zhang 2023).

[8] Personal communication with Hakki Ozdenoren, Revelio Labs.

use for our merge, meaning that a brute force attempt to merge the two datasets would involve approximately 16.5 quadrillion pairwise comparisons.

To surmount these challenges, we begin by partitioning both the Revelio and L2 records by metropolitan statistical area (MSA). Our data snapshots are both from 2022. In L2, we use the most recent available extract of each states' voter file, while in Revelio we make use of employment history covering 2012 to 2022.[9] We exclude records where the work or home address are located outside of an MSA or where location data are missing.[10] This is done for both practical reasons to make the computation feasible, as well as theoretical reasons—MSAs are constructed by the Census Bureau to capture commuting patterns, so there are theoretical reasons to limit our matches to cases where the work location (as reported in Revelio) and home address (as reported in L2) are within the same MSA. We then clean and standardize names by lowercasing and removing punctuation. Our matching procedure is conservative and relies upon exact name matching; we do not attempt to disambiguate between multiple distinct individuals who share a first and last name. To facilitate exact matching on first and last name, we remove duplicate names from both datasets.[11] For MSAs which span multiple states, we attempt to match job positions with voter registrations for all states associated with that MSA (e.g., we would attempt to match job positions located in the Philadelphia-Camden-Wilmington, PA-NJ-DE-MD MSA with voters registered in the MSA in PA, NJ, DE, and MD).

Using this exact matching technique, we are able to match 20,242,686 out of a total of 89,762,656 Revelio records, for a match rate of 23%. When looking at unique individuals (as opposed to job positions), we are able to match 17,092,703 out of 76,430,472 unique workers, or 10% of the entire U.S. workforce (estimated at 163.5 million in 2022). This match rate compares favorably to prior work by Colonnelli, Pinho Neto and Teso (2022), who rely upon administrative data (including a unique tax ID number) and are able to match 7.8% of Brazilian workers with a political party. Since we are interested in the partisan identity of a focal worker's co-workers, our effective sample is necessary limited to those workers for whom we identify at least one coworker, which, as discussed in greater detail below, we define as workers in the same firm and MSA as a focal worker. Applying this restriction, our working dataset contains 16,999,561 positions held by 14,498,184 unique workers. This corresponds to approximately 9% of the U.S. workforce.

Summary Statistics

Table 1.1 presents summary statistics for our working dataset (limited to positions with at least one co-worker that were active at any point prior to end of August 2022), and then separately

[9] This means that we only match workers who were currently registered to vote as of 2022. Most notably, individuals who had work history between 2012–2022 but who died before 2022 are likely removed from the voter rolls and thus not available to be matched.

[10] According to the 2021 American Community Survey (ACS) DP03 table, nearly 88% of Americans in the labor force live within an MSA.

[11] As a result, individual workers with more common names (e.g., "John Smith"), are less likely to appear in our dataset relative to those with more unusual names. This does result in somewhat different match rates for demographic groups, as demographic groups with less unique names (e.g., men, Hispanics) are less likely to be matched than those with more distinctive names (e.g., women, Blacks). Table 1.7 contains match rates by demographic groups.

Table 1.1: Position-level summary statistics for full and partisan samples for the 2022 snapshot

	(1) Mean	(2) Standard Deivation	(3) Min	(4) Max	(5) Obs
For Full Sample					
Democrat	0.43	0.50	0.00	1.00	16,999,561
Republican	0.28	0.45	0.00	1.00	16,999,561
Other Party or Independent	0.29	0.45	0.00	1.00	16,999,561
Woman	0.55	0.50	0.00	1.00	16,999,561
Asian	0.05	0.23	0.00	1.00	13,819,005
Black	0.04	0.18	0.00	1.00	13,819,005
Hispanic	0.11	0.31	0.00	1.00	13,819,005
Other	0.05	0.21	0.00	1.00	13,819,005
White	0.75	0.43	0.00	1.00	13,819,005
Holds Bachelor's	0.89	0.31	0.00	1.00	9,150,523
Holds Graduate Degree	0.35	0.48	0.00	1.00	9,150,523
Percent Dem Coworkers	0.43	0.25	0.00	1.00	16,999,561
Percent Rep Coworkers	0.28	0.23	0.00	1.00	16,999,561
Percent Other Coworkers	0.29	0.22	0.00	1.00	16,999,561
Unique Workers					14,498,184
Unique Positions					16,999,561
For Democrats					
Woman	0.60	0.49	0.00	1.00	7,368,038
Asian	0.06	0.24	0.00	1.00	5,618,874
Black	0.07	0.25	0.00	1.00	5,618,874
Hispanic	0.15	0.36	0.00	1.00	5,618,874
Other	0.05	0.22	0.00	1.00	5,618,874
White	0.67	0.47	0.00	1.00	5,618,874
Holds Bachelor's	0.90	0.30	0.00	1.00	4,126,671
Holds Graduate Degree	0.38	0.49	0.00	1.00	4,126,671
Percent Dem Coworkers	0.51	0.24	0.00	1.00	7,368,038
Percent Rep Coworkers	0.23	0.20	0.00	1.00	7,368,038
Percent Other Coworkers	0.26	0.20	0.00	1.00	7,368,038
Unique Workers					6,244,197
Unique Positions					7,368,038
For Republicans					
Woman	0.52	0.50	0.00	1.00	4,700,918
Asian	0.03	0.17	0.00	1.00	4,072,781
Black	0.00	0.06	0.00	1.00	4,072,781
Hispanic	0.06	0.24	0.00	1.00	4,072,781
Other	0.03	0.18	0.00	1.00	4,072,781
White	0.87	0.33	0.00	1.00	4,072,781
Holds Bachelor's	0.88	0.33	0.00	1.00	2,313,819
Holds Graduate Degree	0.32	0.47	0.00	1.00	2,313,819
Percent Dem Coworkers	0.36	0.24	0.00	1.00	4,700,918
Percent Rep Coworkers	0.38	0.26	0.00	1.00	4,700,918
Percent Other Coworkers	0.26	0.21	0.00	1.00	7,368,038
Unique Workers					4,077,759
Unique Positions					4,700,918

Differences in number of observations across variables within the same subsample reflect missing data. Sample includes jobs active at any point between January and August 2022.

Table 1.2: Employer, occupation, and industry-level summary statistics for the 2022 snapshot

	(1) Full Sample	(2) Democrats	(3) Republicans
Employer Level			
Positions/Employer	19.21	11.15	8.25
	(268.51)	(136.94)	(92.62)
Unique Employers	885,002	660,554	569,567
Positions/Employer-MSA	10.06	5.95	4.40
	(67.88)	(41.66)	(21.49)
Unique Employer-MSAs	1,690,072	1,238,130	1,066,442
Occupation Level			
Positions/Occupation	44,113.88	19,121.22	12,197.27
	(60,940.63)	(28,357.80)	(17,554.17)
Unique Occupations	384	384	384
Industry Level			
Positions/Industry	11,717.86	5,104.94	3183.64
	(49,210.67)	(24,271.28)	(12,097.13)
Unique Industries	1,051	1,048	1,051

Figures are based on the "working sample" which excludes singleton workers (workers without any co-workers in their workplace). Column 1 includes our full sample. Columns 2 and 3 include only Democrats and Republicans, respectively. Sample includes jobs active at any point between January and August 2022.

for Democrats and Republicans within this sample. Table 1.2 presents corresponding summary statistics for this working dataset by employer, occupation, and industry. The sample comprises 885,002 unique employers and 1,690,072 unique employer-MSA workplaces. (Throughout the paper, we refer to employer-MSA combinations as "workplaces".) Consistent with the highly skewed distribution of employees across employers in the United States, most of the employers have a small number of workers, although a small number of employers are very large. The mean workplace in our working dataset has 10.1 positions.[12] 43% of positions are held by Democrats, 27% are held by Republicans, and 29% are held by individuals who are registered with a third party or as independents or non-partisan. While this portion of Democrats appears larger than estimates of the share of the population who identify as Democrats, it is consistent with large-scale survey data which measure partisan *registration* (as opposed to self-identification) among employed individuals. We discuss this at greater length in Section 1.3.

We also use our merged data to create an unbalanced panel of positions spanning 2012 to 2022. We do so by considering all positions active at any point between 2012 to 2022 inclusive (not just those in active employment at the beginning of 2022) held by workers we were able to match to the 2022 L2 state-level voter files. We then use the start and end dates for each position to

[12]The distributions of these two values are reported in Figures A.2 and A.3 in the appendix.

define unique position-year dyads. Here, units of analysis are position-years, rather than simply positions, such that the same position appears separately in each year in which it was active at any time. Naturally, this yields a larger dataset, comprising 121,317,174 unique position-year dyads relating to 32,259,929 unique positions, 14,824,737 unique workers, and 2,246,349 unique employer-MSA dyads. Figure A.4 charts the number of observations by year. Table A.3 shows summary statistics for this unbalanced panel dataset, which are very similar to those for the 2022 data presented in Table 1.1. Since our main objective is to estimate the most current degree of sorting in the United States labor market, we conduct our main benchmarking analysis and sorting estimates with respect to the 2022 data. We return to this panel of workers in Section 1.7 where we examine temporal trends in sorting.

Sample Benchmarking

Next, we benchmark our merged sample to measure the extent to which it is politically representative of the United States labor force. Table 1.3 compares our overall merged sample with employment across industries, as measured by NAICS supersector codes in the 2021 American Community Survey (ACS) from the Census Bureau. Professional sectors in which individuals are more likely to use online job profiles such as LinkedIn—e.g., financial activities (NAICS 52 and 53), government (NAICS 90–93), and information technology (NAICS 51)—are over-represented, whereas more blue-collar industries or industries where online job profiles may be less common—such as leisure and hospitality (NAICS 71, 72), education and health services (NAICS 61 and 62), and retail (NAICS 44–45)—are underrepresented.

We also measure our coverage at the individual occupation level by mapping O*NET occupational codes to O*NET zones. O*NET zones range from one to five and measure the level of preparation required for a position, and broadly map to socio-economic status. Zone one includes jobs which require little education, skills, or experience (e.g., dishwashing or landscaping), while zone five includes jobs which require postsecondary education and extensive knowledge (e.g., lawyer, veterinarian. We present our merged sample versus population estimates of the number of workers in each zone in Table 1.4, along with description of each zone in Table 1.5. Consistent with our benchmarking against industry, our sample includes extensive coverage of workers in higher zones while under-representing those in lower zones relative to their population share.

A potentially much graver threat to our inferences would be if our merged sample were politically unrepresentative. If our sample dramatically overrepresented Democrats relative to Republicans (or vice versa), our inferences about partisan sorting would be suspect. Fortunately, this is not the case.

Because our sample is built off of the L2 voter file of registered voters and the Revelio dataset of employed Americans, we benchmark our sample against the population of working registered voters. Furthermore, because our measure of partisanship is based upon official party registration, we benchmark our sample against official party registration data rather than self-reported partisanship. Accordingly, we compare our sample to the 2022 Cooperative Election Study (CES), a high-quality, nationally-representative, 60,000 respondent survey (Schaffner, Ansolabehere and Shih 2023). Importantly, the CES covers all US adults and features a "validated vote," which con-

Table 1.3: Industry coverage

	Merged Sample	2021 ACS	Difference
Supersector (NAICS Sector #)			
Education and Health Services (61, 62)	0.20	0.23	-0.03
Financial Activities (52, 53)	0.15	0.07	0.08
Manufacturing (31-33)	0.14	0.10	0.04
Professional and Business Services (54-56)	0.13	0.12	0.01
Public Administration (92)	0.09	0.05	0.04
Information (51)	0.08	0.02	0.06
Retail Trade (44-45)	0.08	0.11	-0.04
Leisure and Hospitality (71, 72)	0.04	0.09	-0.05
Transportation, Warehousing, and Utilities (22, 48-49)	0.03	0.06	-0.03
Wholesale Trade (42)	0.03	0.02	0.00
Other Services (except Public Administration) (81)	0.02	0.05	-0.02
Construction (23)	0.02	0.07	-0.05
Natural Resources and Mining (11 and 21)	0.01	0.02	-0.01

Population totals are based upon estimates from the 2021 American Community Survey's DP03 tables.

Table 1.4: Coverage by occupation type

O*NET Zone	Merged Sample	Population	Difference
Zone One	0.02	0.07	-0.05
Zone Two	0.21	0.42	-0.21
Zone Three	0.22	0.22	0.00
Zone Four	0.42	0.25	0.17
Zone Five	0.13	0.05	0.08

Population figures come from Bureau of Labor Statistics (BLS) estimates from the May 2022 BLS data release (`national_M2022_dl.xlsx`), which are merged with O*NET Zone codes from `onetonline.org`. O*NET Zones group occupations based upon the level of education, experience, and on-the-job training required to do a job. O*NET Zone classifications are used in government reporting and for official purposes (e.g., US immigration visa eligibility). See below Table 1.5 for more information about each Zone.

Table 1.5: Details of O*NET Codes

O*NET Zone	Details
Zone One	
Education	High school diploma or GED may be required
Related Experience	Little or no previous work-related skill, knowledge, or experience
Job Training	A few days to a few months
Examples	Food prep workers, dishwashers, landscaping workers, baristas
Zone Two	
Education	Usually require a high school diploma
Related Experience	Some previous work-related skill, knowledge, or experience
Job Training	A few months to one year working with experienced employees
Examples	Orderlies, counter clerks, customer service representatives
Zone Three	
Education	Vocational skills, on-the-job experience, or associate's degree
Related Experience	Previous work-related skill, knowledge, or experience
Job Training	1–2 years of training; both on-the-job and informal training
Examples	Electricians, agricultural technicians, barbers, medical assistants
Zone Four	
Education	Most occupations require a four-year bachelor's degree
Related Experience	Considerable work-related skill, knoweldge, or experience
Job Training	Several years of work-related experience and/or training
Examples	Real estate broker, sales manager, graphic designer, cost estimator
Zone Five	
Education	Most require graduate school
Related Experience	Extensive skill, knowledge, and experience
Job Training	Typically assume that person has required skills and knowledge
Examples	Pharmacists, lawyers, clergy, veterinarians

firms that respondents are registered to vote in their state (Ansolabehere and Hersh 2012). Because our sample relies on state voter files to measure partisanship, this measure on the CES allows us to conduct a like-for-like comparison between our sample and the population of registered voters rather than relying on self-reported partisanship. This comparison is reported in columns 1 and 2 of Table 1.6. For completeness, we also show how our sample compares to the CES's general population of *working* Americans (including those who are not registered to vote) as well as the *overall* general population of Americans, which we characterize in columns 3 and 4.

Table 1.6: Sample demographics vs. 2022 CES and 2020 CCES

	(1) Merged Sample	(2) Working Reg. Voters	(3) Working Gen. Pop.	(4) Gen. Pop.
Party				
Democrat	0.43	0.43	0.34	0.32
Other	0.29	0.25	0.39	0.40
Republican	0.28	0.32	0.27	0.28
Race/Ethnicity				
Person of Color	0.25	0.28	0.31	0.31
White	0.75	0.72	0.69	0.69
Gender				
Man	0.45	0.48	0.52	0.51
Woman	0.55	0.51	0.47	0.48

All columns except "Sample" are calculated using either the 2022 Cooperative Election Study (CES) and are weighted to either the population of registered voters (columns (2) and (3)) or general population (column (4)). Respondents who report working full- or part-time are included in columns 2 and 3. Statistics on party identification differ between registered voters and the general population because partisan identity is measured by registration for the sample and for registered voters but by survey responses (self-reported) for the general population. Race/ethnicity is based upon L2 categories in the sample and upon self-reported race in all other columns. Gender does not sum to 100% because a small portion report a gender other than male or female in the CES.

In both our sample and in column 2, there is a Democratic plurality among registered voters (43% in our sample, 43% in CES), with Republicans comprising a smaller share of the population (28% in our sample, 32% in CES). The share of "others," meaning those that register as non-partisan or under a third party is also similar across the two samples.[13]

Another concern is that our matching technique may result in a sample which is skewed along other demographic dimensions. Because genders and races differ in their tendency to have unique names, our matching strategy could result in a sample which is heavily skewed relative to the underlying population of interest. In Table 1.7 we report match rates by demographic group,

[13]In our data, the vast majority of these "others" (90%) are self-declared non-partisans who are registered as "No Party Preference," "No Party Affiliation," "Independent", or similar label, rather than with a third party such as the Libertarian Party or Green Party.

Table 1.7: Voter file match rates, by demographic

	L2 Records	Matched	Rate
Party			
Democratic	77,628,067	9,713,031	12.5
Republican	56,515,566	6,547,301	11.6
Other	54,732,625	6,709,284	12.3
Gender			
Female	99,211,807	12,946,730	13.0
Male	88,703,294	9,916,414	11.2
Missing	961,157	106,472	11.1
Ethnicity			
White	113,481,101	14,246,435	12.6
Black	20,518,940	2,098,604	10.2
Hispanic	25,625,476	2,020,280	7.9
Asian	7,331,955	982,447	13.4
Other	5,227,502	879,070	16.8
Missing	16,691,284	2,742,780	16.4
Total	**188,876,258**	**22,969,616**	**12.2**

Prior work on demographics note that different groups differ in the uniqueness of their names, and thus our ability to successfully match across datasets. Despite the fact that many married women adopt a new name upon marriage, we have a higher match rate for women due to the fact that women tend to have more unique names then men. Some ethnic groups (Blacks and especially Hispanics) also tend to have less unique names than whites. Overall, we are able to match slightly fewer Republicans. This may be because Republicans tend to be older, whiter, and more male (all characteristics associated with less unique names) and/or because Republicans may be less represented in online professional job postings (e.g., LinkedIn).

and do find that we have slightly lower match rates for men (relative to women) and Hispanics and Blacks (who have fewer unique names and surnames than other ethnicities). Despite the differences in match rates across demographic groups, however, it does not appear that the resulting dataset is overly skewed in relative to the underlying population of interest as seen in Table 1.6.

1.4 Measurement and Main Results

We seek to answer the question: When people go to work, to what extent are they surrounded by members of the same political party? Notably, we are interested in individual-level measurements that capture the experience of the average American *worker*, not firm-level measurements that characterize the condition of the average American *company*. Although measuring the partisan diversity within individual firms is an important question worthy of future research, we focus in this paper on measuring the extent to which individual workers are exposed to partisan diversity among their co-workers. Given this focus, our measurement differs subtly, but importantly, from existing work focused on partisan diversity within boards (Fos, Kempf and Tsoutsoura 2022) or manager-employee dyads (Colonnelli, Pinho Neto and Teso 2022).

Our approach is most closely related to prior work by Brown and Enos (2021) which captures partisans' geographic segregation by using residential addresses in voter file data to construct individual-level measures of the partisan sorting by residence. Brown and Enos measure residential partisanship by looking at the partisanship of the 1,000 nearest neighbors for each American (using residential addresses found in the voter file). We adopt a similar approach: for each worker, we conceptualize his or her co-workers as the set of all positions that share a workplace—i.e., are at the same company and located within the same MSA (excluding himself or herself).[14] Formally we define the share Democratic coworkers for a given position as,

$$ShareDem_{i,e,m} = \frac{\sum_{j=1,j\neq i}^{n_{e,m}} Dem_{j,e,m}}{n_{e,m} - 1}$$

where $ShareDem_{i,e,m}$ is the share of coworkers for a worker's focal position i at employer e and MSA m that are Democrats, $n_{e,m}$ is the number of positions in firm e in MSA m, and $p \in \{Dem, Rep, Other\}$ is the partisanship of coworker j, such that $\mathbb{1}(p_j = Dem)$ is equal to one if coworker j is a Democrat and zero otherwise.[15]

While some companies (e.g., chain retailers) may have many locations within the same MSA, the population of workers within a given MSA represents an estimate of the population of workers with whom a focal worker could plausibly interact with. Workers at the same company but in different MSAs are far less likely to interact. Our main analysis focuses on workplaces with at least two employees (as a focal worker must have at least one colleague for our measure to be defined), but our results are robust to a range of higher number-of-coworker inclusion thresholds. Sub-figures A and B of Figure 1.1 present kernel density plots of the overall distribution of the share of workers that are Democrats and Republicans, illustrating that the average worker (regardless of his or her party) experiences significantly more Democrat coworkers (43%) than Republican coworkers (28%).

Sub-figures C and D of Figure 1.1 provide the first evidence of sorting, showing that the average number of Democrat and Republican coworkers varies significantly according to the partisanship of the focal worker. Under the null of no partisan sorting, Democrats and Republicans would each have the same number of Democratic co-workers. In fact, we observe that the mean Democrat's coworkers are 51% Democrats, while the mean Republican's coworkers are 36% Democrats—a 15 percentage point difference-in-means. Reversing our analysis, and focusing on the percentage of Republicans, we see that the mean Democrat's coworkers are 23% Republican, while the mean Republican's coworkers are 38%— again a 15 percentage point difference-in-means. The symmetry of these results reflects the fact that Democrats and Republicans have, on average, the same share of co-workers that are "Others" as shown in Table 1.1.

[14] Brown and Enos (2021) also weight the 1,000 nearest neighbors by geodesic distance based on address, under the assumption that people are more likely to interact with nearby neighbors.

[15] The share of Republican coworkers is defined analogously.

Figure 1.1: Nationwide distribution of coworkers that are Democrats and Republicans, overall and by focal worker partisanship

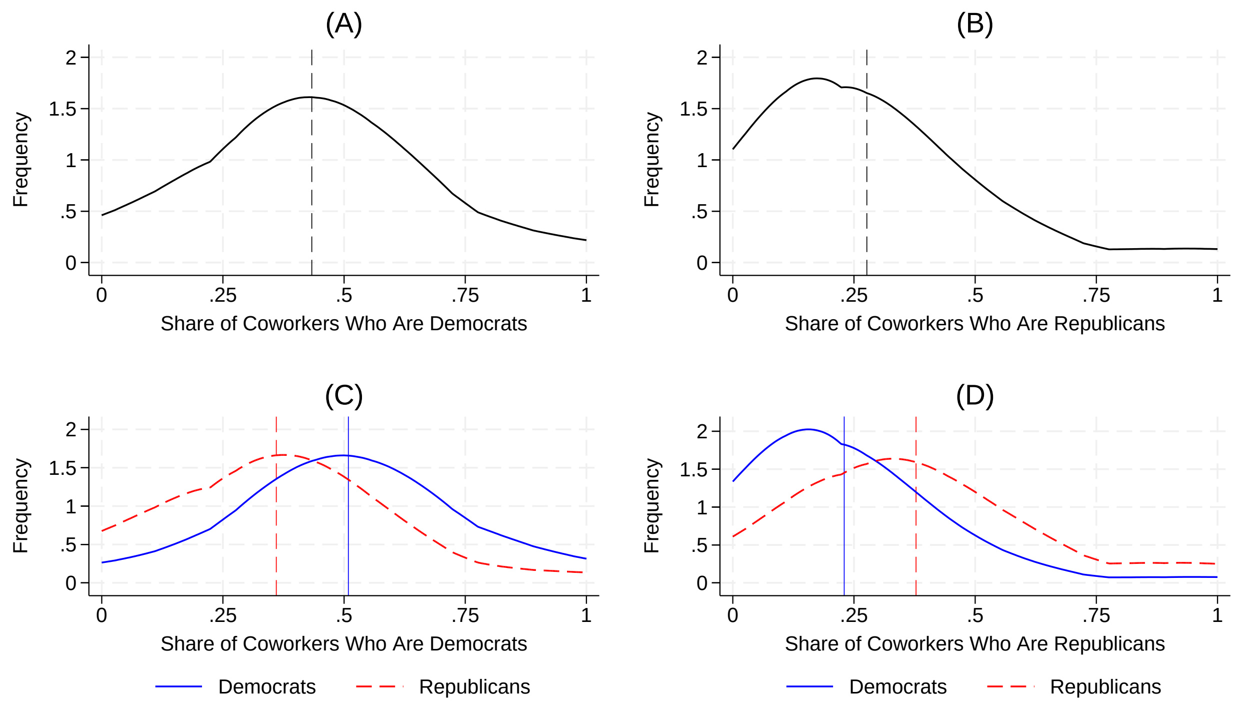

Sub-figures (A) and (B) show the overall distribution in the share of coworkers, by position, that are Democrats and Republicans, respectively. Sub-figures (C) and (D) present these distributions separately for Democrats and Republicans. All plots are kernel density plots with a bandwidth of 0.10. Vertical lines indicate means. N=16,498,184 for Figures (A) and (B). Figures (C) and (D) exclude “Others” and thus have a smaller sample size, N=10,321,956. Figure A.5 in the appendix presents these patterns as histograms rather than kernel density plots.

1.5 Accounting for Geography, Occupation, and Industry

Next, we estimate the extent to which partisan sorting is associated with other important correlates of both work outcomes and partisanship: namely, geography, occupation, and industry. Our objective in this section is descriptive rather than deductive; we do not claim that geography, industry, or occupation are causes of partisan sorting. In actuality, it is very likely that partisanship, workplace sorting, geography, industry, and occupation are co-determined and may mutually reinforce one another. By parsing these different factors, however, we aim to highlight the importance of future research into, for example, why partisans sort into different industries and occupations, since these processes may play an in important role in determining overall levels of sorting. Additionally, we aim to estimate the extent of sorting that remains after we account for these factors, as this residual sorting may be reflective of taste-based sorting on the part of employers and/or job-seekers.

Figure 1.2: Partisanship by MSA

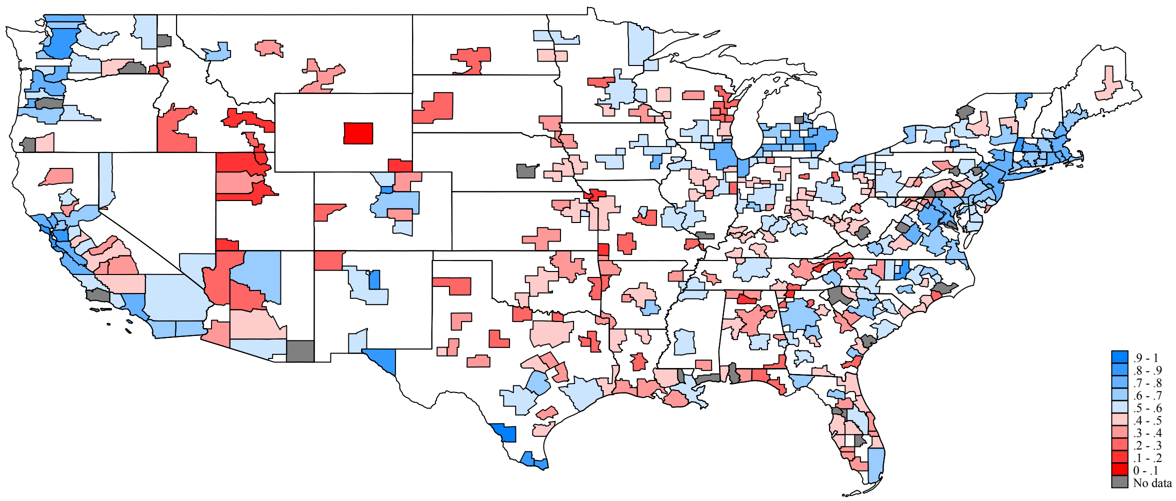

Shaded areas indicate, for our merged sample, the two-party share of workers in the MSA who are Democrats.

Before conducting regression analysis, we provide summary evidence that Democrats and Republican workers are sorted by geography, occupation, and industry. Examining geography, Figure 1.2 maps the share of workers that are Democrats by MSA. The map is consistent with well-documented geographic sorting, with Democrats comprising a larger share of the population in larger cities and on the coasts. [16] Examining occupation, Figure 1.3 plots the distribution of the share of Democrats within each of the 384 occupations within the O*NET-SOC taxonomy. While the modal occupation has roughly the same portion of Democrats as our overall sample, occupations follow a roughly normal distribution, with some being much more Democrat and others much less. Figures A.6a and A.6b show the top-10 most Democrat and Republican occupations. Whereas the most Democrat occupations comprise creative and academic occupations (e.g., editors, English professors, historians, actors) and some blue collar occupations commonly held by workers of color (housekeepers and packagers), the most Republican occupations include commercial airline pilots, and managerial positions in engineering and construction (e.g., bio-fuels, manufacturing engineers, boiler operators). Finally, we plot the share of Democrats within each NAICS (5-digit) industry codes, and then plot this distribution in Figure 1.4. Similarly to occupations, there is considerable variance. Figures A.6c and A.6d lists the top industries by partisanship, showing patterns that mirror those for occupation. Whereas Democrats are most dominant in creative industries and in lower-paying service industries, Republicans are most dominant in natural resources extraction, engineering, and manufacturing.

To quantify how this sorting by geography, occupation, and industry contribute to the overall degree of workplace sorting, we fit a series of regression models in which we sequentially

[16] Reassuringly, our measure of MSA-level partisanship among workers is highly correlated with MSA-level Democratic vote share in the 2020 presidential election ($\rho = 0.69$).

Figure 1.3: Share of partisan workers that are democrat, by occupation

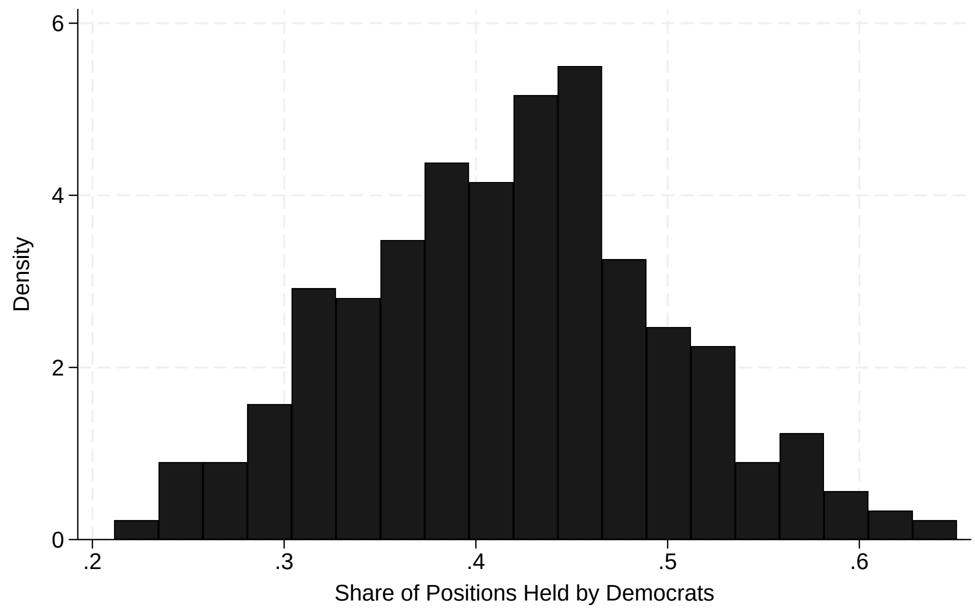

This figure depicts the distribution of the share of Democrats all 384 O*NET occupation designations.

incorporate these dimensions as fixed effects. We first estimate a baseline specification that is the regression analogue to the difference-in-means in the share of Democratic co-workers between Democrats and Republicans. Specifically, we estimate

$$ShareDem_{i,e,m} = \alpha + \beta Dem_{i,e,m} + \epsilon_{i,e,m}, \tag{1.1}$$

where $ShareDem_{i,j}$ is the share of co-workers for worker i in employer e and MSA m that are Democrats, Dem is a dummy equal to one if the worker identifies as a Democrat, and ϵ is the error term. β is the coefficient of interest and represents the difference, on average, in the share of co-workers that are Democrats for Democrat workers compared to Republicans. Under the null of no sorting, our estimate of this coefficient should be statistically indistinguishable from zero, which would mean Democrats and Republicans tend to work in places that have the same shares of partisan co-workers that are Democrats. We cluster standard errors by MSA and estimate the model using Ordinary Least Squares (OLS). These results are reported in the row labeled "Raw" in Figure 1.5.

Consistent with the patterns visible in Figure 1.1 and reported in Table 1.1, Democrats, compared to Republicans, work in firms in which the share of co-workers that are Democrats is 14.8 percentage points greater. Given our large sample, this and all other estimates in this subsection are very precise (p-value < 0.001).

Next, we examine to what extent this sorting can be explained by geographic sorting. Given

Figure 1.4: Share of partisan workers that are Democrats, by industry

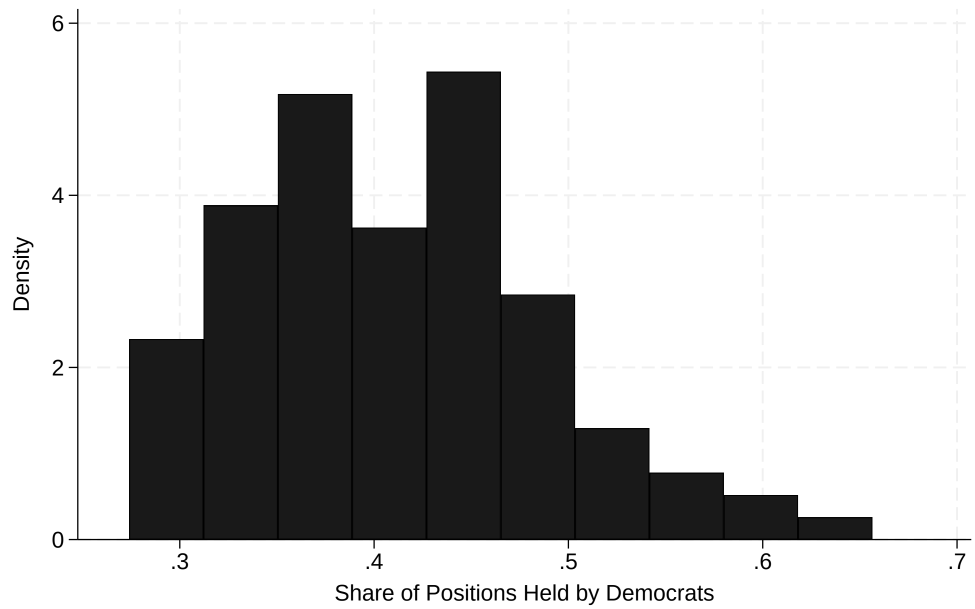

This figure depicts the distribution of the share of Democrats across 1,051 NAICS 5-digit industries represented in the data.

the sorting patterns illustrated in Figure 1.2, it is plausible that copartisan workers tend to work together in part because they live close to one another. To account for geographic sorting, we estimate

$$ShareDem_{i,e,m} = \alpha + \beta Dem_{i,e,m} + \lambda_m + \epsilon_{i,e,m} \tag{1.2}$$

which is identical to Equation 1.1 except for the inclusion of municipality fixed effects, λ_m. This result is reported in the second row of Figure 1.5, labelled "MSA FE." Here, the estimate of β decreases to 9.3 percentage points. This suggests that geographic sorting accounts for about 35% of workplace partisan sorting, but also that workplaces are significantly sorted by partisanship, even after taking into account the partisan composition of their surrounding geographies.

Next, we consider sorting by occupation. It may be the case that workers of certain partisanships are not only more likely to live in certain areas, but also more likely to pursue employment in certain occupations. If this is the case, the observed patterns may arise not only because certain geographies are more likely to be home to certain partisan groups, but also because firms are more likely to employ those within certain occupations. Accordingly, we estimate a model that is identical to Equation 1.2 except we replace MSA fixed effects, λ_m, with occupation fixed effects, λ_o, captured by the O*NET-SOC occupation categories. This result is reported in the third row of Figure 1.5, labeled "Occupation FE." Here, the estimate of β decreases to 13 percentage points. This again suggests that occupational sorting by partisanship explains a meaningful amount of the overall sorting, but apparently less than sorting by geography. Even after taking into account

Figure 1.5: Estimated partisan sorting benchmarked against sorting by race and gender

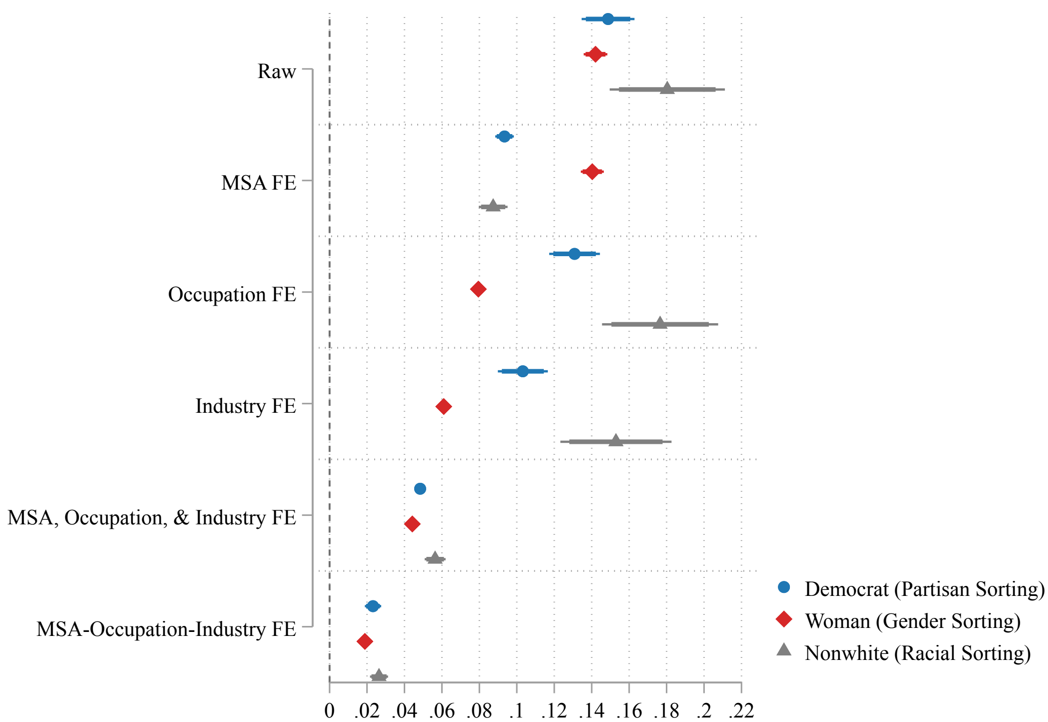

Point estimates are from versions of Equation 1.1 and Equation 1.2 that differ in terms of whether they characterize partisan, gender, or racial workplace sorting and the fixed effects they include. Standard errors clustered by MSA. Bars represent 90% and 95% confidence intervals.

the partisan composition of occupations, workers are still considerably sorted by partisanship across workplaces.

Next, we account for sorting that might occur in terms of industry. We again estimate a version of Equation 1.2 but replace MSA fixed effects λ_m with industry fixed effects λ_d, captured by the five digit NAICS code. This is reported in the fourth row of Figure 1.5, labelled "Industry FE." Here, the estimate of β is 10.3 percentage points. The change in our estimate of sorting is slightly larger than that observed with MSA fixed effects—sorting across industries explains some, but not all, of the partisan sorting in the U.S. work force.

Finally, we consider these explanatory factors together. We first estimate a version of Equation 1.2 but include each of the separate fixed effects λ_m, λ_o, and λ_d. Here, our estimate of β is 4.8 percentage points. We then estimate a version of Equation 1.2 where we include a single MSA-occupation-industry fixed effect λ_{mod}. Here, our estimate of β is 2.3 percentage points ($p < 0.001$). This estimate is 84% smaller relative to our "raw" estimates (without fixed effects), suggesting that geography, industry, and occupation explain a significant portion of workplace sorting. However, a considerable amount of sorting still remains unexplained, which could be consistent with sorting explicitly motivated by partisanship.

Benchmarking Against Workplace Sorting by Gender and Race

To help illustrate the substantive magnitude of these findings, we follow Colonnelli, Pinho Neto and Teso (2022) and benchmark these estimates of partisan sorting with estimates of sorting based on race and gender. Evidence of considerable labor market sorting by gender has inspired extensive research aimed at unearthing its origins, historical development, and present-day consequences (Barbulescu and Bidwell 2013; Blau and Kahn 2017; Goldin 2021; Goldin et al. 2017). To benchmark, we repeat the analysis described in the preceding subsection, but replace (1) the outcome variable $ShareDem$ with the variable $ShareWomen$, which, for the focal position, is the share of women coworkers, and (2) the explanatory variable Dem with $Woman$, which is a dummy equal to one if the focal worker is a woman. These results are reported side by side with our main estimates in Figure 1.5. Overall, workers are about 0.6 percentage points less sorted by gender ($\hat{\beta}$ = 0.079) than they are by partisanship ($\hat{\beta}$ = 0.148). Unsurprisingly, geographic sorting explains more partisan sorting than gender sorting. After accounting for gender sorting by occupation ($\hat{\beta}$ = 0.079), industry ($\hat{\beta}$ = 0.060), as well as as each of these in conjunction with MSA ($\hat{\beta}$ = 0.044 and $\hat{\beta}$ = 0.018), sorting by partisanship is slightly larger in magnitude than sorting by gender.

Next, we benchmark our estimates of workplace sorting by partisanship against workplace sorting by race. Due to America's long history of legal racial segregation and enduring systemic racism, America has high and extensively documented racial sorting in the workplace (Goldman, Gupta and Israelsen 2021; Wilson, Miller and Kassa 2021). To benchmark, we re-estimate Equation 1.1 but replace (1) the variable $ShareDem$ with the variable $SharePOC$, which, for the focal position, is the share of coworkers that are people of color (i.e., not non-Hispanic whites), and the (2) the explanatory variable Dem with POC, which is a dummy to equal one if the focal worker is a worker of color. These results are presented side-by-side with our main estimate in Figure 1.5. Sorting based on race is about 3.2 percentage points greater than sorting based on partisanship ($\hat{\beta}$ = 0.18). However, after taking into account the well-documented fact that the United States is strongly geographically sorted by race, partisan sorting is similar to sorting by race ($\hat{\beta}$ = 0.087). Accounting for occupation and industry, sorting by race is still significantly larger than sorting by partisanship. As with gender, in the multiple fixed effects models the degrees of sorting by race and partisanship are similar ($\hat{\beta}$ = 0.056 and $\hat{\beta}$ = 0.026). Overall, our benchmarking analysis indicates that, after accounting for sorting by MSA, occupation, and/or industry, the estimated degree of workplace sorting by partisanship is similar to workplace sorting by gender and race.

Robustness Specifications

We made several methodological choices in measuring workplace sorting by partisanship, including our decision to exclude independents and non-partisans and clustering standard errors by MSA. To demonstrate that our conclusions are not contingent upon these choices, we re-estimate Equations 1.1 and 1.2, but include workers who identify as independent, those affiliated with third parties, or those who do not disclose their partisanship. Additionally, we demonstrate our findings are robust to excluding firms with fewer than five, ten, and fifty workers. These results are depicted in the specification curves in Figures A.9 and A.10. Furthermore, Figures A.11 and A.12 corroborate the robustness of our results with regressions that include MSA-occupation-industry fixed effects.

1.6 Examining Heterogeneity

Next, we examine heterogeneity in sorting by worker demographics, experience, and political activity, as well as by geography. Characterizing heterogeneity not only provides a finer portrait regarding where and among which groups workplace partisan sorting is most pronounced, but also may shed preliminary light on additional mechanisms driving this sorting. We examine heterogeneity by estimating Equations 1.1 and 1.2 separately for different subgroups of workers. Figure 1.6 summarizes these results.

We first examine heterogeneity by worker demographics. We find little evidence that the degree of sorting varies by worker gender or education. White workers, however, are more sorted than non-white workers of color, who may have reduced labor market flexibility due to discrimination. Similarly, younger workers are slightly more sorted than more elderly workers (Baby Boomers in particular) who may face diminished market power as they age. Consistent with prior work by Fos, Kempf and Tsoutsoura (2022), workers in more senior leadership positions appear more sorted, a pattern which becomes especially stark after including MSA fixed effects. We also explore heterogeneity by job zone, which is a measure developed by the U.S. Department of Labor to classify jobs into five categories based on levels of education, experience, and training necessary to perform the occupation. Occupations in Zone 1 require little or no preparation (e.g., janitorial services), whereas occupations in Zone 5 require extensive preparation (e.g., physicians and physicists). We observe that workers in job zones which require greater skill and experience exhibit more sorting than those in lower zones (two and three). Although by no means conclusive, these patterns are consistent with the idea that workers may have a preference for partisan homophily at work, and that workers with greater market power are better positioned to achieve this.

Second, we examine heterogeneity in terms of the degree to which a focal worker is politically active. We find that sorting is greater among those that voted in the 2020 general presidential election, and especially pronounced among those that voted in the 2020 presidential primary. Furthermore, sorting is significantly greater among that made a political donation at some point in the last four election cycles. These patterns are consistent with the idea that those who are most politically active may place the greatest premium on working alongside those that share their partisanship.

Third, we consider heterogeneity by geography by re-estimating Equation 1.1 separately for each MSA. We map these coefficients in Figure 1.7. Sorting is especially pronounced within the Southern U.S., the “Acela Corridor” running from Boston to Washington, and in the Midwest around Chicago. While further research is necessary to explain these patterns, one possibility—especially in the South—is that racial sorting between political parties and among workplaces may explain higher-than-average rates of partisan sorting in these MSAs.

Figure 1.6: Heterogeneity in Partisan Sorting

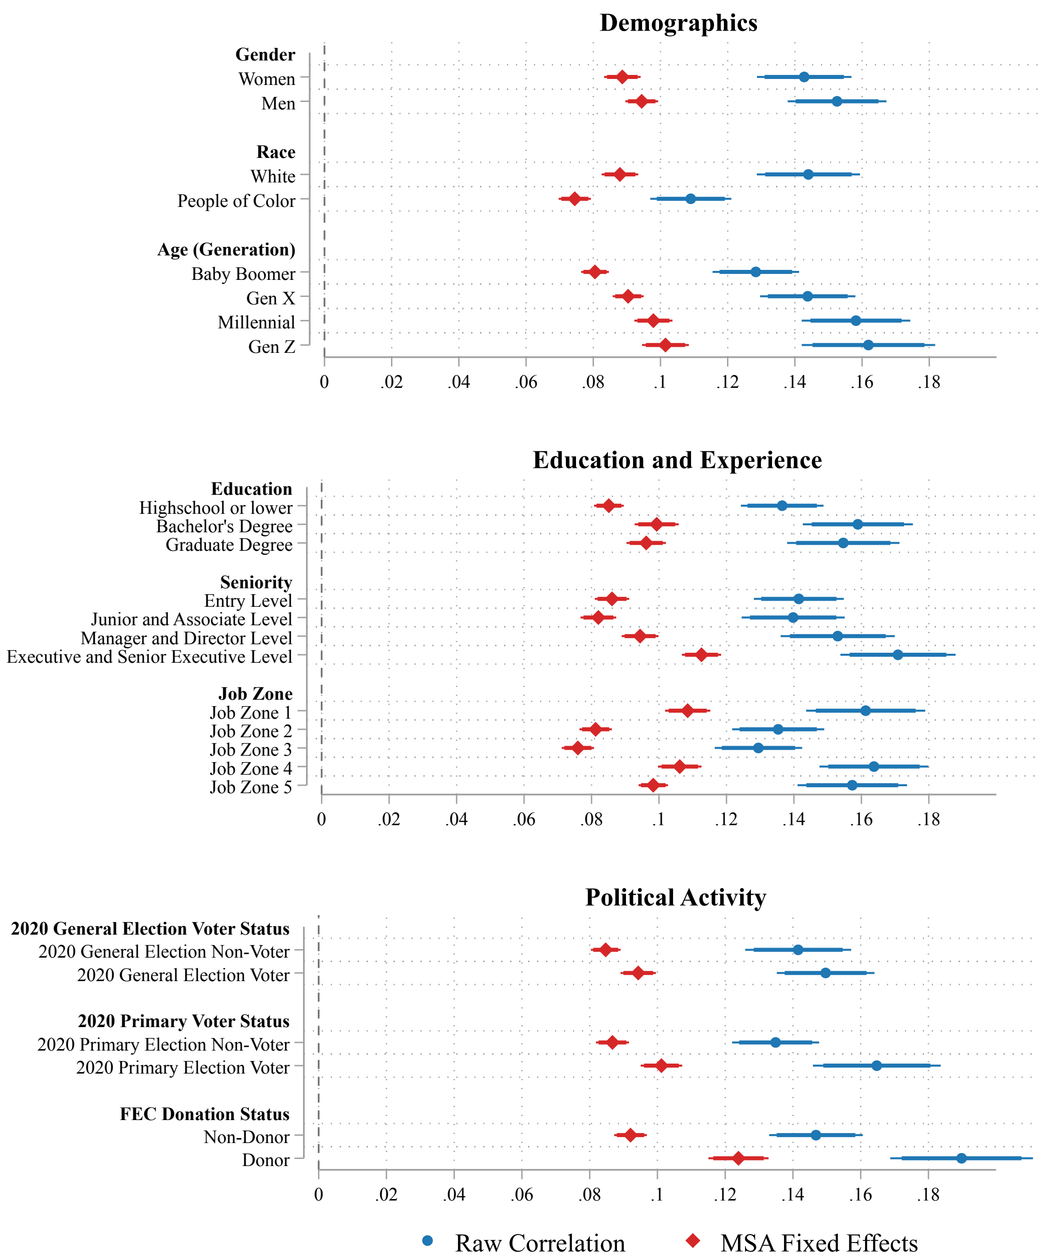

Estimated effect and 90% and 95% confidence intervals of being a Democrat on the Share of Democrats in the employer-MSA. Standard errors clustered by MSA.

Figure 1.7: Partisan sorting by MSA

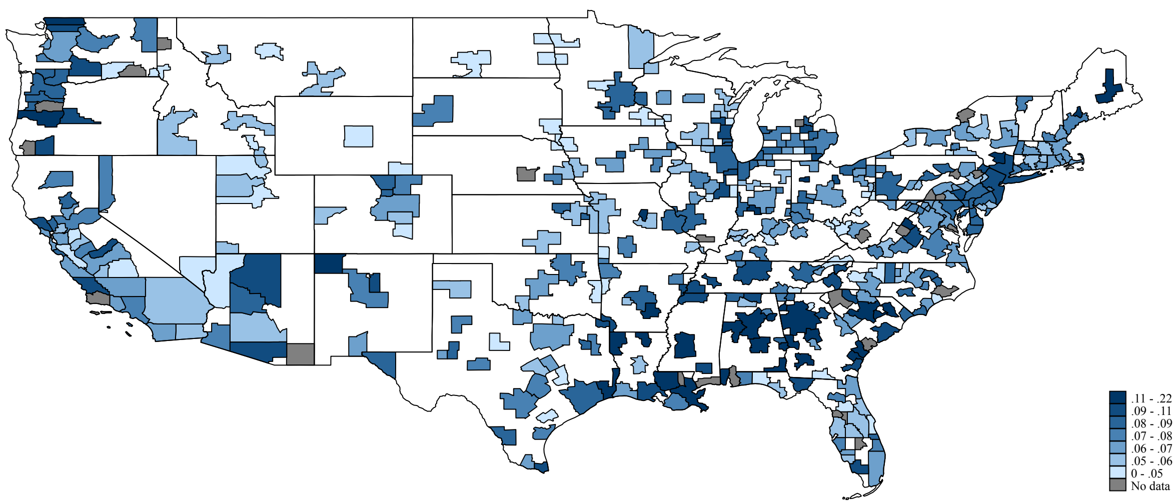

NOTES: Shaded areas indicate the degree of sorting in the indicated MSA, measured as by-MSA estimates of β from Equation 1.1. White designates areas that are not part of any MSA, which tend to be areas with very low population density.

1.7 Longitudinal Analysis

Next, we examine temporal trends in workplace sorting by partisanship. A key aspect of contemporary popular and academic discourse around U.S. politics is that Americans are *increasingly* sorted along partisan lines. We examine temporal trends using the panel described above in which observations are position-years. We then re-estimate Equations 1.1, 1.2, and the variation of Equation 1.2 that includes fully-interacted MSA-occupation-industry, fixed effects. We plot the by-year point estimates and confidence intervals in Figure 1.8a. The results indicate that the degree of overall partisan sorting has increased only very slightly over the past decade. The stability of this estimate over time, however, likely reflects the fact that job changes occur relatively infrequently. Accordingly, we replicate the analysis from Figure 1.8a, but restrict the sample to instances where workers start a new position. This captures workers who are either (1) entering the labor market for the first time or (2) switching employers. These results, reported in Figure 1.8b, indicate that new joiners increasingly sort into employers in which a higher share of coworkers share their partisanship. For the raw estimate, the degree of sorting rises about 10%, from 15.9 percentage points in 2012 to 17.6 percentage points in 2022. For the fully-interacted fixed effects model the estimate rises about 43% from 2.6 percentage points in 2012 to around 3.7 percentage points in 2022.

1.8 Partisan Asymmetry in Out-Party Exposure

Finally, we demonstrate that the degree of out-partisan exposure is asymmetric across the two parties, with Republican workers, on average, experiencing significantly more exposure to Demo-

Figure 1.8: Longitudinal Analysis of Partisan Sorting

(a) Full Sample

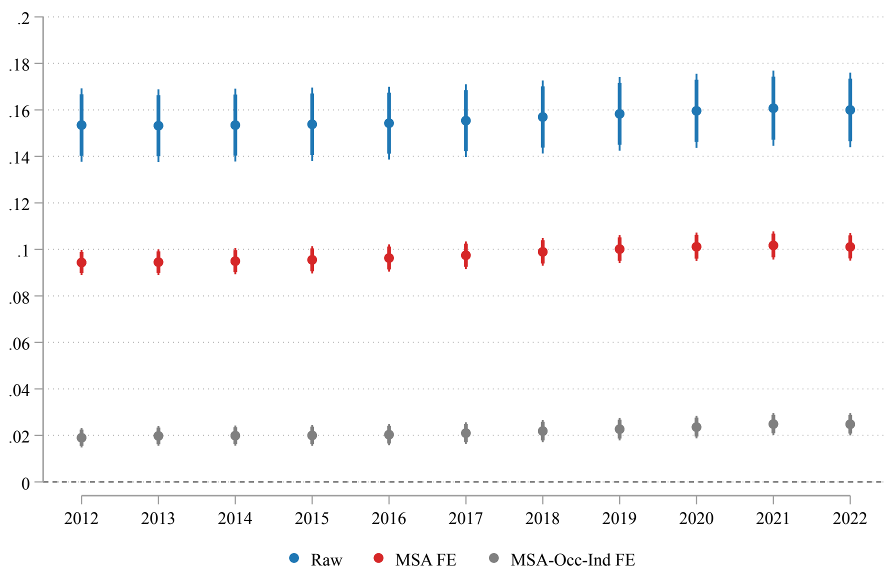

(b) New Positions

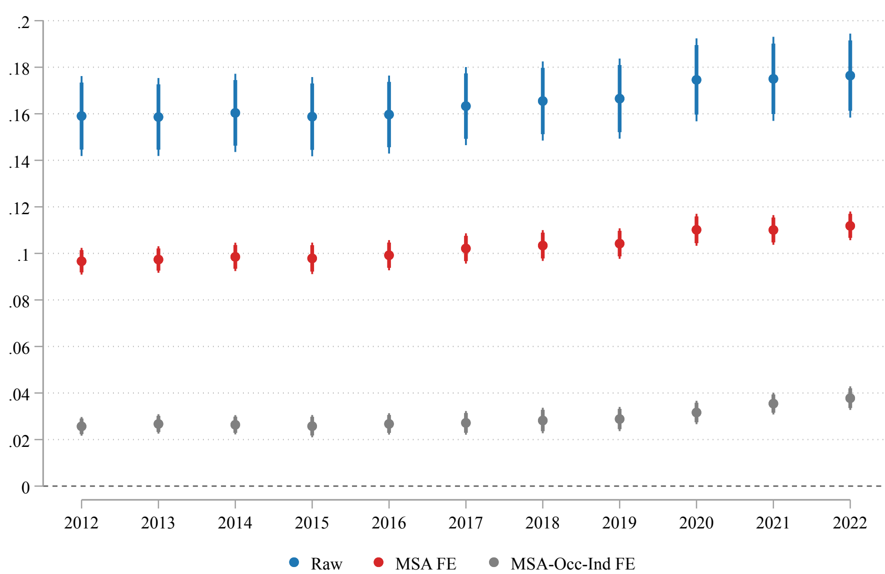

crats than *vice versa*. This is illustrated by Figure 1.1 and summary statistics in Table 1.1. Democrats, on average, hold positions in which 51% of coworkers are Democrats and 23% are Republicans (an in-party to out-party ratio of 2.2). By contrast, Republicans, on average, hold positions in which 38% of coworkers are Republicans and 36% of coworkers are Democrats (a ratio of 1.1). Limiting to two-party share (only consider Democrat and Republican co-workers), the out-party exposure rates for Democrats and Republicans are 37% and 56%, respectively. Moving beyond a simple means comparison, Figure 1.9 presents deciles of out-party exposure by partisanship including all workers. Figure A.13 limits this to two party share. Asymmetric exposure is partly due to the fact that Republican workers are simply less prevalent than Democrat workers, a characteristic of our sample that is consistent with nationally representative samples (see Table 1.6). But this baseline expectation of asymmetric exposure is significantly magnified by the fact that workers are sorted by partisanship.

Figure 1.9: Distribution of in-party and out-party exposure, by partisanship

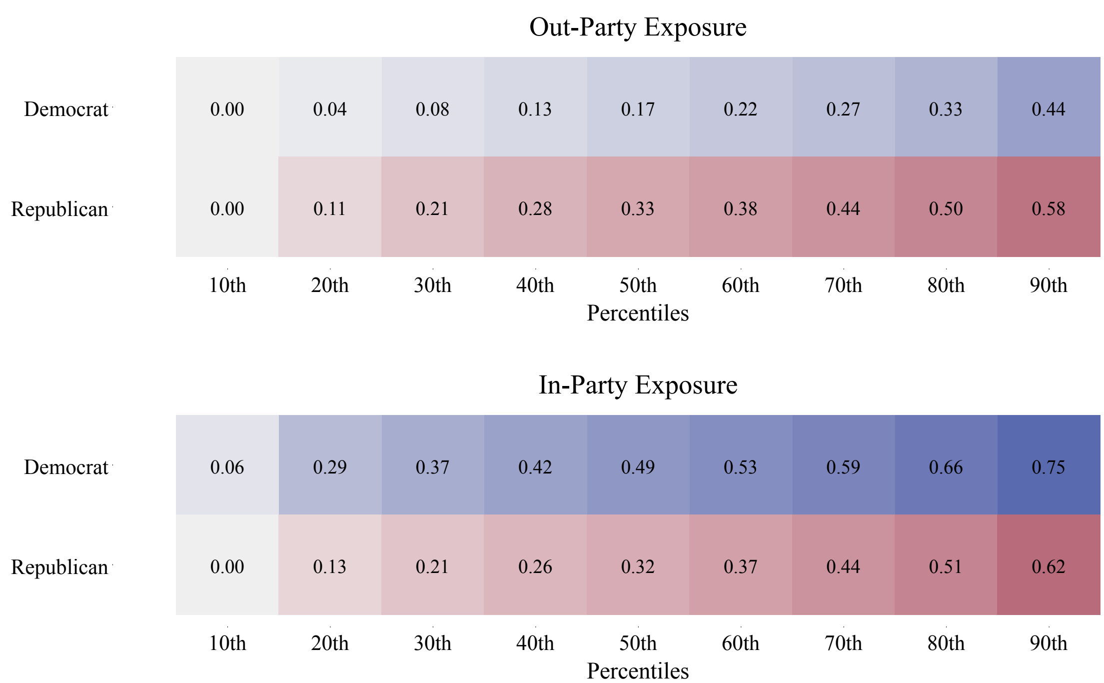

This figure displays the share of in-party and out-party exposure by exposure decile. Democrats have consistently more in-party exposure and less out-party exposure than Republicans. Figure A.13 is an alternative version using only the two-party-share (excluding "others" who are non-partisans, independents, or third-party registrants).

This partisan asymmetry may have significant consequences. Intergroup contact is more likely to reduce prejudice if groups are perceived to hold roughly equal status (Allport 1979). To the extent that demographic majorities are correlated with status, this may not be the case. This asymmetry may also explain the widespread perception, held by both the public and by corporate leaders, that corporate America is moving leftward (Hersh and Shah 2023). If the mean Republican employee's workplace is roughly balanced between partisans but the mean Democrat's workplace is dominated by Democrats, the overall effect may be to shift corporate America towards the Democratic party.

While this exposure may be asymmetric, it is important to note that these figures nevertheless represent meaningful—albeit perhaps lopsided—levels of cross-partisan exposure for most workers. In fact, our two-party-share estimates of out-party exposure are larger than estimates of residential out-party exposure from Brown and Enos (2021), which estimate mean out-party exposure rates of 30% for Democrats and 36% for Republicans.[17] Even if exposure is asymmetric, the workplace does appear to remain a potential locus for cross-partisan contact.

1.9 Conclusion

Merging workers' employment history with voter registration files, we provide the first large-scale estimates of workplace-level sorting by political partisanship in the United States. We present four main findings. First, we provide evidence that workers are sorted by partisanship across different worksites, and that much—but not all—of this sorting is explained by geography, occupation, and industry. Accounting for these factors, the estimated magnitude of sorting by partisanship is comparable to sorting between whites and people of color and between men and women across workplaces. Second, we show that this sorting appears more pronounced among workers in more senior leadership positions, white workers, and workers who are more politically active. Third, we show that partisan sorting appears to be increasing, as new workers increasingly sort into firms where a greater share of co-workers share their partisanship, and that this is true even when controlling for geography, occupation, and industry. Finally, we present evidence that the implications of this sorting appear asymmetric: the average Republican has for more exposure to Democratic co-workers than the average Republican has to Democratic co-workers.

Our findings motivate future research regarding the *measurement*, *origins*, and *consequences* of partisan sorting in the workplace.

Regarding *measurement*, our approach supplements existing campaign finance-based methods of measuring employees' political ideology. While we believe our measure constitutes an important advance, there is ample room to further improve and expand our measurement of cross-partisan contact. For example, our data proxy workplace-level contact by treating employer-MSAs dyads as the unit of analysis, but subsequent inquiry could use surveys or observational data to identify conversation networks and explore the extent to which workplace partisan diversity is actually associated with greater rates of cross-partisan interaction.

Regarding *origins*, future work can disentangle to what extent sorting is driven by worker supply, employer demand, or even socialization. Sorting may be driven by supply side forces

[17] Brown and Enos (2021) use a slightly different methodology. They impute Democratic or Republican partisanship for registered voters who do not identify as either of these and weight neighbors' partisanship according to their proximity to the focal voter.

to the extent employees prefer and are able to sort into jobs with more copartisan coworkers. This might occur as a part of the hiring process, where employees discern partisan "fit" based on employer characteristics with partisan connotations (Burbano 2021). This seems especially plausible given that employers increasingly take explicit stances on partisan issues (Larcker, Miles and Tayan 2018; Bondi, Burbano and Dell'Acqua 2022). Moreover, sorting could occur via differential attrition where members of the minority party are more likely to leave their job (Bermiss and McDonald 2018). Our evidence of heterogeneity in the tendency of partisans to sort suggests several specific hypotheses for supply-side sorting. For example, we find that sorting is more pronounced among workers with more market power, such as top executives, who have more financial resources and employment capital, and white workers, who may face fewer obstacles to employment than workers of color. Similarly, our finding that partisan sorting is greater among those who are more politically active suggests that those who place a higher premium on having co-partisan coworkers may be more likely to sort.

Sorting at work, however, may also be driven by demand-side discrimination in hiring (Gift and Gift 2015; Colonnelli, Pinho Neto and Teso 2022). In many U.S. states, it is not illegal for employers to discriminate based upon political partisanship. Employers or hiring managers may prefer to work with co-partisans for personal or ideological reasons. Even managers who do not otherwise care about the political beliefs of their subordinates may prefer to hire politically-like minded candidates to limit employee turnover and promote a harmonious workplace.

Third, patterns of partisan homophily across coworkers might reflect socialization. There is evidence of peer-to-peer partisan socialization among neighbors Brown (2023) and college classmates (Firoozi 2023), as well as in partisan workplaces (i.e., politicians' offices) (Jones 2013), so it seems possible that sustained workplace exposure could have a similarly influential effect on partisan identity. In this telling, sorting occurs not (or at least not only) because workers choose to work alongside those that share their partisanship, but because they come to take on the partisanship of those they work with. Such socialization processes might also be enhanced by the shared perception among coworkers that their employer is more likely to benefit from the policies of one party over the other. Such perceptions might even be shaped by communication from top managers (Hertel-Fernandez 2018), which might further motivate coworkers to converge on a shared partisan identity.

Regarding *consequences*, future work can examine the effects of partisan sorting for democratic society writ large, firms' engagement with social issues, and even firm performance. First, and perhaps most importantly, research is urgently needed to understand the consequences of rising workplace partisan sorting for major questions of political polarization and democracy. The workplace context is perhaps uniquely well-suited to foster the type of intergroup contact which can potentially reduce prejudice (Allport 1979; Mutz and Mondak 2006). Given growing concerns about sharp rises in political discord and affective polarization in the United States, it is imperative to understand whether the workplace can still serves as a locus of prejudice-reducing cross-partisan contact.

Second, workplace sorting may be linked to the growing tendency of businesses to take public positions on controversial social issues (Larcker, Miles and Tayan 2018; Bondi, Burbano and Dell'Acqua 2022). These stances have often appeared puzzling, given that the benefits from appealing to like-minded stakeholders are often outweighed by the negative reactions from opposing stakeholders (Burbano 2021; Hou and Poliquin 2023). One suggestion is that employees' partisanship may be correlated with firm engagement on controversial political and social controversies (e.g. Li and Disalvo 2022), and thus, as workplaces become increasingly sorted by partisanship, firms will confront increasingly homogeneous workforces and may thus face greater incentive to engage with controversial issues. Research able to map over-time trends in sorting to corporate sociopolitical activism may be able to elucidate disagreement regarding whether corporate sociopolitical positioning is primarily a strategic response to workers' demands as opposed to a unilateral expression of top managers' beliefs or values (Hambrick and Wowak 2021; Mohliver, Crilly and Kaul 2023; Hurst 2023).

Finally, workplace partisan sorting may have consequences for firm performance. On one hand, political homogeneity arising from partisan sorting might enhance firm performance. Employees who feel politically misaligned with upper management or their peers may be less economically productive (Besley and Ghatak 2005; Spenkuch, Teso and Xu 2023; Burbano 2021; Carpenter and Gong 2015) and more likely to exit the firm (Bermiss and McDonald 2018). Moreover, if workers value politically homogeneous workplaces where they are in the majority, they may even accept a lower wage for these jobs (Burbano 2016; McConnell et al. 2018). On the other hand, political sorting might erode firm performance. Studies of historical labor force segregation based upon gender and race have shown that resulting (mis)allocations of talent across industries and occupation can dramatically hinder economic growth (Hsieh et al. 2019). To the extent that partisan sorting may also result in a sub-optimal allocation of talent relative to individuals' innate talents and abilities (Roy 1951), sorting may have significant economic consequences. in this vein, work by Colonnelli, Pinho Neto and Teso (2022) in Brazil points to Becker's (1957) theory of taste-based discrimination to explain a documented negative correlation between indulging a taste for partisan homophily and firm profitability. Understating the relationship between partisan sorting and firm productivity, innovation, and profitability represents an important frontier for future scholarship.

Chapter 2

Do Conservative Policies Harm Access to Employee Talent? Evidence from the North Carolina Bathroom Bill

> We write with concerns about legislation you signed into law last week, HB 2, which has overturned protections for LGBT people and sanctioned discrimination across North Carolina… The business community, by and large, has consistently communicated to lawmakers at every level that such laws are bad for our employees and bad for business. This is not a direction in which states move when they are seeking to provide successful, thriving hubs for business and economic development. We believe that HB2 will make it far more challenging for businesses across the state to recruit and retain the nation's best and brightest workers and attract the most talented students from across the nation.
>
> Letter to North Carolina Governor Pat McCrory signed by over 100 CEOs
> (Human Rights Campaign 2016)

2.1 Introduction

In recent years, companies and their leaders have increasingly taken public stances on sociopolitical issues such as abortion, climate change, gun control, LGBT rights, racial inequality, and voter identification laws. While business has long sought to influence politics (Werner 2018; Vogel 1978), scholars distinguish contemporary forms of "corporate sociopolitical activism" (CSA) from more traditional "corporate political activity" (CPA) in that CSA focuses on issues that are not directly related to core business activities (Chatterji and Toffel 2019; Dodd and Supa 2014; Larcker, Miles and Tayan 2018). As some scholars note, however, this distinction between "core" and "non-core" issues is not always straightforward in practice (Hambrick and Wowak 2021). Companies that engage in CSA often claim that the issues on which they speak out are in fact closely linked to their business performance through their impact on stakeholder relationships (Bondi, Burbano and Dell'Acqua 2022). Given the importance of stakeholder relationships to firm success (Freeman 1984), these companies argue that taking stances that appeal to their stakeholders is beneficial—if not strictly necessary—for their performance.

While these arguments can apply to any stakeholder group, they are especially frequent with regard to employees. Sometimes, companies' explanations for their employee-motivated activism appear consistent with the idea that they view these issues as important for the stakeholders, but "non-core" with regard to business interests. For instance, in 2021, business leaders spoke out against changes to Georgia state election laws, arguing that these laws would disproportionately impact Black voters. Delta Airlines CEO Ed Bastian explained his company's decision to criticize the law by referencing conversations with "employees in the Black community." Similarly, Microsoft President Brad Smith argued that "a healthy business requires a healthy community… and a healthy community requires that everyone has the right to vote conveniently, safely, and securely" (Gelles 2021). While these business leaders recognized that these issues were important to their employees, they did not clearly articulate how the proposed election law changes would directly impact their business.

Yet in many other cases, business have made a stronger and more explicit claim that government policies directly harm their business, especially by making it harder for them to attract and retain employees. This claim is frequently made when companies oppose conservative government policies on issues such as abortion and LGBT rights. After the Supreme Court's draft decision in *Dobbs v. Jackson Women's Health Organization* was leaked in 2022, *Business Insider* published an article dramatically entitled "The country's biggest employers, including Walmart and Amazon, should 'say goodbye to attracting top female talent' in abortion 'trigger law' states" (Cain et al. 2022). Similar arguments are also made with regard to states which have passed conservative policies related to LGBT rights. In 2023, the Society for Human Resource Management (SHRM) published an article entitled "Many workers are leaving states that pass anti-LGBT laws" (Ladika 2023).

While these claims are supported by vivid anecdotal reporting, we lack large-scale scientific evidence to evaluate whether or not these conservative policies are actually detrimental to companies' ability to hire and retain workers. Despite growing scholarly attention to various aspects of CSA, including the influence of employee preferences over companies' CSA (Li and Disalvo 2022; Maks-Solomon and Drewry 2021), the reasons why business leaders feel the need to signal a sociopolitical stance to appeal to employees (Appels 2023; Hurst 2023), and how current and prospective employees actually react to CSA (Appels 2023; Burbano 2021; Wowak, Busenbark and Hambrick 2022), we have little evidence that evaluates the major argument that companies often use to explain employee-centric CSA: namely, do government policies actually affect companies' ability to attract and retain employees?

To address this gap, I examine a highly publicized episode where a state adopted a conservative social policy over the opposition of many prominent voices in the business community. In 2016, North Carolina's Republican legislature passed HB2, a controversial "bathroom bill" which strictly required transgender people to use the bathroom consistent with the sex listed on their birth certificate.[1] As part of the broad national backlash, over 100 CEOs signed an open letter opposing the bill in which they argued that the bill would make North Carolina a less attractive destination for employees (Wowak, Busenbark and Hambrick 2022). This bill was nevertheless enacted over the objections of the business community. Using large-scale administrative data from the US Census Bureau's Longitudinal Employer-Household Dynamics (LEHD) program that has near-universal coverage of turnover in the US labor force (Hyatt et al. 2017), combined with recent advances in causal inference (Xu 2017), I assess the impact of HB2 on worker migration to and from North Carolina.

I find no evidence that HB2 had a negative North Carolina employers' ability to attract and retain employees. Across all educational categories—including college-educated workers—in-migration from other states to North Carolina did not decrease and out-migration from North Carolina did not increase. These null findings are precisely estimated. Further evidence that HB2 had minimal impact comes from the wages of migrating workers. Prior research suggests that employers in "stigmatized" locations must offer higher wages to attract employees (Hurst 2023). If North Carolina were less attractive to workers as a result of HB2, we would expect that wage differential for workers moving to North Carolina would increase, as employers offer higher wages to compensate for the fact North Carolina was less appealing than job options in other states (Smith 1776). I find no evidence of a positive wage differential, and weak evidence to suggest that relative wage differentials may have actually *decreased* for workers coming to North Carolina.

Together, these findings suggest that the common argument that conservative state-level policies harm businesses' access to workers may be overstated. This does not necessarily mean that there is no business case for CSA in opposition to conservative social policies, but it does suggest that one plank of this business case may be weaker than is commonly believed. While previous scholarship has documented that businesses may be worried about the risks of inaction—i.e., guilt-by-association if they do not take stances in opposition to conservative legislation passed in their state (Hurst 2023)—this research suggests that businesses' subjective concerns about the risks of sociopolitical stigma may be overly pessimistic.

On the other hand, there may be other reasons for companies to engage in CSA. Even if conservative policies do not harm access to talent overall, some research suggests that stance-taking may nevertheless serve as a useful signal to potential employees (Appels 2023), possibly leading to better alignment between specific employees and firms and thus increasing productivity (Barber and Blake 2023). Some specific firms or organizations that rely upon highly mobile elite workers or where workers may be especially concerned about sociopolitical stigma may face specific idiosyncratic hiring challenges. These findings may also strengthen arguments about whether CEO's activism should be viewed more as a credible signal of CEO's personal beliefs and ideology or instead as a reflection of other stakeholders' preferences (Appels 2023; Branicki et al. 2021; Nalick et al. 2016).

These findings also may have important limitations. While transgender rights are an important and highly controversial ongoing political issue, the number of transgender people in the

[1] "Bathroom bills" and transgender rights more broadly have featured in a number of studies, including Burbano (2021); Appels (2023) and Wowak, Busenbark and Hambrick (2022).

United States remains extremely low relative to the size of the overall workforce. It is possible that we would observe larger migration effects on issues which have the potential to directly impact a broader swatch of the population—such as abortion restrictions, which potentially impact all women of reproductive age (as well as their male partners).

2.2 Theory

Government Policies and Access to Talent

Why do firms engage in corporate sociopolitical activism? A common argument is that CSA is a reaction to stakeholder preferences. In the most dramatic cases, this may be done in response to pressure campaigns from stakeholders, such as boycotts (from customers), proxy resolutions (from investors), or strikes and walkouts (from employees). In other cases, companies may respond to demands from stakeholders that are conveyed in more subtle and less visible efforts at persuasion. Finally, companies may decide to proactively engage in CSA in an attempt to appeal to stakeholder preferences (or what they imagine stakeholder preferences to be), even in the absence of explicit pressure from these stakeholders (Bondi, Burbano and Dell'Acqua 2022).

In this paper, I start by taking firms at their word. Firms frequently argue that the speak out against (typically conservative) policies that they argue will harm their ability to hire and retain talented workers.[2] For instance, in the quote reproduced at the beginning of this article, over 100 CEOs co-signed in a letter in which they argued that North Carolina's HB2 "bathroom bill" regarding transgender bathroom access would "make it far more challenging for businesses across the state to recruit and retain the nation's best and brightest workers" (Human Rights Campaign 2016).

How might government policies harm businesses' ability to attract and retain employees? In extreme cases, policies may directly restrict employers' access to hire otherwise qualified employees. For example, under regimes of legalized racial discrimination or strictly enforced gender segregation, employers would find it illegal or prohibitively difficult to hire racial minorities or women, shrinking the size of the labor market and making it more difficult for them to hire. Less stringent policies may also have a similar effect if they make it dangerous or risky for workers to accept jobs within a certain geography. For instance, employers located in a jurisdiction that strictly enforces sodomy laws against same-sex relations will likely struggle to hire gay employees. In the current United States context, this might mean that companies located in states which adopt legislation that makes life difficult, unpleasant, or risky for LGBT people will have a harder time hiring LGBT workers because there will be fewer of them available to hire in the local talent market.

While this is a major concern for specific populations (such as transgender people), the overall direct effect on employers is likely to be minimal in many cases. At the time North Carolina passed

[2]There does not seem to be a theoretical reason why firms could not speak out against liberal policies, but in practice, there have been virtually no large mainstream firms that that take explicitly conservative stance on issues (Bhagwat et al. 2020; Hambrick and Wowak 2021; Maks-Solomon 2020).

HB2, estimates put the total number of transgender adults in North Carolina at only 22,000 out of a total population of 10.2 million (Mallory and Sears 2016). But in addition to the immediate effects on those directly targeted (i.e., transgender people), employers must also consider how these laws these impact their ability to hire and retain employees overall. Accordingly, the CEO's open letter makes it clear that the business community was concerned with the signal that the law would send to prospective employees in general. The letter did not argue that the bill would harm North Carolina's ability to attract and retain transgender (or LGBT) workers specifically, but rather that it would negatively impact the states' ability to attract the "best and the brightest workers" and "most talented students"—workers and students of whatever gender and sexual orientation (Human Rights Campaign 2016).

Accordingly, these types of laws may harm companies' ability to attract and retain workers in several ways (apart from outright bans). First, laws may influence companies' access to talent by serving as a signal regarding the conditions of employment. Employment with a company is often conceived of as an incomplete contract that involves an experience good. While prospective employees can negotiate many aspects of their job (e.g., wages and benefits), it is often difficult for them to understand important non-pecuniary aspects of their job, such as organizational culture. Because employees must make these decisions under some uncertainty, they rely upon signals and heuristics that may convey useful information about culture and other difficult-to-observe aspects of the job (Bangerter, Roulin and König 2012; Connelly et al. 2011). In turn, companies can attempt to send signals that convey useful information about what it is like to work at the company (Bangerter, Roulin and König 2012). (While this information may indeed be useful, job seekers will also rationally discount this information, as companies have an incentive to portray employment conditions in the most positive possible light.) Specifically, these signals may convey information about what types of employees are likely to fit in and be successful at the company (Chatman 1989; Schneider, Goldstiein and Smith 1995).

Recent work exploring the role of signaling in person-organization (PO) fit on has shown that women respond differently to company signals such as prosocial claims by management (Abraham and Burbano 2022) or job advertising claims to have a "flat" organizational structure (Hurst, Lee and Frake 2024). Similarly, workers with different political beliefs or policy preferences may also view signals differently as either attractive or unattractive. When it comes to companies hiring in jurisdictions with a certain set of government policies, job seekers may infer that the internal company culture reflects the politics of the local area, and their co-workers are likely to agree with these policies. This is not an unreasonable assumption. Not only are workers drawn from the local area, but recent work suggests that companies strategically choose where to locate facilities to minimize political distance with existing facilities—in other words, a company headquartered in conservative Alabama may foresee personnel issues if it opens a new facility in liberal Massachusetts, and instead prefer to expand in a more ideologically similar area (Barber and Blake 2023).

When given a signal of the likely partisan tilt of their prospective employer and co-workers, these job seekers may thus use these signals to infer that they are unlikely to be a good fit for the workplace (Appels 2023). While the PO fit literature has long examined demographic factors such as gender and race, recent work examines political partisanship and finds evidence of partisan homophily in both the United States labor market (Chapter 1, this volume) as well as elsewhere (Colonnelli, Pinho Neto and Teso 2022). This research also suggests that that employers may actively discriminate against members of opposing parties (Colonnelli, Pinho Neto and Teso 2022; Gift and Gift 2015) and that partisan "misfits" are more likely to depart their firm (Bermiss and McDonald 2018), further suggesting the value of political alignment.

Second, potential workers may adopt a similar lens focused on not employment with a specific company, but rather to the state (or other geography) as a whole. Just as it is not fully possible to understand what it is like to work at a specific company *ex ante*, it may not always be possible to understand what it is like to live in an unfamiliar state. The political voting patterns and policies of the state are a highly visible and potentially useful proxy for potential migrants in determining if they would enjoy the experience of living in a certain state.

Scholars in labor economics, political science, geography, and other disciplines have long debated whether this type of "Tiebout" (1956) sorting along political lines actually occurs, given the substantial expense involved in moving and the myriad other factors that influence geographic mobility. While there is evidence that this politically-motivated migration does occur, the effects are quite small. Americans are increasingly geographically polarized along political lines (i.e., Republican "red" rural areas and Democratic ''blue" cities) (Brown et al. 2023; Brown and Enos 2021) but it appears that this is almost entirely driven by conservatives and liberals holding different preferences over neighborhood amenities rather by partisans trying to live alongside fellow partisans or partisans preferring jurisdictions that enact their preferred policies (Martin and Webster 2020; Mummolo and Nall 2017). Recent work does suggest that younger college-educated Americans—who, on average, hold far more liberal social policy attitudes that older or less-educated citizens—are less likely to move to states with Republican control of state government, although the overall effect is not huge in absolute terms (a 0.4% annual reduction in migration rates) (Downey and Liu 2023). A persistent finding is that the largest determinants of interstate migration rates are not political but economic (i.e., relative wages and house prices) (Olney and Thompson 2024), for more local moves, the major factors are economic and lifestyle considerations such as proximity to jobs, school quality, housing affordability, and neighborhood safety (Martin and Webster 2020; Mummolo and Nall 2017).

How Might CSA Help with Employee Talent?

Firms' access to talent may be affected by social policies in their local areas. First, employees may observe government policies and infer that employers in the area are likely to support those policies. Second, government policies may influence the general attractiveness of the region and result in population growth or less—although most research suggests this effect is likely to be small.

This dynamic may help explain—at least in part—why firms frequently engage in CSA even on issues where these stances may lead to a backlash. While early work on CSA was optimistic about its potential influence public opinion without provoking negative reactions (Chatterji and Toffel 2019), more recent work shows that CSA may result in backlash from customers (Hou and Poliquin 2023; Korschun, Martin and Vadakkepatt 2020; Liaukonytė, Tuchman and Zhu 2022), employees (Burbano 2021), and investors (Bedendo and Siming 2021; Bhagwat et al. 2020; Pasirayi, Fennell and Follmer 2023). There have also been prominent incidents in which CSA appears to lead to backlash from governments and politicians. In 2021, after Coca-Cola and Delta Airlines expressed concerns that changes to Georgia's election law were intended to disenfranchise Black voters, Republican Senator Mitch McConnell warned the companies of "serious consequences" and suggested that Republicans might engage in a boycott. Since then, Republican-led states have carried out punishments against companies that oppose their policies. These include Texas, which has banned banks which adopt ESG policies from doing business with state and local governments, and Florida, which responded to Disney's opposition to the "Parental Rights in Education Act" (which critics argued was anti-LGBT) by stripping the company's *de facto* right to control local government.

Given this potential for backlash from customers, employees, investors, and government regulators, scholars have asked why firms nevertheless engage in risky corporate activism (Bondi, Burbano and Dell'Acqua 2022; Hurst 2023; Mohliver, Crilly and Kaul 2023). In many cases, it seems that firms weigh the risks of speaking out or remaining silent and the resulting benefits of pleasing some stakeholders while potentially upsetting others (Bondi, Burbano and Dell'Acqua 2022; Melloni, Patacconi and Vikander 2023). This has long been an element of niche business strategies for smaller "companies with a conscience" where sociopolitical alignment is a core differentiating factor (Vogel 2006), but as sociopolitical activism has become more common, large mainstream firms are also becoming active on these issues. This may because of an overall shift in both social pressure and strategic equilibria strategies as more firm speak out, which in turn force other firms to speak out lest their silence be taken as a tacit endorsement of a given side (Melloni, Patacconi and Vikander 2023).

These trade-offs across multiple stakeholders may be especially likely to pit employees against other stakeholders. While the direction of these preferences could in theory push firms in any direction, in practice, it seems almost universally the case that employees push their firms towards more liberal directions (Maks-Solomon 2020; Hersh and Shah 2023). First, research in political science and other disciplines shows that younger, more highly-educated individuals and people residing in dense urban areas hold very liberal and progressive views (Downey and Liu 2023; Mason 2015, 2018; Zingher 2022). Second, these younger generations of workers may hold different attitudes about the appropriateness of companies and other organizations to public stances on issues relative to older generations (Barzuza, Curtis and Webber 2021). Third, these especially liberal contingents—young, educated, and urban workers—may hold disproportionate leverage over companies. In general, employees are critical for firm performance (Coff and Kryscynski 2011; Cowgill and Perkowski 2020), but high-skill employees are especially important from a value creation standpoint given their contribution to firm growth and innovation (Florida 2014), which afford them greater bargaining power, which in turn makes CSA more likely (Melloni, Patacconi and Vikander 2023). High-skill employees are also increasingly difficult for companies to find (Black, Hasan and Koning 2024), especially given long-term decreases in on-the-job training (Cappelli 2015, 2019). This high criticality and scarcity affords these (disproportionately liberal) high-skill workers a high degree of leverage; scholars have noted that employees are especially able to persuade management, given their greater knowledge of internal culture and power structures (Briscoe, Chin and Hambrick 2014; Briscoe and Safford 2008; Morrill, Zald and Rao 2003; Raeburn 2004). From a more sociological perspective, these employees may also have more influence over firm leadership because of cultural similarity (e.g., through shared educational backgrounds) as well as physical proximity (e.g., they work at company headquarters rather than in a branch location). For all of these reasons, firms may be especially responsive to the concerns of employees.

The composition of younger and better-educated employees can help explain why companies may take public stances in opposition to socially conservative policies (and conversely, why they rarely take stances in opposition to socially liberal policies). Firms may not be able to actually influence whether these policies are enacted. While older political economy theorists—many influenced by Marxian analysis—argue that business holds almost unparalleled influence over American society (e.g. Lindblom 1982), more recent empirical work finds that the ability of business to shape public opinion through CSA is quite modest (Chatterji and Toffel 2019) and that business is frequently disunified (Smith 2000) and often does not get its preferred policies on controversial and highly publicized issues (Culpepper 2010). Yet even if firms are not able to directly influence government policy through CSA, they may reason that these stances serve as effective signals to stakeholder groups such as employees. From this perspective, the fact that these stances may engender backlash may actually strengthen the signal—classical game theory holds that a "costly" signal should be viewed more credibly than one that is merely "cheap talk" (Appels 2023).

Does This Logic Hold?

While theoretically plausible, prior studies have suggested that this logic may not actually hold in the real world. To-date, studies have considered whether corporations' use of CSA is beneficial in attracting and retaining employees. These studies on the effectiveness of CSA as a signal to employees have found mixed results. Wowak, Busenbark and Hambrick (2022) argue that sociopolitical activism polarizes employee organizational commitment. They find that among firms with a greater share of Democratic donors among their employees, Glassdoor ratings improved after CEOs took a stance in support of transgender rights, while the reverse was true for firms with a greater share of Republican donors. Hurst (2023) looks at employers in Charlottesville, Virginia, who faced a "sociopolitical stigma" after a white supremacist march in their town. Despite investing significantly in making countervailing prodiversity claims (e.g., in recruitment language), these employers still had to offer quite significant positive wage differentials to successfully recruit employees, suggesting that the effect of the signals was limited. Finally, Burbano (2021) actually finds that sociopolitical signals may be counterproductive. In a field experiment involving a firm hiring employees to complete an online task, potential employees who agreed with the firms' stance (on transgender rights) were actually *less* effective.

To-date, however, no study has investigated the implied underlying mechanism: namely, the idea that (conservative) policies that (liberal) employees oppose actually have a negative effect on access to employee talent. Despite how common these claims are, the political science literature on residential sorting suggests reasons for skepticism on the tendency of employees to actually relocate in response to government policy shifts. While there is significant evidence that Republicans and Democrats are highly sorted (Brown et al. 2023; Brown and Enos 2021), our best evidence also suggests that this is *not* because partisans explicitly move to be close to co-partisans (Martin and Webster 2020; Mummolo and Nall 2017). Within the workplace, Bermiss and McDonald (2018) do find that workers sort over-time to workplaces with more politically-like minded coworkers, but this finding comes from highly mobile, politically involved, and economically well-off workers in the private equity industry—and even then, the pace at which partisans sort themselves alongside like-minded coworkers is slow and gradual.

Recent work on partisan sorting within corporations also suggests similar dynamics to those found in interstate migration. In experimental settings, both employers and employees appear to prefer to work alongside co-partisans (Colonnelli, Pinho Neto and Teso 2022; Gift and Gift 2015; Lelkes and Westwood 2017). In real-world settings, however, their ability to sort based upon these preferences appears more limited. While there is substantial homophily within the workplace (i.e., Democrats tend to work alongside other Democrats and Republicans with other Republicans), this appears mostly driven by partisan differences in geography, industry, and occupations rather than by partisans explicitly changing jobs to work alongside fellow partisans (Chapter 1, this volume)—a finding that parallels similar results in geographic sorting.

For all these reasons, it is not clear whether we should expect that policies will affect companies' access to talent. In the next section, I describe the empirical setting I use to test these claims: namely, HB2, a highly-publicized "bathroom bill" in North Carolina that became one of the largest political stories of 2016.

2.3 Empirical Context: North Carolina HB2

In February 2016, the city council in Charlotte, North Carolina passed a local ordinance banning discrimination based upon sexual orientation or gender identity in public accommodations, for-hire vehicles, and city contracting. One month later (March 2016), the Republican-controlled state legislature called a one-day special legislative session to overrule the Charlotte ordinance. The resulting bill was formally known as the "Public Facilities Privacy & Securities Act" or House Bill 2 (HB2), but quickly became known as the "bathroom bill." The bill not only required people to use the bathroom consistent with the sex listed on their birth certificate, but also prohibited local governments (such as Charlotte) from passing their own anti-discrimination legislation.

The bill was widely criticized by LGBT organizations and Democratic politicians. The Obama administration issued administrative guidelines that directly contradicted the bill, putting the state at risk of loosing billions of dollars in Federal grant funding for education, housing, and other government areas. A number of Democratically-controlled state and local governments banned publicly-funded travel by state employees to North Carolina, and media reports suggested that a number of conferences scheduled to take place in the state were cancelled or relocated from the state. Musicians including Bruce Springsteen, Ringo Starr, and Pearl Jam cancelled scheduled concerts. Finally, the NCAA voted to remove college sports championships scheduled to take place in the state—a major concern in North Carolina, which is known for its avid college basketball fandom.

Members of the business community also publicly condemned the bill. A small number of companies announced they would abandon plans to open new facilities in North Carolina. Most prominently, Paypal announced in early April that it was cancelling plans to open a facility in Charlotte that would have created an estimated 400 jobs (Rothacker, Portillo and Peralta 2016). In addition to individual business decisions, the Human Rights Campaign, a large non-profit focused on LGBT rights, invited the CEOs of over 400 companies to co-sign an letter opposing the bill; over 100 did so (Wowak, Busenbark and Hambrick 2022).[3] The open letter explicitly argued that the bill would be detrimental to business by harming their ability to attract and retain talented employees to North Carolina, stating that "…HB2 will make it far more challenging for businesses across the state to recruit and retain the nation's best and brightest workers and attract the most talented students from across the nation" (Human Rights Campaign 2016).

The story of HB2 became one of the most widely-covered national news stories of 2016.[4] In a December 2016 public opinion poll by the Associated Press asking about ten significant news events of the year, two-thirds of Americans stated that it was at least "somewhat" important to them personally (see Figure 2.1).

The bill became a major issue in the 2016 gubernatorial race in the state, in which Attorney General Roy Cooper, a Democrat, ultimately defeated Republican incumbent Pat McCrory after a close election. After being sworn in, Cooper was able to negotiate a partial repeal of the bill in early 2017. Almost exactly a year after it was enacted, the bill was partially repealed on March 2017 (seemingly in response to a 48-hour ultimatum sent by the NCAA which threatened to relocate all college sport tournaments from the state, including March Madness). The partial repeal (HB142) removed the bathroom-related provisions of the law, but did not remove the state preemption of local authorities' ability to pass anti-discrimination ordinances. While the partial repeal largely removed the risk of losing Federal funding (and NCAA March Madness), it was criticized by both pro-LGBT activists and conservative defenders of the bill as an unsatisfactory compromise (Avery 2020), with organizations such as the American Civil Liberties Union referring to HB1424 as a "fake repeal" (Rho 2017). The local preemption provisions of HB142 were not fully repealed until HB2 was allowed to expire in December 2020, at which point multiple local authorities in North Carolina immediately enacted their own local anti-discrimination provisions akin to the initial measure passed by the Charlotte city council.

Given the extent to which HB2 became a national political issue, there was large-scale debate over the extent to which the bill had a negative economic impact on the state. One early analysis, published by the UCLA School of Law's Williams Institute, detailed a range of potential costs

[3]Among the many notable signatories were Lloyd Blankfein (Goldman Sachs), Marc Benioff (Salesforce), Brian Chesky (Airbnb), Tim Cook (Apple), Paul Graham (YCombinator), Travis Kalanik (Uber), Marissa Mayer (Yahoo), Doug Parker (American Airlines), Dan Schuman (Paypal), Howard Schultz (Starbucks), Harvey Weinstein (The Weinstein Company), and Mark Zuckerberg (Facebook).

[4]A Nexis Uni (formerly LexisNexis) search for "North Carolina bathroom" between March 1, 2016 and April 1, 2017 returned over 10,000 news results, including 2,229 results from CNN television transcripts, 2,122 results from the *New York Times*, and 1,787 results from the *Associated Press*.

Figure 2.1: Public opinion of the importance of the most important news stories of 2016

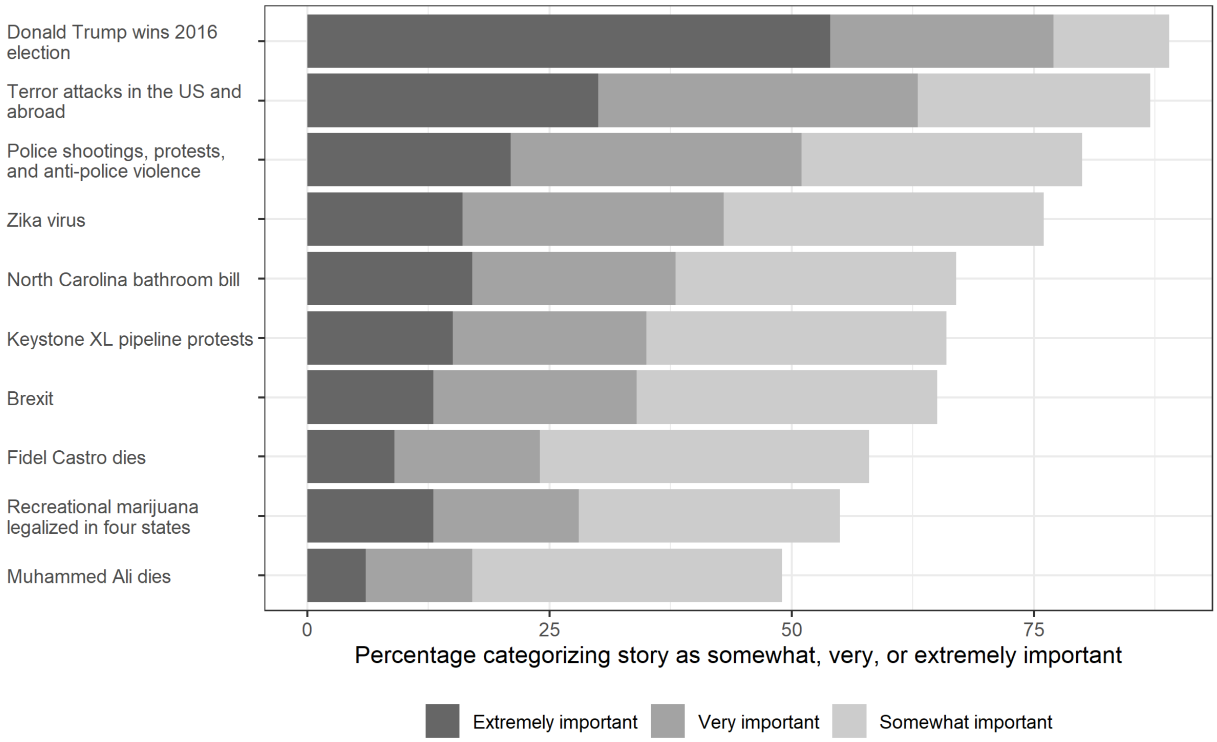

In December 2016, an Associated Press poll provided respondents with ten news stories from the year and asked respondents ”For each [story], please indicate how important that story was to you personally.” 67% of respondents indicated that the North Carolina bathroom bill story was at least “somewhat” important to them personally, comparable to other major stories from the year such as the Brexit referendum and protests against a natural gas pipeline in South Dakota. The exact text used to describe the story in the questionnaire was “North Carolina enacts law curbing LGBT (lesbian/gay/bisexual transgender) rights and transgender bathroom access, provoking economic and political backlash” (Associated Press 2016).

from the bill, including the potential loss of grant money from the Federal government to support North Carolina state-administered programs in education, healthcare, job training and workforce development, housing, and criminal justice (Mallory and Sears 2016). The Williams Institute analysis totalled these potential losses at nearly $5 billion, but nearly all of this was from the potential loss of Federal grant dollars rather than from a negative impact to state businesses from lost employment. A later, widely-publicized analysis by the Associated Press (AP), based upon companies’ publicly-released statements from companies, found that these companies claimed that North Carolina lost over 3,000 jobs as a result of the bill (Dalesio and Drew 2017). The analysis also suggested that this figure was likely an undercount of the true loss in employment, as it only totalled potential job losses from company press releases and public records that explicitly tied lost jobs were attributable to HB2.

Yet while the AP suggested that the losses are far larger, there are several reasons to be skeptical that this analysis should be the final word on the employment impact of HB2. First, the employment impact of proposed projects such as PayPal's proposed 400-job facility in Charlotte come from economic forecasts generated by the North Carolina Commerce Department for the purpose of granting tax breaks for job creation in the state. Researchers who have conducted causally-identified research unto the efficacy of these tax breaks have found that these economic impact assessments often rely upon unrealistically optimistic assumptions and dramatically overstate the fiscal and employment impacts of development projects (Jensen 2017*a*,*b*; Jensen and Malesky 2018; Jensen and Thrall 2021). Second, this analysis simply summed self-reported job cancellations without any attempt to define or construct a counterfactual in which the bill was not passed. While this reporting represents an accurate reporting of journalistic facts, social scientists using modern statistical techniques for causal identification can (and should) more carefully construct an estimate of the true impact of HB2. In the next section, I describe the data which I use to perform such an analysis.

2.4 Data

The data used in this analysis come from the Longitudinal Employer-Household Dynamics (LEHD) Program from the US Census Bureau, which compiles statistics related to employment and job mobility. Specifically, I make use of the LEHD Job-to-Job (J2J) flows data, which report quarterly statistics related to employment transitions. These J2J data are based upon state unemployment insurance (UI) filings, which are then linked with establishment-level data collected as part of the Quarterly Census of Employment and Wages (QCEW). Because all employers must submit this information to their respective states' UI agencies, coverage is nearly universal; the Census Bureau estimates that LEHD data cover 95% of private-sector employment, as well as employees in state and local government (Hyatt et al. 2017). While the data are publicly available from the US Census Bureau, they are spread across a large number of individual state files which requires significant manipulation by users to prepare the data for analysis.

The J2J data capture the number of workers moving from job to job. Relative to other commonly-used data on interstate migration, the LEHD J2J are unique in that they focus specifically upon workers, rather than the population as a whole. Other Census Bureau products, such as the American Community Survey (ACS) are commonly used in studies of migration (e.g., Downey and Liu 2023). The ACS is based upon a nationally-representative 1% sample of the U.S. population. Because it is not limited to those in employment, it covers a broader swath of the population relative to the LEHD. On the other hand, because it relies upon a small probability sample, researchers using the ACS must take into account potentially large sampling errors. This is especially concerning when looking at relatively small cell counts—such as moves involving uncommon origin-destination pairs (e.g., the population of workers moving from Vermont to Wyoming in any given period is likely to be quite small and so a 1% sample may measure this population with large error). While researchers have deployed statistical techniques to attempt to quantify and account for this potential sampling error (see Downey and Liu 2023), my reliance on the LEHD J2J obviates the need to rely upon these models to quantify sampling error, as the J2J product includes the nearly the entire employed population. The LEHD J2J figures also have the advantage of more frequent periodicity (quarterly) relative to the annual coverage of the ACS. This is also an advantage of the LEHD J2J over the IRS Statements of Income (SOI), which provide annual figures on the number of tax filers who relocate from state-to-state (see Olney and Thompson 2024). The SOI's focus on tax returns also means that the IRS cannot precisely measure the number of individuals moving as they do not necessarily know the exact number of household members (dependents) associated with a given tax filer, nor does it cover individuals who do not file tax returns (e.g., those with very low incomes). The SOI also does not include individual demographic information (including employment status) of filers and their dependents.

Because the J2J data capture employment transitions, they focus on a subset of workers in fairly regular employment rather than workers (re-)entering the labor force from longer spells outside of work. For individuals who may receive positive earnings from multiple employers, the J2J data focus on primary (dominant) jobs, which are defined as the job with the greatest earnings in a given quarter (Hyatt et al. 2017).[5] While the broader LEHD does include data on workers who are less attached to the labor market, the J2J flows only capture workers whose new job is separated by the prior job by at most one full quarter of labor market inactivity. Within this set of transitioning workers, the J2J data distinguish between transitions of slightly longer and shorter durations. Employer-to-employer (EE) flows measure employment transitions which occur entirely within a calendar quarter and so do not include any measured period of labor market activity. Adjacent-quarter flows (AQHires) include a single quarter of labor market inactivity, and so may include a transition period of between three and six months, depending upon calendar timing. Because interstate moves may reasonably involve some period of labor market inactivity as workers move, I sum EE and AQHires into a combined measure of all migration.

Worker mobility patterns in the J2J can be further decomposed by individual worker characteristics such as age, sex, and education (captured from survey, Census, and other administrative data sources), as well as employer characteristics (in both the origin and destination) such as the location of the employer and its industrial sector (NAICS 2-digit code), firm age, and firm size (captured in the QCEW). Finally, the data also contain average earnings in both the origin and destination. For our purposes, these data can be used to track the quarterly flows of migrants to (and from) North Carolina in aggregate, as well as highly specific types of workers (e.g., differences by educational attainment, age, sex,[6] race and ethnicity, industry of employment, and state or metropolitan area of of origin).

The LEHD J2J data product begins coverage in 2000 Q2, and attained full national coverage in 2010 Q2 when the last state (Massachusetts) began sharing data. My analyses cover the period 2011 to 2020, which avoids periods of significantly reduced job transitions during the Great Financial Crisis of 2008–2009 and its immediate recovery as well as the onset of COVID-19 lockdown

[5] Transitions involving short-duration jobs that last less than a full quarter or secondary jobs are thus not captured. For the purposes of this paper—focused on interstate moves—this is unlikely to be an issue. While there may be workers who hold multiple simultaneous jobs, it is reasonable to assume that most workers are unlikely to simultaneously hold jobs across multiple states.

[6] The LEHD product measures sex as either male and female.

in early 2020. During the 2011–2020 period, three states (Alaska, Arkansas, and Mississippi) suspended sharing data with the LEHD. Because these are relatively small states, I drop them from the analysis to avoid an unbalanced panel. In some additional analyses and robustness checks, I drop certain states (and DC) with very volatile or otherwise unusual migration patterns, such as DC (which has very high migration rates due to high transience and a small population) as well as Wyoming and North Dakota, which have volatile migration rates due to the large role of oil in the states' job market.

My main measure of interstate migration focuses on worker migration rates rather than absolute number of migrants to normalize for differences in state population size. I calculate this measure analogously to the manner in which the Census Bureau calculates quarterly hiring rates within a state, but subset only to only interstate job transitions. This measure divides the total number of migrating workers to or from a given state in a given quarter by the average number of unique workers employed in the state during the relevant quarter. Data on the total number of workers (not just movers) come from the Quarterly Workforce Indicators (QWI), another package available as part of the LEHD program. For inbound migration, I define the inbound interstate migration rate as

$$\text{inbound}_{dtp} = \frac{\sum_{o \neq d} \left(\text{EE}_{dtp,o\neq d} + \text{AQHire}_{dtp,o\neq d} \right)}{\left(\text{Emp}_{dtp} + \text{EmpEnd}_{dtp} \right) \times 0.5} \times 100, \tag{2.1}$$

for destination state d, origin state o, time (year-quarter) t and optional sub-population p (i.e., specific subgroups by educational attainment, age, or industry of employment). *EE* are employer-to-employer flows and $AQHire$ are adjacent-quarter flows from the J2JOD files and *Emp* and *EmpEnd* are the beginning- and end-of-period employment totals from the QWI.[7]

In other words, I take the sum total of all employer-employer and adjacent-quarter hires moving to a state in a given quarter that do not originate in the same state, and then calculate the percentage of the average employment in that quarter. When analyzing specific specific subpopulations, such as college-educated workers, I take the inbound migration as a percentage of the average employment in the destination state for this subpopulation.

In some additional exhibits and analyses in the appendix, I focus instead on raw (or logged) number of movers to better illustrate the scale of workers involved. Additionally, the LEHD figures are not seasonally adjusted and exhibit highly seasonal patterns as moves are most common during the summer months (Q3), and so I seasonally adjust the figures in some supplemntary analyses.[8]

[7]Out-migration is defined analogously for origin state o as

$$\text{outbound}_{otp} = \frac{\sum_{o \neq d} \left(\text{EE}_{otp,o\neq d} + \text{AQHire}_{otp,o\neq d} \right)}{\left(\text{Emp}_{otp} + \text{EmpEnd}_{otp} \right) \times 0.5} \times 100. \tag{2.2}$$

[8]For ease of interpretation in some exhibits, I seasonally adjust the time series of *EE* and $AQHire$ flows for state

2.5 Analysis of Migration Rates

Descriptive Data

Before conducting any statistical analyses, I begin by plotting the raw data. In the years prior to the enactment of HB2, North Carolina—like much of the "Sun Belt"—had fairly high net migration rates. Figure 2.2 is a map depicting the quarterly average worker in-migration, out-migration, and net migration rates for all states during the four years 2012–2015 (i.e., before HB2). During this period, states throughout the Northeast and much of the Midwest experienced net worker outflows, while states in the West and Sunbelt experienced net worker inflows, consistent with long-run population dynamics (Olney and Thompson 2024).

Figure 2.3 depicts also the in-migration and out-migration of workers to and from North Carolina compared to other states both before and after the passage of HB2.[9] Several features of Figure 2.3 are worth noting. First, there is strong seasonality in quarterly interstate migration figures, as workers more typically move during the summer months (Q3). Second, North Carolina appears roughly average among all states in terms of in-migration rates, somewhat below average in terms of out-migration, and somewhat above average in terms of overall net migration. Third, and more relevant for the topic at hand, there do not appear to be major breaks in North Carolina migration trends following the enactment of HB2 in March 2016. While I proceed to conduct a variety of statistical tests to show the (lack of an) effect of HB2 on migration, these raw figures should constitute initial evidence that there was not such a large trend that it is easily visible to the naked eye.

Estimation Strategy

HB2 was enacted and signed into law in late March 2016. Given that LEHD data are reported quarterly, I treat 2016 Q2 (beginning April 1, 2016) as first quarter in which North Carolina was "treated" by the effect of the law, and then analyze subsequent changes in North Carolina migration patterns relative to other geographies. Given the lack of variation in treatment timing (only one state "unit" is treated and only once), I attempt to test for a causal effect using the canonical 2×2 difference-in-difference setup, which can be estimated using a two-way fixed-effects model

$$y_{st} = \alpha_s + \gamma_t + \beta_1 \mathrm{NC} + \beta_2 \mathrm{POST} + \beta 3(\mathrm{NC} \times \mathrm{POST}) + \epsilon_{st}, \tag{2.3}$$

where y is a measure of migration flows, s indexes geographies (i.e., states), t indexes time periods (quarters), *NC* is a state indicator for North Carolina, and *POST* is an indicator of whether the time period is after 2016 Q1, when the law was enacted. α_s and γ_t are state- and time-fixed

using the `X13ARIMA-SEATS` adjustment program developed by the US Census Bureau for this purpose (US Census Bureau 2017) and then calculate separate seasonally-adjusted per-capita figures, which I present in some supplemental analyses for greater graphical clarity. To do this, I make use of the `seasonal` package, which provides an R interface to `X13ARIMA-SEATS` (Sax and Eddelbuettel 2018, 2022).

[9] In addition to Figure 2.3 shown here, Appendix **??** displays the data state-by-state and in terms of absolute number of migrants instead of rates. It also includes seasonally-adjusted figures.

Figure 2.2: Average quarterly migration rates by state, 2012–2015

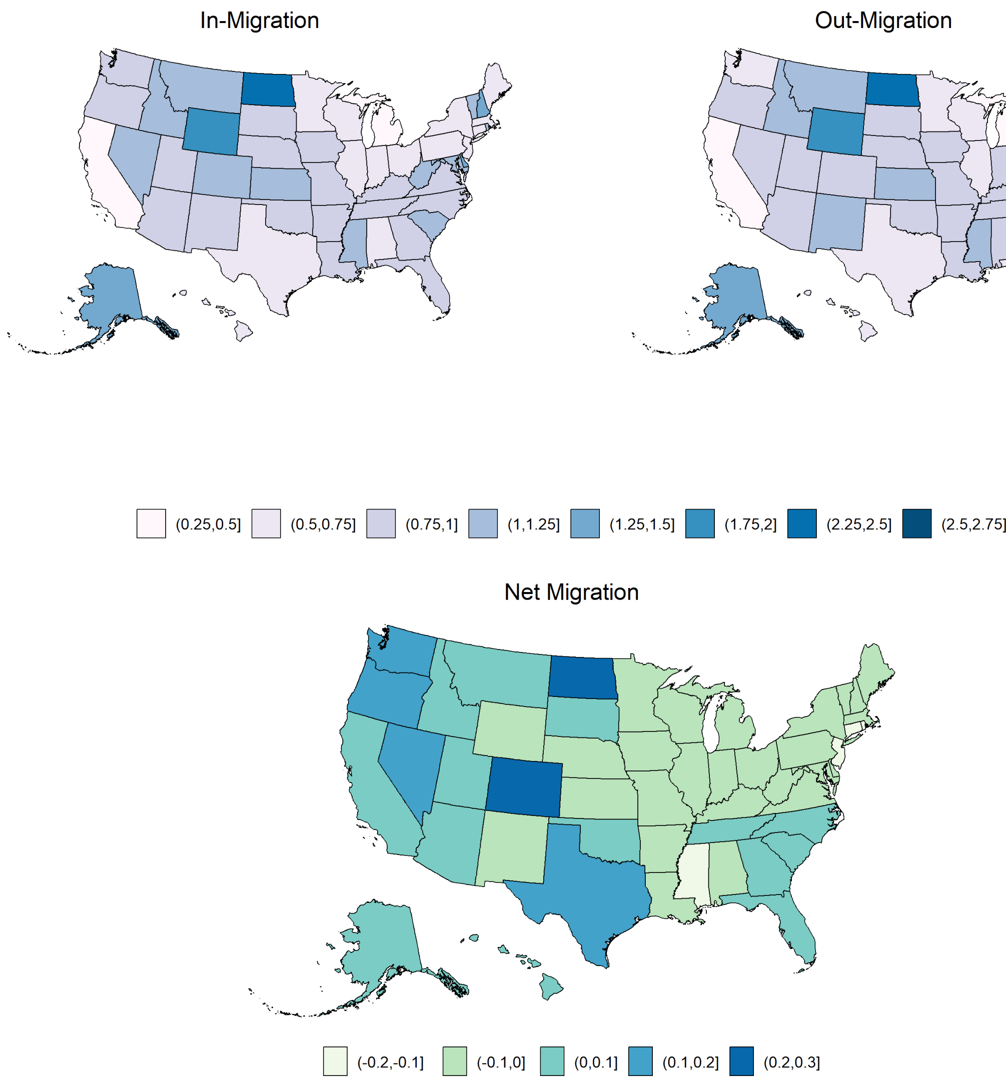

Figure includes average quarterly migration rates from 2012 to 2015 (number of workers migrating per quarter divided by average total workers in the state during the quarter). Note different scales for in- and out-migration (top row) and net migration (bottom).

Figure 2.3: Quarterly interstate worker migration rates, 2011–2019, NC vs. other states

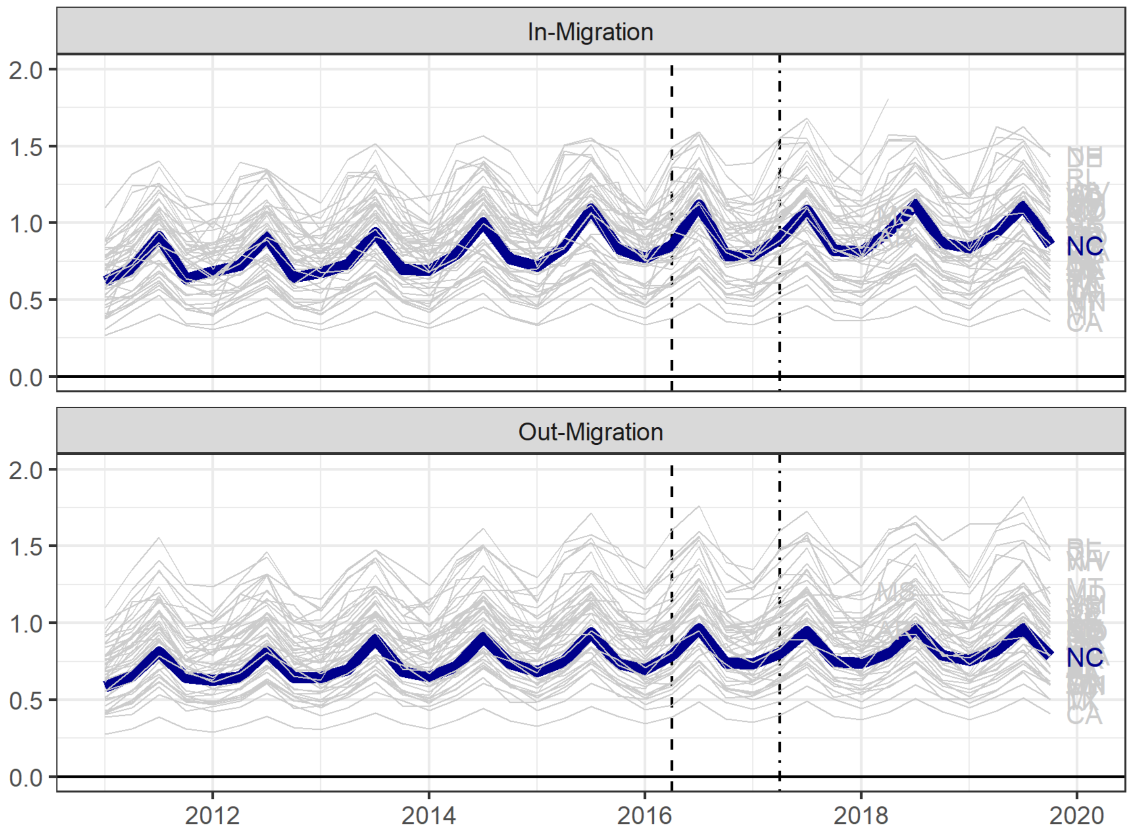

Migration rates are defined as the number of workers entering or leaving the state over the average number of workers employed during the quarter. Includes both employer-employer (EE) and adjacent-quarter (AQHires) from LEHD J2J data. Quarterly rates are defined by number of migrants over the workers total number of individual workers employed in the quarter, as reported in LEHD QWI data. First vertical dashed line is 2016 Q2, when HB2 first entered into effect, and second dot-dash line is when it was partially repealed in 2017 Q2 (HB142). Figure excludes small, natural-resource states with highly volatile migration flows (AK, ND, WY) as well as DC (which has very high worker migration rates owing to a high degree of transience and low overall population). AR and MS data series terminate in 2018.

effects, respectively. The two sets of fixed effects account for time-invariant characteristics of specific destination states (such as proximity to other highly-populated states and geographic amenities such as weather and natural beauty) as well as time-varying factors which affect all states equally (e.g., the overall macro-economy).

The DiD setup requires an assumption that in the absence of treatment, the differences in potential outcome between treated and control units would remain the same. While this counterfactual outcome is unobservable, researchers typically test to see if treated and control units moved in parallel prior to the treatment (Egami and Yamauchi 2023; Xu 2017). A common test is to regress the outcome on a treatment group indicator interacted with multiple pre-treatment "lead" time indicators. This is a necessary but not sufficient test; passing this test does not guarantee that the (post-treatment) parallel trend assumption holds, but rejecting the null of zero pre-treatment effects is a likely sign that the parallel assumption trend does not hold (Kahn-Lang and Lang 2020).

Figure 2.4: Event study of "leads" prior to HB2 in North Carolina

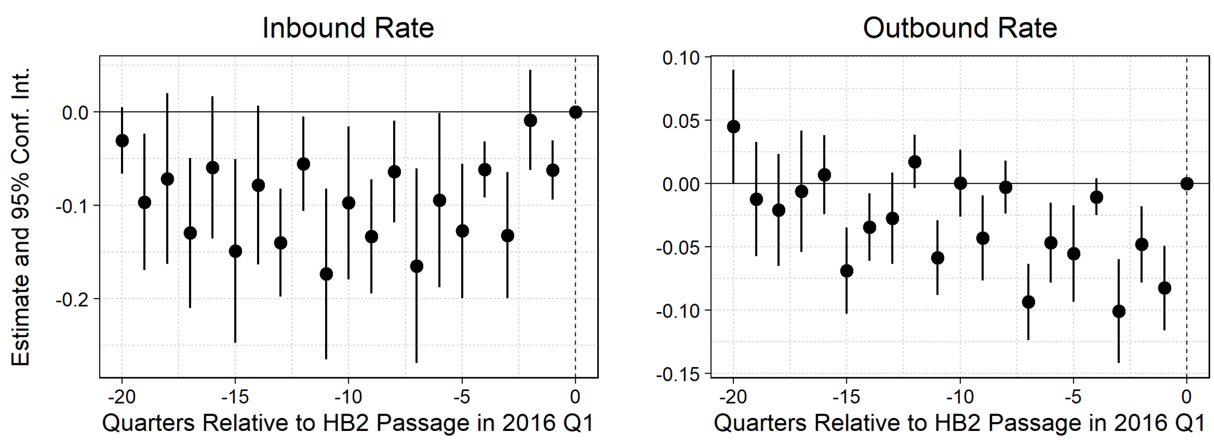

The plot of "leads" suggests that the assumption of parallel trends is unlikely to hold in this context. Models include state and year-quarter fixed effects.

In this case, the assumption of parallel pre-trends appears not to hold, as seen in Figure 2.4. In a series of unreported additional analyses, I also test for to see if the event study test of parallel trends holds for other model specifications and various transformations of the dependent variable, as the existence of parallel trends is not necessarily robust to functional form transformations of the dependent variable (Kahn-Lang and Lang 2020; Roth and Sant'Anna 2023). These include seasonally-adjusting the data and also using the absolute or logged level of interstate migrants instead of the rate. I also origin-destination pairs as the unit of analysis (instead of aggregating across origins) and consider shortened pre-treatment periods. Finally, I also test a triple-difference (DDD) specification in which I test for the existence of parallel trends between the difference in within-state migration rates between North Carolina and other states and inter-state migration rates to North Carolina versus all other states. I consistently reject the null that trends are parallel prior to treatment.

While it is possible that further manipulation of the model (i.e., inclusion of control variables) could help to make the parallel trends assumption more plausible, I proceed with a different modeling technique that is specifically designed to deliver robust causal inference in settings where the parallel trend assumption appears implausible. Specifically, I make use of the generalized synthetic control (GSC) method developed by Xu (2017). Unlike DiD, synthetic control methods do not require researchers to invoke a (unobserveable) parallel trend assumption. The synthetic control method relies upon pre-treatment outcomes (and optionally, pre-treatment covariates) among a set of control units (the “donor pool”), which are then re-weighted to generate a “synthetic control” that serves as the counterfactual for the unit under treatment (Abadie, Diamond and Hainmueller 2010, 2015). This method is thus an extension of earlier matching approaches for causal inference (e.g., Abadie 2005), but unlike matching approaches, the weighting method of the generalized synthetic control is both more accurate and more transparent and avoids the potential for specification searches (Xu 2017). The GSC method (Xu 2017) builds upon the interactive fixed effects approach (Bai 2009), which derives unit-specific intercepts and time-varying coefficients. This approach builds upon the quantitative finance literature on factor models, familiar to scholars who have conducted asset price event studies using familiar factor models of asset pricing (e.g., the Fama-French three-factor model).

Overall migration

In-Migration I first examine overall migration to North Carolina. As the GSC method requires a fairly large number of pre-treatment periods, the analysis begins in 2011 Q1 (20 quarters prior to treatment) and extends through the end of 2019 (15 quarters post-treatment). The generalized synthetic control method uses a weighted average of control units to generate a “synthetic” treated unit that estimates what the outcome for the treated units would have been under the counterfactual of no treatment (Abadie, Diamond and Hainmueller 2010). For in-migration to North Carolina, the implied weights are presented in Figure 2.5.

The largest positive weight is assigned to DC and the largest negative weights[10] are assigned to North Dakota and Wyoming. These states (and DC) have several unusual features. All are small in population, which results in high per capita migration rates. DC, given its small population and highly transient workforce has very high per capita migration rates; North Dakota and Wyoming also have unusually volatile migration patterns owing to their small populations and economic dependence upon commodities. In appendix Figure B.6, I show that the results I present here are robust to the exclusion of these two states and DC.

[10] Although there have been recent concerns about the use of implied negative weights, these concerns about negative weights arise when using traditional two-way fixed effect (TWFE) estimators of differences-in-differences in the presence of either heterogeneous treatment effects and/or staggered rollout designs in which not all units are treated at the same time (de Chaisemartin and D’Haultfœuille 2023; Roth et al. 2023). This concern about negative weights does not apply in the GSC design (Xu 2017).

Apart from these outliers, the weights appear consistent with broad intuition about states which might seem comparable to North Carolina. The states which are assigned the highest weight are all "Sun Belt" states within the same broad region of the country as North Carolina. The remaining states with the largest weights are Tennessee, Alabama, South Carolina, Louisiana, and Georgia, all of which are located alongside North Carolina in the South Census Region. The (remaining) states with the lowest (negative) weights are states located in the Northeast and/or upper Midwest (Montana, Maine, Vermont, New Hampshire, and South Dakota) that are geographically, demographically, and economically quite distinct from North Carolina. While the weights are derived through a cross-validation algorithm rather than set by researchers on the basis of some theoretical similarity, we can take comfort from the fact that the assigned weights appear consistent with our general intuition.

Figure 2.5: Implied weights for estimated effect on in-migration to North Carolina

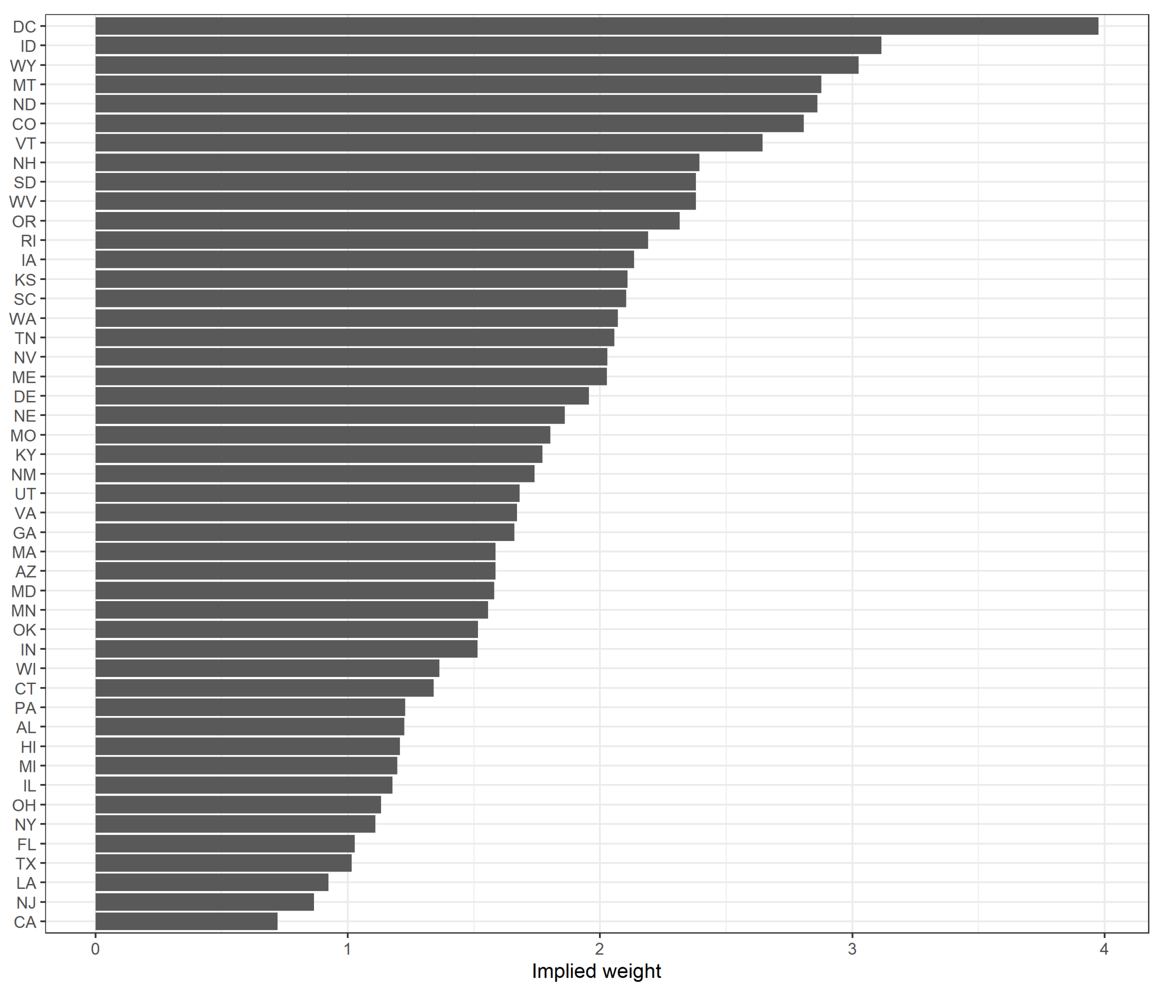

Implied weights for generalized synthetic model fit in Figure 2.6.

Figure 2.6: Estimated effect of HB2 on in-migration to North Carolina

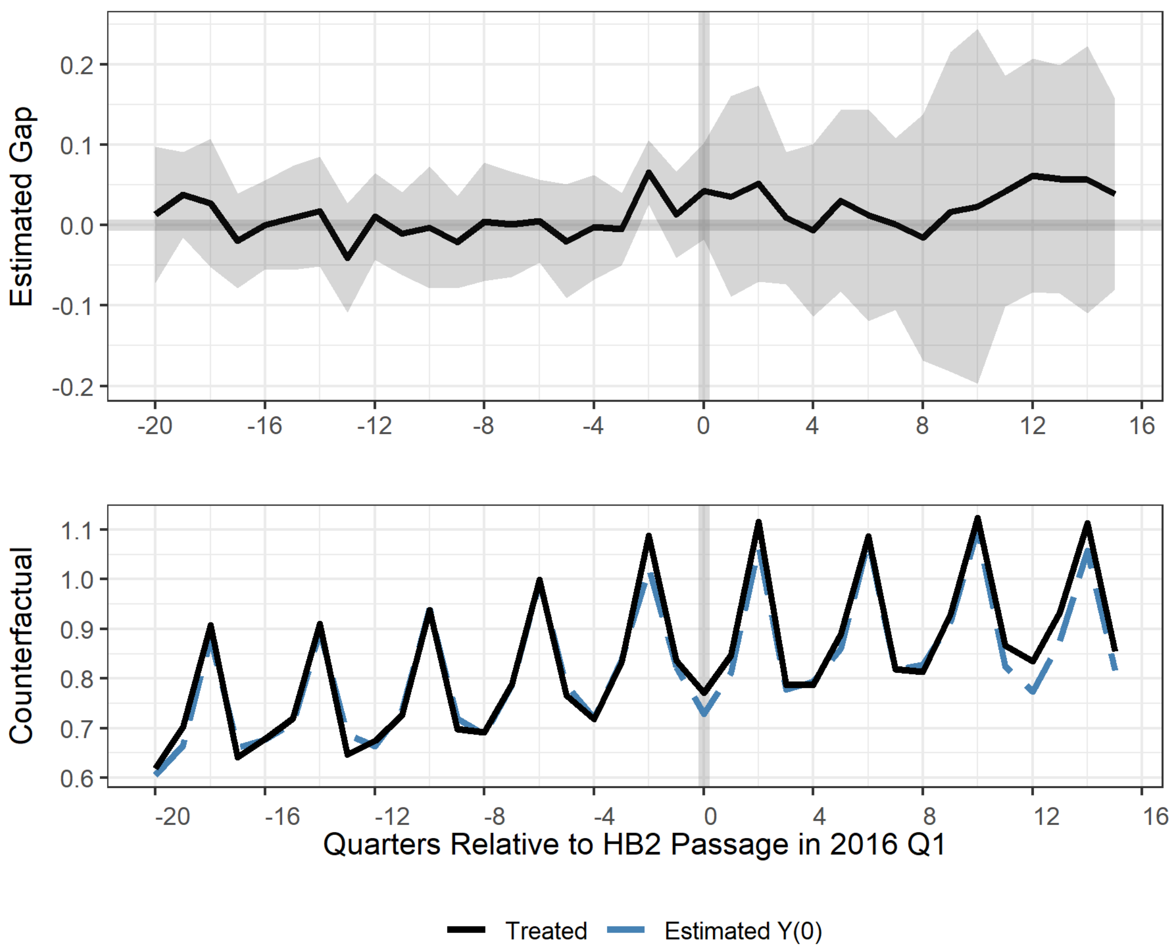

Top panel is a "gap" plot and bottom panel depicts the actual observed values for North Carolina versus the generalized synthetic control estimate of North Carolina under control. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure 2.6 contains the results of the GSC model. The top panel is a "gap" plot which displays the difference between the actual observed values and the our estimated counterfactual "synthetic" North Carolina (i.e., what we would have expected employment to be in the absence of HB2), while the bottom panel depicts the actual observations versus the counterfactual synthetic North Carolina. We see no evidence of a negative impact on in-migration in either the four quarters immediately following the enactment of HB2 but before it was partially repealed, or in the full analysis period through the end of 2019.

Figure 2.7: Estimated cumulative average treatment effect (CATT) of HB2 on in-migration to North Carolina

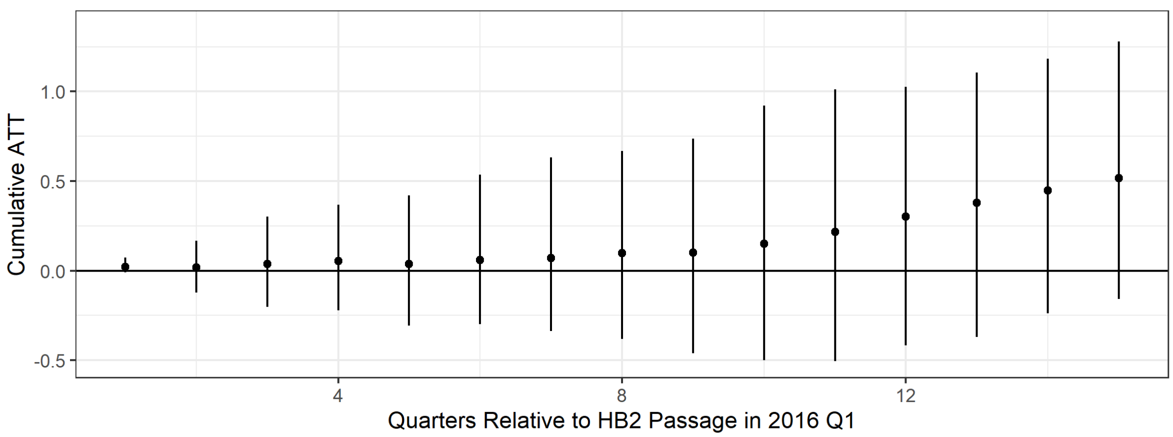

Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure 2.7 depicts the cumulative average treatment effect (CATT) estimates of the effect of HB2 for each quarter post-treatment. After four quarters of treatment and prior to the partial repeal of HB2 (i.e., 2016 Q1–2017 Q1), there appears to be no effect on in-migration to North Carolina. After four quarters of treatment, the cumulative average treatment effect for the treated (cumulative ATT) is fairly precisely estimated around 0. A 95% confidence interval for the CATT after four quarters is $(-0.19, 0.26)$, meaning we can largely rule out any negative effect larger than a 0.2% change in per capita migration to North Carolina overall. For reference, the total working-age employment in North Carolina during this period was about 4.15 million, meaning that we can likely rule out that the law led to a decline in migration larger than 8,000 workers during the year that it was in full effect.

Out-Migration Another possible impact of HB2 on North Carolina employers' ability to access talent might occur if workers increase the rate at which they leave the state. I thus repeat the same exercise for migration leaving North Carolina. Figure 2.8 depicts the cumulative average treatment effect on out-migration from North Carolina; I find no significant increase in out-migration from North Carolina in the four quarters following the passage of HB2 or in the period after the partial repeal. Figure B.8 depicts the weights uses for this calculation and Figure B.9 depicts the "gap" and "counterfactual" plots; these results are robust to excluding outlier states.

Figure 2.8: Estimated cumulative average treatment effect (CATT) of HB2 on out-migration from North Carolina

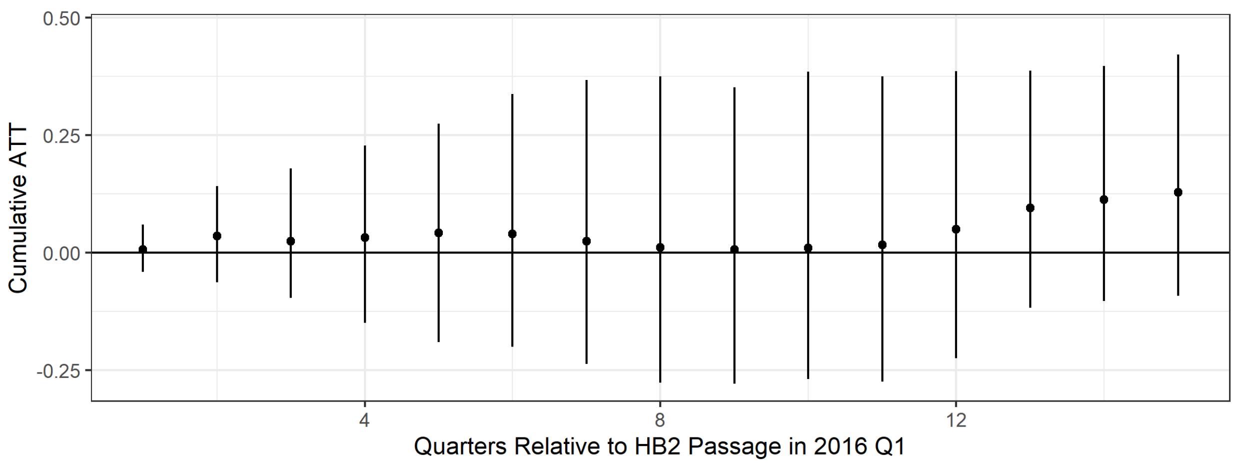

Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Subpopulations

Education The effects of policies such as HB2 may be heterogeneous by educational category. Better-educated workers may have more labor market power and income, and thus more geographically mobile. They may also hold more liberal attitudes (Downey and Liu 2023; Zingher 2022) and so be more responsive to conservative policies. Finally, better-educated workers may be of particular concern for employers. The open letter (Human Rights Campaign 2016) specifically mentioned the "best and brightest workers" and "most talented students."

The LEHD J2J data disaggregate workers by educational category. The categories are workers with less than a high school degree (category E1), a high school degree or equivalent (e.g., GED) (E2), some college or a two-year associate's degree (E3), and a four-year degree or higher (E4). Finally, workers under 24 who may still be completing their education are modelled separately (E5).

Figure 2.9 displays the "gap" plot for in-migration rates, with separate GSC models fit for each educational category. To account for differences in educational attainment across states, the denominator for the quarterly migration rate is the average number of workers of that specific educational category during the quarter. As was the case for overall migration, there is no detectable decrease in in-migration to North Carolina for any educational cohort. This is true for those with four-year college degrees and younger respondents (those under 24). In fact, we observe an *increase* in migration to North Carolina relative to the counterfactual for those with lower education, although takes place significantly after the enactment of the law in 2018–19 and may be attributed to other factors.

Figure 2.9: Estimated effect of HB2 on in-migration to North Carolina, by education category

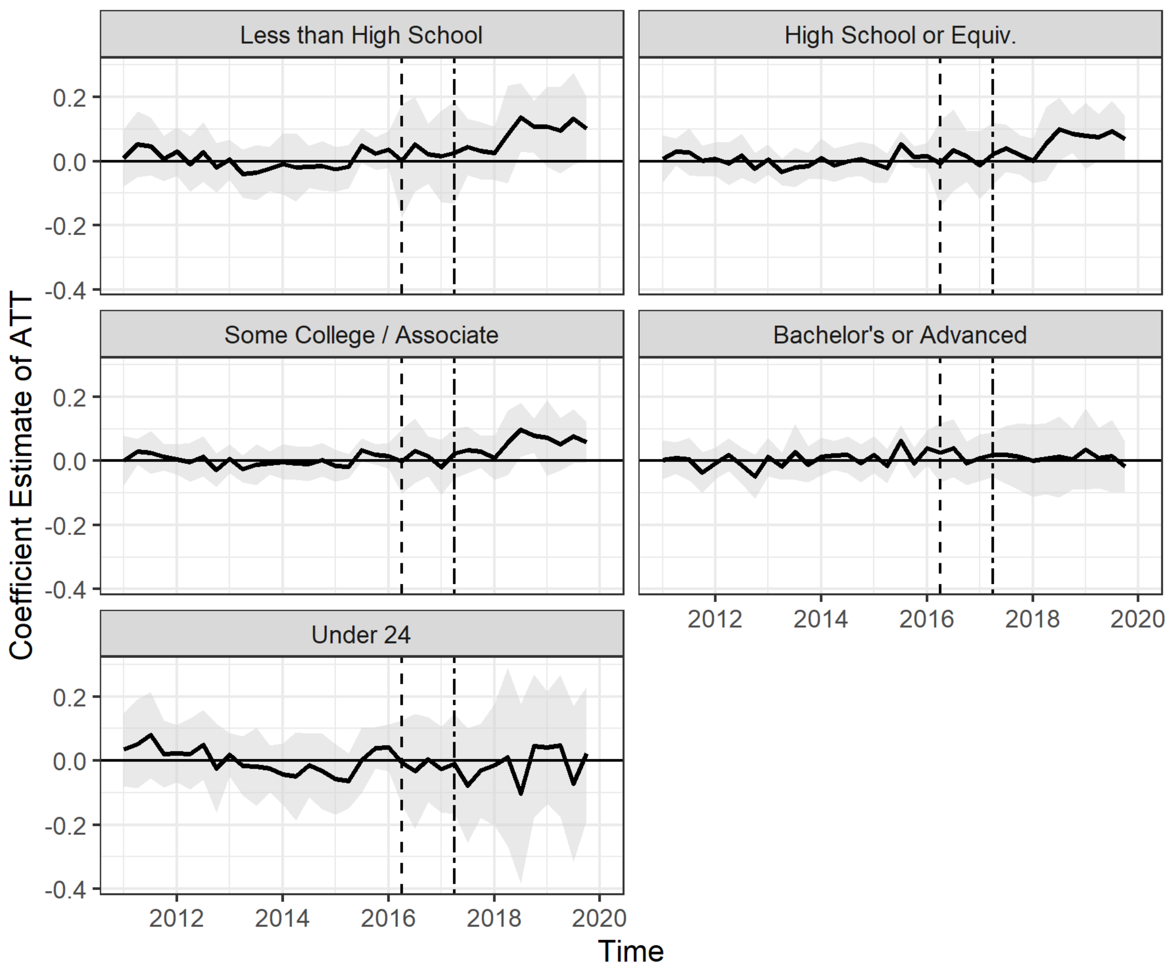

"Gap" plots display the difference between observed migration and the counterfactual "synthetic" North Carolina, separately for each educational category. The rate is defined relative to the number of workers of that specific educational category to normalize for differences across states. The first vertical dashed lined is at 2016 Q2, the first quarter where HB2 was in effect. The second dot-dash line is at 2017 Q1, when the law was partially (but not completely) repealed. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Age cohort Policies such as HB2 may also have heterogeneous effects by age. Younger people may be especially likely to hold pro-LGBT beliefs and may be more geographically mobile. The LEHD J2J data disaggregate workers into eight categories. As in prior sections, Figure 2.10 displays the "gap" plot for in-migration rates separately for each age category. Although some younger age categories may have low overall employment rates (i.e., 14–18 year olds), these rates are calculated as a percentage of the workers of that specific age category working in the state to normalize for these differences. Once again, I find no evidence of a change in migration rates to North Carolina; 95% confidence intervals for all age cohorts except (highly mobile) 19–21 year olds are extremely precisely estimated around zero.

Industry Finally, I look at the destination industry—i.e., the industry where interstate migrants are employed in their destination state. The results are shown in Figure 2.11. Agriculture (11), extractive industries (21), and construction (23) exhibit higher volatility. In industries that might be expected to suffer from conservative policies such as HB2, including high-skill industries such as IT (51), finance and insurance (52), and professional services (54), as well as in industries with a high proportion of liberals such as education (61) and arts and entertainment (71) (see Chapter 1, this volume), we still find no negative effects.

2.6 Alternative Explanations

In the previous section, I found no evidence that HB2 resulted in decreased in-migration to North Carolina. In this section, I consider—and reject—two potential explanations for why this might be.

Shifts in Partisan Composition of Migration

First, it is possible that the overall number of migrants remained constant, but that different groups responded differently and offset one another. Specifically, it is possible that one group responded negatively by becoming less likely to move to North Carolina (and more likely to leave the state), while a roughly equivalent number of some other group responded positively by becoming more likely to move to North Carolina (and less likely to leave the state).

In this context, I use the partisanship of the origin states and destination states for migrants to and from North Carolina, respectively. Partisanship is an effective proxy for attitudes towards transgender rights, with Republicans far more likely to be supportive of measures such as HB2 than Democrats. While this is a rough proxy and by no means fully determinative, significant shifts in the origins and/or destinations of migrants to and from North Carolina in the post-HB2 may indicative the potential for offsetting and polarized responses.

Figure 2.12 depicts the quarterly number of inbound and outbound migrants to and from North Carolina during the analysis period, grouped by the 2012 Presidential vote margin in the origin or destination state. Migration to North Carolina from solidly Republican states did increase after the passage of HB2, but appeared to have already been increasing prior to 2016, and

Figure 2.10: Estimated effect of HB2 on in-migration to North Carolina, by age category

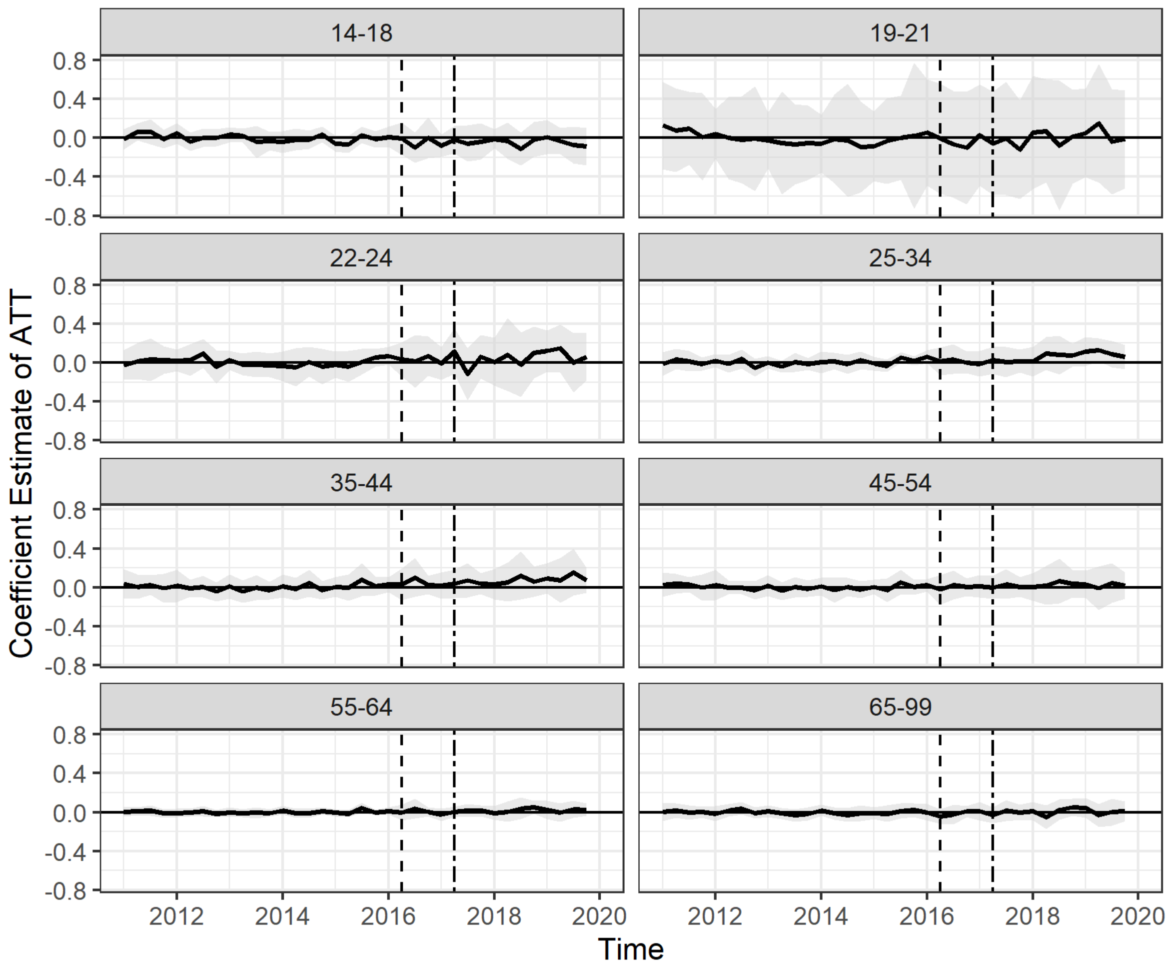

“Gap” plots display the difference between observed migration and the counterfactual “synthetic” North Carolina, separately for each age category. The rate is defined relative to the number of workers of that specific age category to normalize for differences across states. The first vertical dashed lined is at 2016 Q2, the first quarter where HB2 was in effect. The second dot-dash line is at 2017 Q1, when the law was partially (but not completely) repealed. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure 2.11: Estimated effect of HB2 on in-migration to North Carolina,, by destination industry

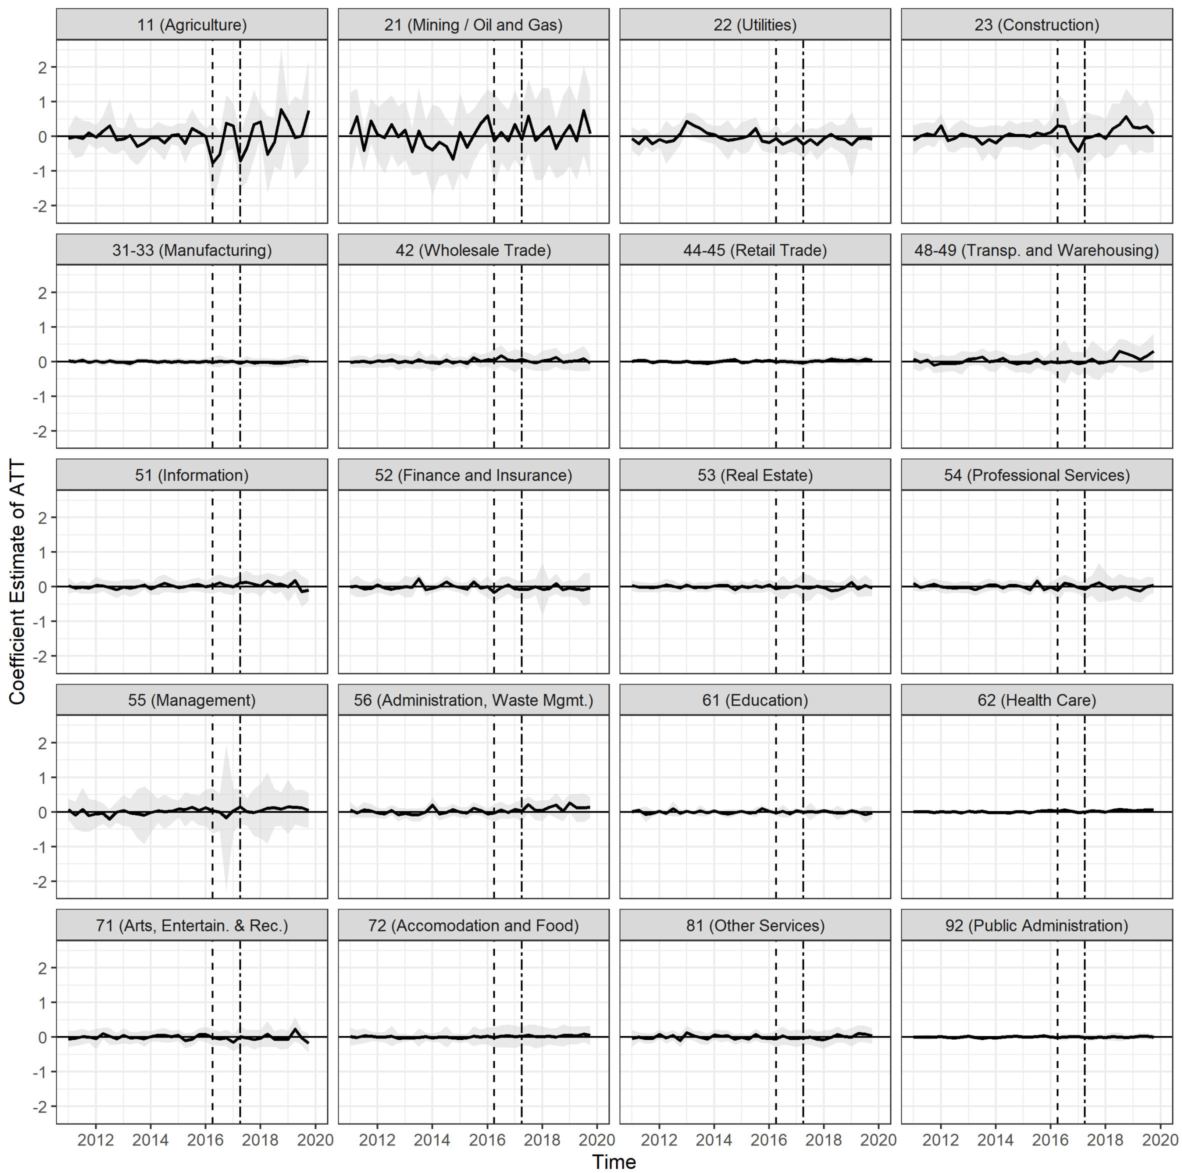

"Gap" plots display the difference between observed migration and the counterfactual "synthetic" North Carolina, separately for each industry. The rate is defined relative to the number of workers of that specific industry to normalize for differences across states. The first vertical dashed lined is at 2016 Q2, the first quarter where HB2 was in effect. The second dot-dash line is at 2017 Q1, when the law was partially (but not completely) repealed. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure 2.12: Origins and destinations of migrants to and from North Carolina by 2012 election returns

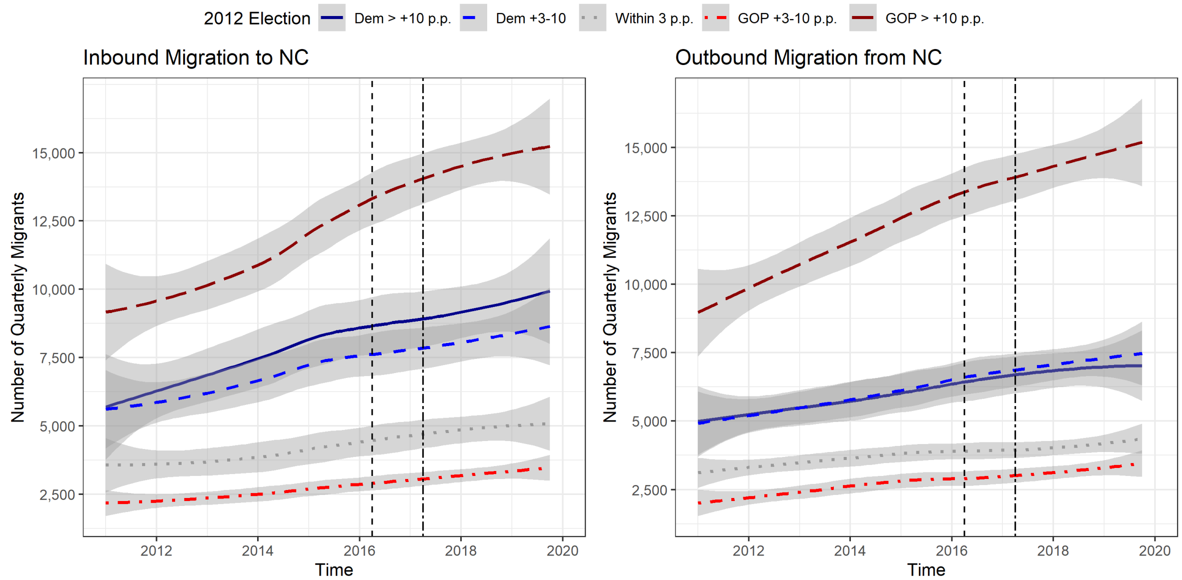

Left-hand-side plot depicts the quarterly number of migrants (smoothed using LOESS) to North Carolina based upon the two-party vote margin in the 2012 election. The right-hand-side plot depicst the quarterly number of migrants (smoothed using LOESS) from North Carolina based upon the two-party vote margin in the 2012 election. Solidly Democratic states (dark blue solid line) with a two-party Democratic vote margin of greater than 10 p.p. were CA, CT, DE, DC, HI, IL, ME, MD, MA, NJ, NM, NY, OR, RI, VT, and WA. Leaning Democratic states (blue dash line) with a Democratic two-party vote margin of between 3 and 10 p.p. were CO, IA, MI, MN, NV, NH, PA, VA¡ WI. Swing states (grey dotted line) were decided by less than 3 p.p.; these were FL, NC (not included in plot), and OH. Leaning Republican states (red dot-dash line) had a Republican two-party vote margin of between 3 and 10 p.p. and were AZ, GA, MO. Finally, solidly Republican states (dark red dash line) had a Republican two-party margin of over 10 p.p. and were AL, AK, AR, ID, IN, KS, KY, LA, MS, MT, NE, ND, OK, SC, SD, TN, TX, UT, WV, and WY.

the rate of increase actually moderated after the passage of HB2. Migration from Democratic states continued to increase after the passage of 2016 without any significant appreciable decline.

States of origin for in-migration may not represent a perfect proxy for partisanship. Migrants leaving heavily Republican (or Democratic) states may be likely to represent the party that predominates in their origin state, but it is also possible that the people leaving a state come from the minority party in the state.

Given this, the right-hand side of Figure 2.12 may be more compelling. It depicts the destination states for migrants leaving North Carolina. If migrants are leaving North Carolina in response to HB2, it seems unlikely that they would move to a Republican state where support for conservative policies is likely stronger. Once again, we find no evidence that HB2 had an effect on the destinations of choice for those leaving North Carolina: there is no noticeable increase in migrants leaving North Carolina to Democratic states or decrease in migrants leaving North

Carolina in favor of more Republican states.

Compensating Salary Differentials

Another potential explanation for this null finding may be employers' responses to the law—specifically, employers in North Carolina may have increased wages in an attempt to offer a compensating differential for the reduced attractiveness of jobs in the state. Prior work has found that employers in Charlottesville, Virginia, advertised higher wages after an infamous white supremacist rally in the town (Hurst 2023). To test whether this was the case after HB2, I take advantage of salary data also present in the LEHD J2J. For each origin-destination pair and defined sub-population, the J2J report the average quarterly salary in the destination and origin. These data can be used to measure the extent to which the premium or wage differential for workers moving to North Carolina may have grown relative to the premium paid to workers relocating to other states.

I can thus test for whether North Carolina employers responded to HB2 by increasing wages by using a difference-in-difference specification:

$$\Delta EESEarn_{odtp} = \alpha_{od} + \gamma_t + \beta_1 \text{NC} + \beta_2 \text{POST} + \beta_3 (\text{NC} + \text{POST}) + \varepsilon_{odpt}, \tag{2.4}$$

where $\Delta EESEarn$ is the change in quarterly average wages for employer-employer movers who moved from origin geography o to destination d at time t within subpopulation p (i.e., subpopulations partitioned by educational attainment). α_{od} is an origin-destination fixed effect that absorbs time-invariant characteristics of each state origin-destination pair, such as physical distance and cultural similarity. γ_t is a quarter fixed effect that accounts for nationwide trends that may influence migration, such as the overall macroeconomic climate.

Unlike for overall migration levels, the "event study" plot of quarterly wage differentials for movers suggests that parallel trends are more plausible in this context (see Figure 2.13), so I continue with a standard difference-in-difference (DiD) model. Because wages and hiring dynamics may differ across different segments of the labor market, I fit a model separately for each sub-population defined by educational attainment. As $EESEarn_{odtp}$ is an average, I weight the regression by the number of workers in each cell defined by $o, d, t,$ and p. I also winsorize the values of $EESEarn_{odtp}$ at $[0.025, 0.975]$ to account for implausibly large salary changes.

Table 2.1 portrays the estimates from this model, with a separate model for each category of educational attainment. Across all sub-populations, I find no evidence that offered wages in North Carolina increased post-HB2, relative to control states. In fact, our results suggest the opposite: after North Carolina passed HB2 in the first quarter of 2016, the wage premium for interstate migrants *decreased* over the next 15 quarters (through the end of 2019) relative to other states, although this effect is small (see Figure 2.13). This decrease was found across nearly all categories of workers, and it was largest for college-educated workers (category E4).

In the appendix, I conduct the same analysis for workers *leaving* North Carolina. A compensating differential would suggest that these workers would be willing to accept lower wages in order to leave North Carolina. Table B.1 and the event study plot in Figure B.13 show that this was not the case: relative to movers leaving other states, movers from North Carolina did not experience a change in the salary differential from moving out of state post-HB2.

Table 2.1: Difference-in-difference of change in quarterly earnings for movers, by education category

	Change in Quarterly Earnings for Movers (Dollars)				
Education	E1	E2	E3	E4	E5
	(1)	(2)	(3)	(4)	(5)
NC × Post HB2	-299.1***	-222.0**	-272.0**	-468.8**	-69.56
	(85.20)	(103.0)	(114.5)	(195.7)	(58.08)
Observations	70,765	75,041	75,950	75,776	68,640
Number of Workers	2,561,367	4,687,650	5,302,797	5,082,694	4,769,683
Origin-Destination FE	✓	✓	✓	✓	✓
Year-Quarter FE	✓	✓	✓	✓	✓

Difference-in-difference (DiD) analysis where the dependent variable is the change in quarterly wages for employer-employer movers (i.e., those who move directly from one job to another without an quarter of labor market inactivity) between their old job (in quarter $t-1$ and their new job in quarter t. Standard errors are clustered by origin-destination pair and regression is weighted by the number of individual workers used to compute the average wages in each cell defined by origin, destination, quarter, and subpopulation (education category). Observations are at the origin-destination-quarter level; missing observations differ across categories due Census Bureau data depression rules for cells with small counts. Each educational category is modeled separately: E1 (less than high school), E2 (high school or equivalent, no college), E3 (some college or Associate's degree), E4 (Bachelor's or advanced degree), and E5 (workers 24 years of age or younger, for whom educational attainment is not available). Time period begins in 2011 (20 pre-treatment periods), treatment begins in 2016 Q2, and time period ends in 2019 Q4 (15 post-treatment periods).

2.7 Conclusion

In this paper, I use generalized synthetic control (GSC) methods to show that, contrary to predictions and widespread reporting, the passage of North Carolina's "bathroom bill" in 2016 did not result in large-scale changes to migration. My estimates of the effect of this bill on migration to North Carolina are fairly precisely estimated around zero in the short- to medium-term, and in the long-term (up until the onset of the COVID-19 pandemic), migration to the state appears may have actually accelerated relative to other states.

Why did HB2 not result in decreased access to employee talent for North Carolinian employers, despite widespread claims that it would do so? There are several potential hypothesis.

Figure 2.13: Event plot of DiD Estimates for salary differences for movers to NC

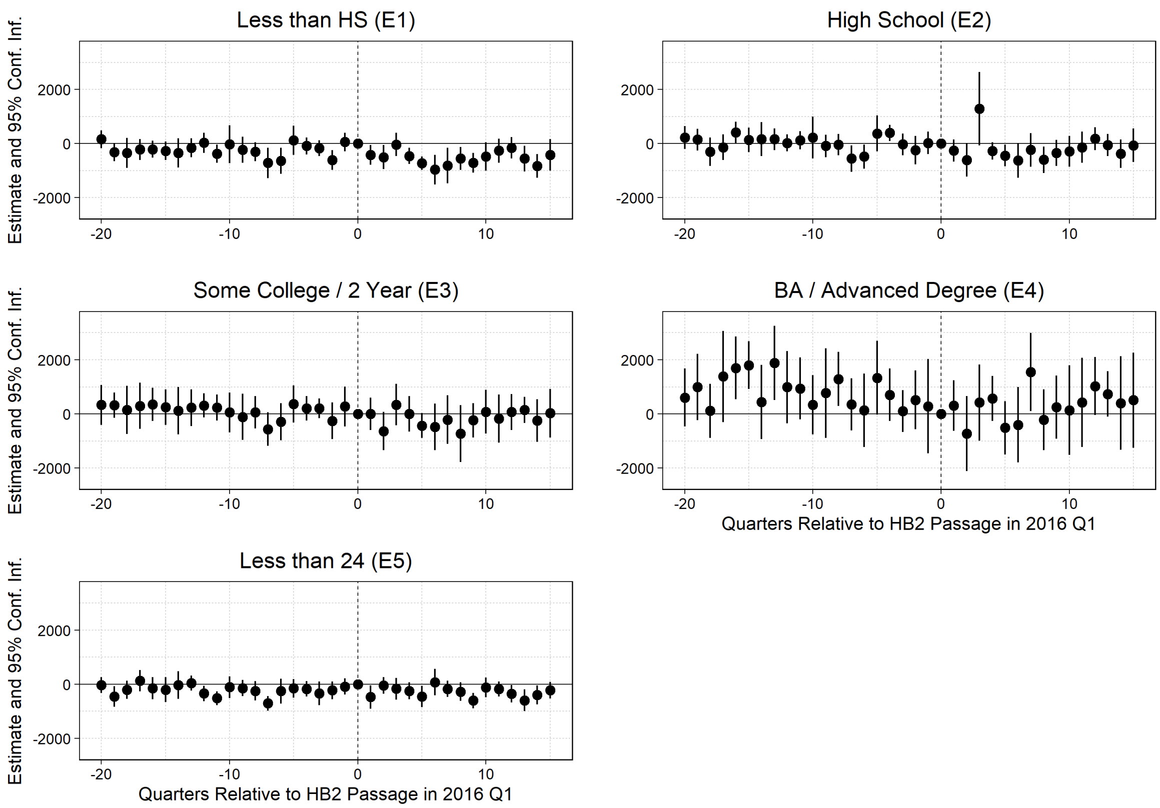

Difference-in-difference (DiD) analysis where the dependent variable is the change in quarterly wages for employer-employer movers (i.e., those who move directly from one job to another without an quarter of labor market inactivity) between their old job (in quarter $t-1$ and their new job in quarter t. Standard errors are clustered by origin-destination pair and regression is weighted by the number of individual workers used to compute the average wages in each cell defined by origin, destination, quarter, and subpopulation (education category). Observations are at the origin-destination-quarter level; missing observations differ across categories due Census Bureau data depression rules for cells with small counts. Each educational category is modeled separately: E1 (less than high school), E2 (high school or equivalent, no college), E3 (some college or Associate's degree), E4 (Bachelor's or advanced degree), and E5 (workers 24 years of age or younger, for whom educational attainment is not available). Time period begins in 2011 (20 pre-treatment periods), treatment begins in 2016 Q2, and time period ends in 2019 Q4 (15 post-treatment periods).

The first possibility is perhaps the most obvious: this issue is not important enough to spur workers to move across state lines. Those most directly affected by the bill—namely, transgender people—are a very small minority group. While transgender identification has grown significantly in recent years, at the time of the law's passage, there were only an estimated 22,000 transgender adults in North Carolina out of a total 10.2 million (Mallory and Sears 2016). Even those directly targeted by these types of bills may not respond by relocation. In a small-n indicative survey (n = 113) of the parents of LGBTQ+ children in Florida regarding the Parental Rights in Education Act (or "Don't Say Gay" Bill), 88% of these parents expressed concern, but only 56% were considering leaving the state and only 16% were actually taking steps to relocate (Goldberg 2023). While cisgender people who support the transgender community likely viewed the law unfavorably, it is possible that this specific bill did not pose enough of a threat to influence migration decisions.

Second, it is possible state-level policies may have little effect beyond pre-existing partisanship due to ceiling effects. In a recent working paper, Downey and Liu (2023) use a difference-in-difference design around party transitions in state governorships to argue that Republican control of state government is associated with declining in-migration rates among college graduates (but not among non-college graduates). One possibility, then, is that we fail to see a response to state policy changes such as HB2 because potential migrants make decisions based upon state party control; they are already in effect "pre-treated" and incremental policy changes do not affect the attractiveness of potential state destinations. In this case, North Carolina Republicans achieved a state government trifecta (i.e., control of both houses of the legislature and governorship) in 2012. By the time HB2 was passed in March 2016, the North Carolina GOP had enacted a broad agenda of policy changes (Fausset 2014).

Third, it is possible that laws such as HB2 do have effects on small, highly specific labor markets that are not detectable in aggregate data, but are important to employers seeking to fill certain roles. News reporting in this vein suggests that red states are losing professionals in certain medical fields (i.e., obstetrics and gynecology) and education (both K-12 and higher education) (Noah 2023). Once again, existing evidence for these claims is mixed. For instance, after the *Dobbs* decision, an analysis by a lobby group, the Association of American Medical Colleges, argued that states which restricted abortion had fewer medical school graduates apply for residencies in obstetrics and gynecology (OB/GYN) (Orgera, Mahmood and Grover 2023), but a subsequent analysis by the National Resident Matching Program (NRMP) found that there were no shifts in medical school graduates' location preferences or in restrictionist states' ability to fill OB/GYN residencies (The National Resident Matching Program 2023).

Fourth and finally, it is possible that because the widespread and highly-public opposition within North Carolina—especially in more urban areas that might be disproportionately likely to attract in-migrants—may have neutralized the negative impact of the bill itself. In related work, Hurst (2023) finds that firms in Charlottesville, Virginia deployed pro-diversity claims to counteract the stigma that area firms faced after the Unite the Right white supremacist rally in their town. It is possible that the sustained anti-HB2 activism—including opposition by the business community—may have fulfilled this role by reassuring workers that many North Carolinians opposed the bill.

Chapter 3

This Bud's for You? The Effect of Partisan Cues on Politically Polarized Brands

3.1 Introduction

On April 1, 2023, Dylan Mulvaney, a social media influencer, posted two short videos featuring Bud Light. Mulvaney had previously come out as a transgender woman a year prior. An advertising manager at Bud Light had sought to take advantage of this anniversary as an opportunity for "sponsored content" promoting the brand to Mulvaney's young, socially progressive, and LGBT-friendly audience—a group not typically associated with Bud Light's blue-collar reputation.

A major backlash soon followed. Conservative figures ranging from Republican presidential candidates to country music stars quickly responded with their own social media posts condemning the brand and pledging not to buy Bud Light. Anheuser-Busch, Bud Light's parent company, quickly began various efforts to limit negative publicity. The advertising executive responsible for endorsement was placed on leave, and the CEO issued a public apology. These public relations efforts did little to contain the damage to the brand. By June, industry sales figures showed that Bud Light had lost its longstanding position as America's best-selling beer.[1]

This event represents a dramatic manifestation of a broader phenomenon. Brands which become associated with a particular partisan stance—whether deliberately or inadvertently—often see a polarization in consumption patterns, where opponents of the party or political stance conduct a boycott, while supports conduct a simultaneous "buycott." Studies of real-world businesses show the potential for customers' buying behavior to polarize along partisan lines after brands take political stances or are spoken about (favorably or unfavorably) by politicians (Hou and Poliquin 2023; Liaukonytė, Tuchman and Zhu 2022, 2023; Endres, Panagopoulos and Green 2021; Chatterji and Toffel 2018). Experimental evidence also shows that informational cues linking companies to a specific political party can induce these polarized patterns in purchase intentions

[1] Moreno, J. Edward. June 14, 2023, "Bud Light is No Longer America's Top-Selling Beer After Boycott. *The New York Times*. `https://www.nytimes.com/2023/06/14/business/bud-light-lgbtq-backlash.html`.

or favorability towards companies (Dodd and Supa 2014; McConnell et al. 2018; Panagopoulos et al. 2020).

These studies have shown that companies that neutral or apolitical companies can acquire a partisan reputation in response to side-taking or information associating them with a particular party, which in turn leads to a polarization in consumer preferences. However, because of this focus on neutral or apolitical brands, it is unclear whether strongly political cues or information can redefine corporate reputations in cases where attitudes have already polarized. If deliberate stance-taking (Hou and Poliquin 2023), political endorsements or condemnations by prominent politicians (Endres, Panagopoulos and Green 2021), or the provision of information about political spending (Panagopoulos et al. 2020) can shift preferences to induce polarization, could similar political cues *reduce* existing polarization?

This study tests the limits of partisan cues in shaping consumer preferences. It focuses on Bud Light, a brand which has become highly politicized since spring 2023. Using a pre-registered survey experiment carried out in June 2023 (two months into the boycott of Bud Light), I first confirm that beer consumption has polarized along partisan lines. Regardless of party, American beer drinkers tend to favor similar brands—with the notable exceptions of Bud Light and Budweiser, which Democrats are far more likely to consume. I then expose respondents to a treatment which consists of a short vignette in the style of a news story, ostensibly for the purposes of asking about their awareness of stories in the news. Different treatments emphasize different parties (Democrats or Republicans) and the nature of their stance towards Bud Light (negative or positive). Following the treatment, respondents were given a series of choice-based conjoint tasks asking them to choose between two six-packs of beer at various price points, in different formats (can vs. bottle), and at different price points. They were also asked about whether or not they were participating in the boycott of Bud Light.

Partisan political cues had no effect on Republican political attitudes. Republicans already appeared highly aware of the Bud Light story and entrenched in their preferences. Cues that might be expected to increase Republican antipathy towards the brand (information about Republican boycotts or Democratic support for the brand) had no effect, but neither did cues that might be expected to reduce antipathy (i.e., information about Republicans' support for the brand or information about Democratic boycotts). While the brand's earlier implicitly partisan stance in support of the LGBT community clearly diminished the brand's appeal for Republicans, this reputation appears to have been firmly entrenched by the time the survey took place (approximately two months into the boycott). Among Democrats, however, negative in-party cues (suggesting that Democrats were boycotting Bud Light) and positive out-party cues (suggesting that Republicans were supporting Bud Light) resulted in a decreased probability of selecting Bud Light. Democrats exposed to these treatments were also more likely to self-report boycott participation, although this appears possibly driven by Democrats misreporting their own past behavior.

There are two important lessons we can draw from this study. First, these findings suggest important limits to the ability of partisan cues to shape consumer preferences. For instance, Endres, Panagopoulos and Green (2021) build upon an extensive literature on the role of elite cues in political science (e.g., Lenz 2012; Zaller 1992) to argue that partisan cues can polarize perceptions of major brands. To the extent that scholars have found that elite cues can mold supporters' political preferences, they tend to focus on policy positions that are relatively unfamiliar or novel, rather than core political values or highly-charged partisan issues (Lenz 2012, 2009). Similarly, these findings suggest a limit to the ability of partisan cues or information to re-shape consumer preferences. Partisan cues appear to exert a one-way ratchet effect: it is far easier to take a previously neutral or apolitical brand and induce polarization, but partisan cues appear less effective at undoing existing polarization.

In addition to the impact for scholars, this has an important implication for managers. It suggests first that managers should be cognizant that any politically salient actions, even inadvertent ones, may trigger a polarization in consumer preferences that is difficult to reverse. It also suggests that brands, once polarized, may find it difficult to undo this polarization by trying to triangulate across different political groups or attempting to tack back to the center.

Second, this study reinforces that self-reported data on political consumerism is highly susceptible to motivated reasoning and inaccurate recall. When exposed to a treatment that suggested that Democrats were boycotting Bud Light, Democrats in the survey were over 50% more likely to report that they had been engaging in a boycott. Notably, this question was *not* about future boycott intentions, but about past behavior, so there is no reason why a survey treatment should have affected past behavior. Furthermore, nearly two-thirds of Democrats who claimed to be boycotting post-treatment had earlier (pre-treatment) stated that they had recently consumed Bud Light in the last 30 days, casting doubt about the accuracy of their reports. This finding is consistent with studies which show that the volume of social media posts related to political consumerism vastly outweigh any measurable economic impact (Liaukonytė, Tuchman and Zhu 2022). It also suggests the need for caution when relying upon self-reported survey data when studying political consumerism.

3.2 Theory and Related Literature

Political consumerism refers to individual group or consumption decision to boycott or buycott specific products or services out of political, social, or moral considerations (Bennett and Entman 2000; Newman and Bartels 2011; Micheletti 2003). While the phenomenon dates back centuries (the term "boycott" comes from Charles Boycott, a 19th century Anglo-Irish landlord), the term "political consumerism" gained traction in the 1990s, in response to events such as the environmentalist boycott of Shell Oil (Micheletti 2003). While long-run surveys of political consumerism are limited, the phenomenon appears to have increased over time (Endres and Panagopoulos 2017; Stolle and Micheletti 2013).

Earlier scholarship on political consumerism tended to emphasize instrumental motivations for political consumerism, in which social movements or activists target companies to directly influence corporate conduct or indirectly influence government policy (Soule 2009; King 2011; Vogel 2006). More recent scholarship has posited a theory of expressive political consumerism (Endres and Panagopoulos 2017), building on a large body of work in political science which emphasizes how political partisanship serves a social identity (Iyengar et al. 2019; Druckman and Levendusky 2019; Mason 2015). Recent growth in political consumerism coincides with a rise in *affective* polarization between partisans. Affective polarization is rooted in social identity (Tajfel and Turner 1979) and an antipathy towards the opposing party that is based upon in-group/out-group distinctions, rather than necessarily being based upon disagreements over political parties. Rising affective polarization has been shown in influence partisans' decision-making in a range of economic domains, including in hiring (Gift and Gift 2015), job acceptance and effort (Burbano 2021), and decisions about with whom to do business in general (McConnell et al. 2018).

Given the role of partisanship, prior work has found that companies which become associated with a particular partisan or ideological position may find their customer base becomes polarized. This has been found both experimentally (Banda, Carsey and Severenchuk 2020; Panagopoulos et al. 2020) and using observational data of real-world episodes of political consumerism and measures of brand perceptions (Endres, Panagopoulos and Green 2021; Chatterji and Toffel 2018), sales figures (Liaukonytė, Tuchman and Zhu 2022, 2023), or store visits (Hou and Poliquin 2023).

In all these episodes, respondents receive new information about a previously neutral or apolitical brand. This partisan cue may arise in multiple ways that are not necessarily mutually exclusive. In some cases, this partisan cue may come from a public stance made by company itself (Chatterji and Toffel 2019). For instance, following a mass shooting at one of its stores, Walmart publicly announced that it would no longer sell assault weapons and some types of ammunition. Given the nature of debates about gun control in the United States, this specific policy was viewed unfavorably by Republicans, who responded by becoming less likely to visit Walmart stores (Hou and Poliquin 2023). Another case is when Apple CEO Tim Cook spoke out against a proposed religious freedom bill in Indiana (Chatterji and Toffel 2019). A similar example comes from the large number of companies which announced changes to their political giving following the January 6 Capitol riot (Li and Disalvo 2022; Yang and Jia 2023).

In other cases, the partisan cue arises from events not within the company's control. This may be because of an external call for a boycott, such as when the musician Neil Young and activists launched a boycott campaign against Spotify to protest COVID-19 misinformation on the Joe Rogan Show, a podcast available on its platform (Liaukonytė, Tuchman and Zhu 2023). Brands may acquire a partisan reputation as a result of the dissemination of information about political spending. While in some cases, this spending may be secret (e.g., Werner 2017), most members of the public lack awareness of even publicly-available information about political spending by companies they patronize (Panagopoulos et al. 2020). When this information is made public, as in an experiment by Panagopoulos et al. (2020), attitudes towards the brand polarize.

Finally, brands may acquire a partisan reputation as a result of a political cue that does not necessarily contain specific information about the brand, but is merely an invocation of partisan loyalties. In an experiment, Banda, Carsey and Severenchuk (2020) ask respondents about hypothetical products and services which are more favored by members of one specific party. They find that hypothetical products (e.g., a soft drink or vacation resort) which are described as being typically purchased or patronized by members of the out-party are seen as less desirable. In addition to cues about fellow partisans, these political cues may come from political elites. While President, Donald Trump tweeted about a number of brands, including L.L. Bean (which he spoke of positively) and Macy's and Nike (which he criticized) (Endres, Panagopoulos and Green 2021), causing polarization in these brands' reputations among partisans. Even if devoid of specific information about the brands, these cues from a political elite may induce changes in partisans' attitudes. This finding is related to research on political belief formation among the public shows that partisans often "follow the leader" by adopting the policy positions of their favored politicians (Lenz 2012; Broockman and Butler 2017; Zaller 1992), and that there may be backlash in which members of the opposing party may adopt a contrary position (Zaller 1992). In many cases, this may be because policy issues are complicated or difficult-to-understand, so partisans rely upon elite cues as a heuristic. However, research in social psychology also posits a simpler explanation, based upon the need for groups to coordinate upon adopting symbols as clear markers of in-groups and out-groups (Tajfel and Turner 1979).

While research has established that partisan cues can induce polarization among brands, other studies have suggested that there may be limits to this effect. In many cases, boycotts and buycotts prove to be short-lived, with effects dissipating quickly as news cycles move on (Liaukonytė, Tuchman and Zhu 2022, 2023). Some brands and services may be more susceptible to political consumerism. When switching costs are high or there are few existing alternatives, customers are less likely to sustain boycotts (Liaukonytė, Tuchman and Zhu 2022). Reactions may also be asymmetrically balanced, with either boycotts or buycotts being larger. Chatterji and Toffel (2019) found that gay marriage supporters were more likely to state an intention to purchase Apple products after being informed CEO Tim Cook opposed state legislation that would limit LGBT protections, while there was no statistically significant backlash among those who opposed gay marriage (Chatterji and Toffel 2019). At the level of individual psychology, negative cues towards boycotting may be more potent than positive cues to induce buycotts (Kam and Deichert 2020). At the aggregate economic level, however, there are limits to the magnitude of boycotts—only existing customers are in a position to boycott, whereas anyone can begin to participate in a buycott (Liaukonytė, Tuchman and Zhu 2022).

To date, however, nearly all studies have focused on whether partisan cues have the potential to *increase* political polarization. As such, both experimental and observational studies focus on brands which did not have a strong partisan tilt prior to a particular incident (or experimental treatment). This study, by contrast, focuses on a brand which already has a sharp divide in brand perceptions among partisans, and is currently being boycotted by members of one party. It asks whether political cues can potentially unwind existing polarization. This is not of purely academic interest. Many brands, when faced with a boycott from one particular political party, may attempt to tack back to the center by trying to curry favor or mollify members of the opposing party.

3.3 Methods

Sample

The survey was fielded from June 19–June 20, 2023 on Lucid Marketplace. The survey was generically described as a study of purchasing behavior and attitudes towards current events. It was open to respondents of legal drinking age (21) who stated that they regularly consumed beer, and who reported having consumed at least one of nine popular American beer brands (e.g., Bud

Light, Budweiser, or a competitor). Respondents also had to pass a series of pre-treatment attention checks. After excluding a small number of respondents who have low-quality responses,[2] I was left with 3,000 respondents. Of these, 1,415 (47%) were Democrats or Democratic leaners,[3] 1,114 were Republican or leaners (37%), and 411 (14%) were (true) independents. While it is not necessarily the case that beer-drinking population will mirror the overall U.S. population, the sample does not appear to differ in large ways from overall U.S. demographics on dimensions including gender, LGBT identity, education, and income. As with many online samples, it does skew younger than the U.S. population.[4]

After completing a screener and demographic information section, respondents were asked about their beer consumption habits (see Table C.2). At the time the survey was fielded in mid-June, the boycott had been occurring for roughly two months, and had been widely covered in the media. Given this, I anticipated that a majority of respondents would already be aware of the boycott, and that attitudes towards Bud Light would have polarized among partisans. The survey confirmed that this was the case. When asked about recent beer consumption (pre-treatment), Republicans were about 16 p.p. less likely than Democrats to report that they had consumed Bud Light in the last 30 days (see Table C.2).[5] Republicans were also more likely to report that they had consumed Bud Light in the past but not in the past 30 days ($\sim$ 11 p.p.), indicative of a greater likelihood that they are boycotting the brand.

Treatment

Respondents were then told that the next section would ask them about their news consumption habits, and asked how they follow the news (i.e., what types of media they consume). They were then asked to read two short vignettes written in the style of a news story, and then asked if they had previously encountered the story. The first story neutral "distracter" story was a short blurb about the recent coronation of King Charles III in England. The second story constituted the main treatment. After each story, respondents were asked a simple multiple-choice question on the story's topic, and if they had previously heard about the story. Across all six stories (the first "distracter" task and the five treatments), over 97% answered these correctly.

The main treatment consisted of one of five vignette treatments written in the style of a news report. The four main treatments report different aspects of the Bud Light boycott. While the bulk of the real-world news reporting emphasized that conservatives and Republicans were boycotting the brand, some Republicans did not support the boycott and instead spoke out in support of the brand. For their part, some Democrats and liberal groups[6] supported the brand in the face of

[2]234 respondents were excluded for straight-lining, e.g., giving the same ordinal response level to 5 or more questions in sequence.

[3]Following standard practice, respondents were asked about their partisan identification using a branching series of questions that result in a 7-point scale. Those who initially identify as neither party but then "lean" towards a party in the second question are grouped with partisans.

[4]Full descriptive statistics are available in Appendix

[5]Similar effects of a smaller magnitude ($\sim$ 10 p.p.) were also found for Budweiser, the sister brand to Bud Light.

[6]For the purposes of this paper, I treat ideology (liberal vs. conservative) and partisanship as synonymous, although

Republican boycotts, while were upset that the brand attempted to distance itself from Mulvaney and did not issue a statement supporting the LGBT community.

The treatments emphasized stances taken by both specific partisan elites as well as rank-and-file partisans. The vignette treatments included a picture of some Bud Light beer and the headline "[Republicans/Democrats] [Boycott/Show Support for] Bud Light After Controversy," and ended with the line "While it remains to be seen how much these events will affect the company, it is clear that many [Republicans/Democrats] [will not be drinking Bud Light any time soon/are prepared to move on and share a Bud Light with their fellow conservatives/may soon be cracking open a Bud Light to share with their fellow liberals]" (full treatment language is available in Appendix C.1). While each treatment covers different aspects of Bud Light's recent brand difficulties, all vignettes are based on true reporting; no deception was employed.

The first treatment ("Republican Boycott") summarized the boycott by conservatives and Republicans. It specifically mentioned Florida Governor Ron DeSantis' critical comments about the beer, as well as comments made by (unspecified) Republican congressmen. It also mentioned a video posted by Kid Rock, a 1990s musician, in which he destroyed cans of Bud Light with a submachine gun.

The second treatment ("Republican Support") emphasized Republicans who did not take part in the boycott or who attempted to support Bud Light. Most prominently, Donald Trump, Jr., the son of the former President, used his podcast to express support for Anheuser-Busch in light of the brand's longstanding support for the Republican party and its subsequent apology for the Mulvaney promotion. The treatment also mentioned that Bud Light remained available at Trump hotels and resorts.[7]

The third ("Democratic Support") treatment mentions a social media post by Rep. Ted Lieu, a California Democratic representative who posted a picture of himself with fellow Democratic colleagues drinking a Bud Light. It also reports that members of the Facebook group "Occupy Democrats" were going out of their way to purchase Bud Light to support the brand and show their opposition to Republicans.

Finally, the fourth ("Democratic Boycott") treatment explained that some gay bars were conducting a boycott of Anheuser-Busch after the company attempted to distance itself from the transgender community,[8] and that the Human Rights Council, a prominent LGBT advocacy organization, had removed the company from its corporate equality index.

A final neutral ("Control") condition is headlined "New Beer Industry Report Shows $409 Billion in Economic Impact" and is excerpted from a press release from the Beer Institute, an

they are theoretically distinct. I also assume that officially nonpartisan organizations—such as gay bars or the Human Rights Council, a nonprofit that supports LGBT rights—are viewed by the public as being liberal and aligned with the Democratic party, in keeping with prior work which demonstrates that members of the public are readily able to assign these groups to the appropriate ideology and party (Elder and O'Brian 2022).

[7]Fung, Katherine, June 9, 2023. "Trump's Still Selling Bud Light. *Newsweek.* `https://www.newsweek.com/trumps-still-selling-bud-light-1805565`

[8]Valle, Jay, May 10, 2023. "Chicago gay bars boycott Anheuser-Busch for distancing itself from Dylan Mulvaney." *NBC News.* `https://www.nbcnews.com/nbc-out/out-news/chicago-gay-bars-boycott-anheuser-busch-distancing-dylan-mulvaney-rcna83537`

industry body. It addresses the same broad topic (beer), but without any partisan or ideological content.

Respondents were told that the purpose of the survey was to understand their purchase behavior as well as gather information related to their news consumption. Respondents were asked several questions about their news consumption, and then told to read two short simulated news vignettes. The news stories were created for the purpose of this survey based upon actual news reporting and written in a standard news wire style, with the help of Open AI's Chat GPT.

Dependent variables

One concern with research on political consumerism is that survey respondents may overstate the extent of their participation in political consumerism (Endres and Panagopoulos 2017). Recent studies which compare social media data with actual purchase data also suggest that many individuals may overstate the extent to which they are actually engaged in political consumerism (Liaukonytė, Tuchman and Zhu 2022). To better capture the real-world effects of political consumerism, prior studies have relied upon observational data such as store footfall or product sales (Hou and Poliquin 2023; Liaukonytė, Tuchman and Zhu 2022). Experimental studies also attempt to focus on costlier measures of actual behavior rather than pure attitudinal measures, such as a choice between gift cards with real economic values (Panagopoulos et al. 2020).

To attempt to accurately measure consumer attitudes, I do two things.

First, I ask about boycott intentions in two ways. Combined, these allow me to generate rough estimates of the extent to which stated consumer preferences may be overstated. As part of the initial screener for the survey, I ask about respondents' recent consumption of different beers. Respondents were given a list of nine beers and asked if they had consumed the beer (a) in the last 30 days, (b) at some point prior to the last 30 days, or (c) never. Crucially, these questions are asked *prior* to any treatment. Later on in the survey, *after* the treatment, respondents were asked if they had been participating in the boycott of Bud Light. Both of these questions focus on past behavior, so there is no reason why treatment assignment should alter responses. However, by asking both before and after treatment, I am able to compare responses for consistency. A respondent who states both that they are boycotting Bud Light and that they consumed it in the last 30 days may likely be misreporting their own past behavior.[9]

Second, my main measure comes from a choice-based conjoint task. Conjoint studies are widely used in marketing, political science, and other social science disciplines to study situations in which respondents must select between different options that differ along multiple dimensions, such as choosing a candidate in an election or a product to buy. Because each option (or "profile") contains multiple features, conjoint studies allow researchers to study how respondents trade-off between multiple conflicting preferences. While this technique was initially developed to help marketers understand how customers value product features, social scientists have also

[9]Respondents may also have begun to boycott Bud Light more recently than the last 30 days. I believe this is unlikely, as the boycott had already been occurring for two months at this point, and most boycotts tend to diminish rather than grow over time. Furthermore, even if this is the case, there is no reason why the likelihood of this possibility should differ across treatment groups.

found that the technique is useful in avoiding social desirability bias and elucidating more reliable measures of respondents' true preferences.

Figure 3.1: Example of choice-based conjoint task

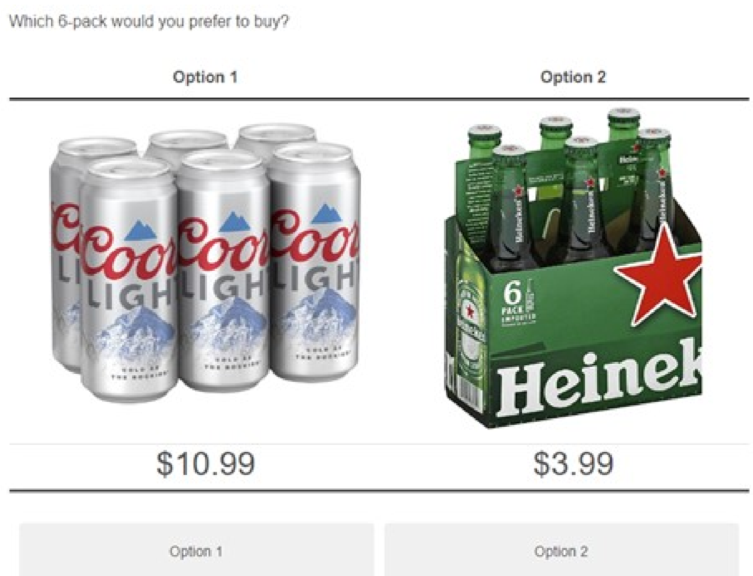

The conjoint in this study asks respondents to choose between two six-packs of beer, each of which differs along three dimension: (a) brand, (b) format (i.e., can or bottle), and (c) price (see Figure 3.1).[10] To make the experiment more naturalistic and easier for respondents, they are shown an image of a six-pack which clearly conveys both brand and format. Given the strong use of color and other branding elements in beer packaging, a picture is likely immediately recognizable, making the overall task easier as tabular text-based designs often used in choice-based conjoints. The prices are randomly-generated dollar amounts ending in $.99, ranging from $3.99 to $14.99. While prices of a six-pack of beer differ substantially by brand, geography, and venue, this represents a realistic range of standard beer prices. Respondents were given ten tasks (choices) to complete. While this number may seem high, Bansak et al. (2018) find little evidence of respondent fatigue in up to 30 tasks. Relative to other conjoint tasks which may incorporate numerous features and present data in a tabular text-based format, this is a very simple and thus relatively cognitively undemanding task, making respondent fatigue less of a concern.

In designing the survey, one concern was that this conjoint experiment focuses on buying beer "off-premises" (i.e., from a store), whereas some respondents may more typically purchase beer in bars or a restaurant. Another concern was that some respondents may drink beer at home, but are not themselves the decision-maker in the household (i.e., someone else does the

[10] According to data from the Beer Institute, an industry body, aluminum cans (57%) and glass bottles (32%) constitute the most common formats for beer sales, and the six-pack is the most popular size. https://www.beerinstitute.org/trends-beer-packing/

shopping). Both concerns proved to be unfounded. 90% of those in the sample drank beer at home, versus only 58% who drank beer elsewhere. 96% stated that they regularly purchased beer "off-premises," and all but 29 (99%) stated that they were either the sole or joint decision-maker for beer purchases.

3.4 Results

Conjoint

There are several methods of estimating consumer preferences using choice-based conjoint tasks. Business practitioners typially use widely-available commercial software models which fit multinomial or conditional logistic regression models. More recent social scientific methodologists suggest instead relying upon the average marginal component effect (AMCE) (Hainmueller, Hopkins and Yamamoto 2014). This estimator can be fitted nonparametically without reliance upon arbitrary functional form assumptions, or indeed any assumptions save for those which are guaranteed to hold from the conjoint setup itself. It also has an intuitive causal interpretation—for any given element, the AMCE can be interpreted as the change in the likelihood that a given profile is selected.

Figure 3.2 shows differences in marginal means (MM), a theoretically related estimator. Because I compare across multiple treatment and party subgroups, I show marginal means or differences in marginal means (clustered by respondent) rather than average component marginal effects (ACMEs), which are sensitive to the researchers' selection of baseline reference catgory (Leeper, Hobolt and Tilley 2020). First, it is reassuring to note that respondents' general preferences suggest that respondents took the conjoint task seriously. As one would expect, respondents are more likely to choose cheaper options.[11] Consistent with industry statistics, bottles were slightly more likely to be chosen than cans. Brand choice also appeared consistent with trends reported in the initial pre-treatment questionnaire: Corona was the second-most popular beer in the screener (45.1% reported drinking it within the last 30 days) and Busch Light the least (17.5% drank it in the last 30 days); these were also the most- and least-commonly selected in the choice-based conjoint. All this suggests that the the choice-based conjoint setup was taken seriously by respondents and their stated preferences in the experimental task bear a resemblance to their actual (reported) behavior.

Among Democrats, there were notable and significant differences in the probability of selecting Bud Light depending on treatment assignment. In the "Republican Support" and "Democratic Boycott" conditions—both of which should signal to Democrats not to buy Bud Light—the brand was selected 53.3% and 52.6% of the time, respectively. These are statistically significantly different than both the the "Democratic Support" treatment (which provided information about how certain Democratic Congresspeople as well as rank-and-file Democrats were going out of their

[11] In a supplemental analysis, not reported here, this tendency was less pronounced among those who reported household incomes of greater than $150,000, which further suggests respondents answered these choice tasks realistically.

Figure 3.2: Conjoint marginal means by treatment and party

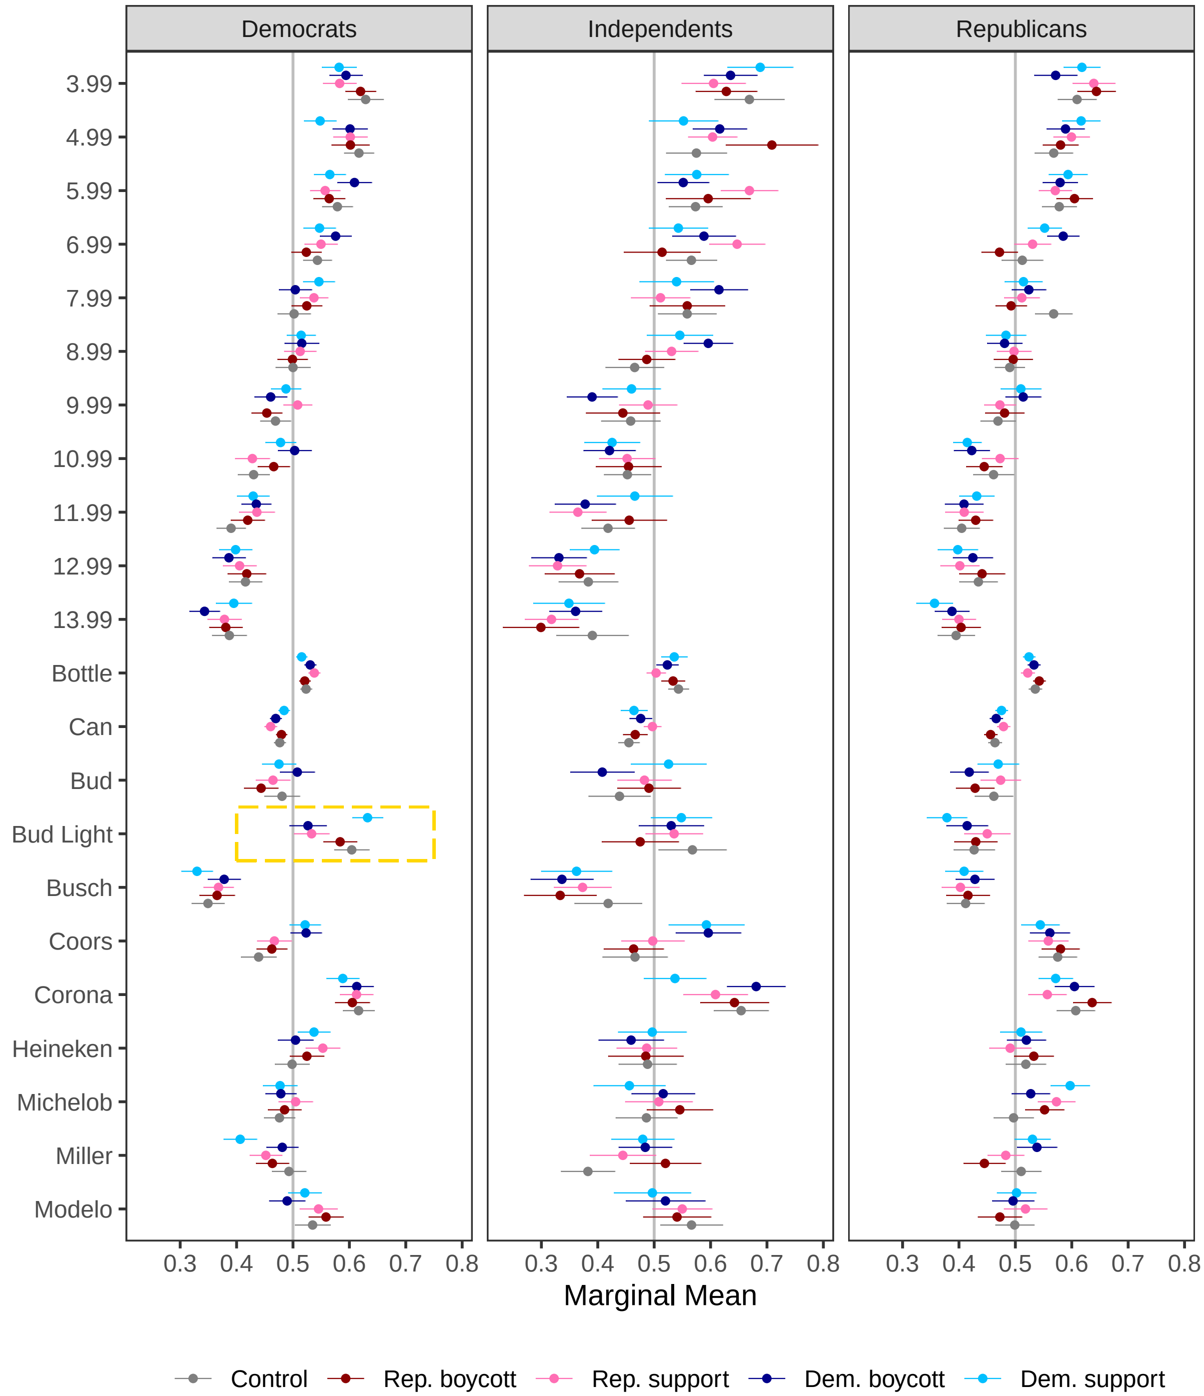

Conjoint marginal means by treatment and party. Standard errors are clustered by respondent. Bars depict 83.4% confidence intervals, such that non-overlapping lines represent a significant difference in marginal means at the $\alpha = 0.05$ level of significance; see Knol, Pestman and Grobbee (2011).

way to purchase Bud Light) and the control treatment, which were selected by Democrats 63.2% and 60.4% of the time, respectively.

The differences are presented more clearly in Figure 3.3, which shows the *differences* in marginal means by treatment and party. Relative to the neutral control treatment, both the "Republican Support" and "Democratic Boycott" treatments significantly reduce the changes that Bud Light is chosen. The "Democratic Support" treatment is positive but not significant. Finally, the Republican boycott treatment—which many respondents may already have been aware of, given the volume of news coverage—had no significant effect.

Among Republicans and independents, there were no effects. Independents are unlikely to be swayed by partisan appeals by either party, so this is not surprising. Republicans appear also unconvinced by partisan treatments. Likely this is because attitudes towards Bud Light are are already fairly hardened. A large number of Republicans reported that they were already aware of the boycott, and so any additional information received in the treatment vignette does not appear to change strongly-held attitudes towards the brand.

Nevertheless, given the wide array of media coverage towards the Bud Light boycott, it is remarkable that I was able to observe *any* persuasive effect from the treatments. The findings for Democrats suggest two conclusions. First, even in already highly-contested partisan contexts, partisans may change their consumption behavior to comport with group norms. Democrats who were told that other Democrats were boycotting Bud Light were less likely to choose the brand. Second, out-party cues appear just as strong as in-party cues. Democrats were also less likely to choose Bud Light when told that notable Republicans—Donald Trump and Donald Trump, Jr.—were standing by the brand and speaking out in support of it. This suggests that partisan cues can and do contribute to political consumerism.

Boycott participation

In addition asking respondents to choose between brands in the conjoint task, respondents were also asked if they had been engaged in boycotting Bud Light. Specifically, respondents were asked the following question:

> A "boycott" is when someone refuses to buy a product as a protest. As you may have heard, many people are now boycotting the beer brand Bud Light.
>
> Have you personally participated in the boycott of Bud Light?

This language was designed to measure *past* behavior. As such, there is no reason to expect that treatment assignment should have affected accurate reporting of respondents' past behavior.

Table 3.1 reports the results of a linear probability model (OLS) of the probability of stating that participants were boycotting, by treatment. Respondents are split by party. Among those assigned to the neutral control condition, 33.6% of Republicans reported boycotting, versus only 11.9% of Democrats and 13.2% of Independents.

However, the "Democratic Boycott" treatment appeared to have a significant effect on inducing Democrats to report that they had previously been boycotting Bud Light. Given the wording

Figure 3.3: Differences in conjoint marginal means by treatment and party

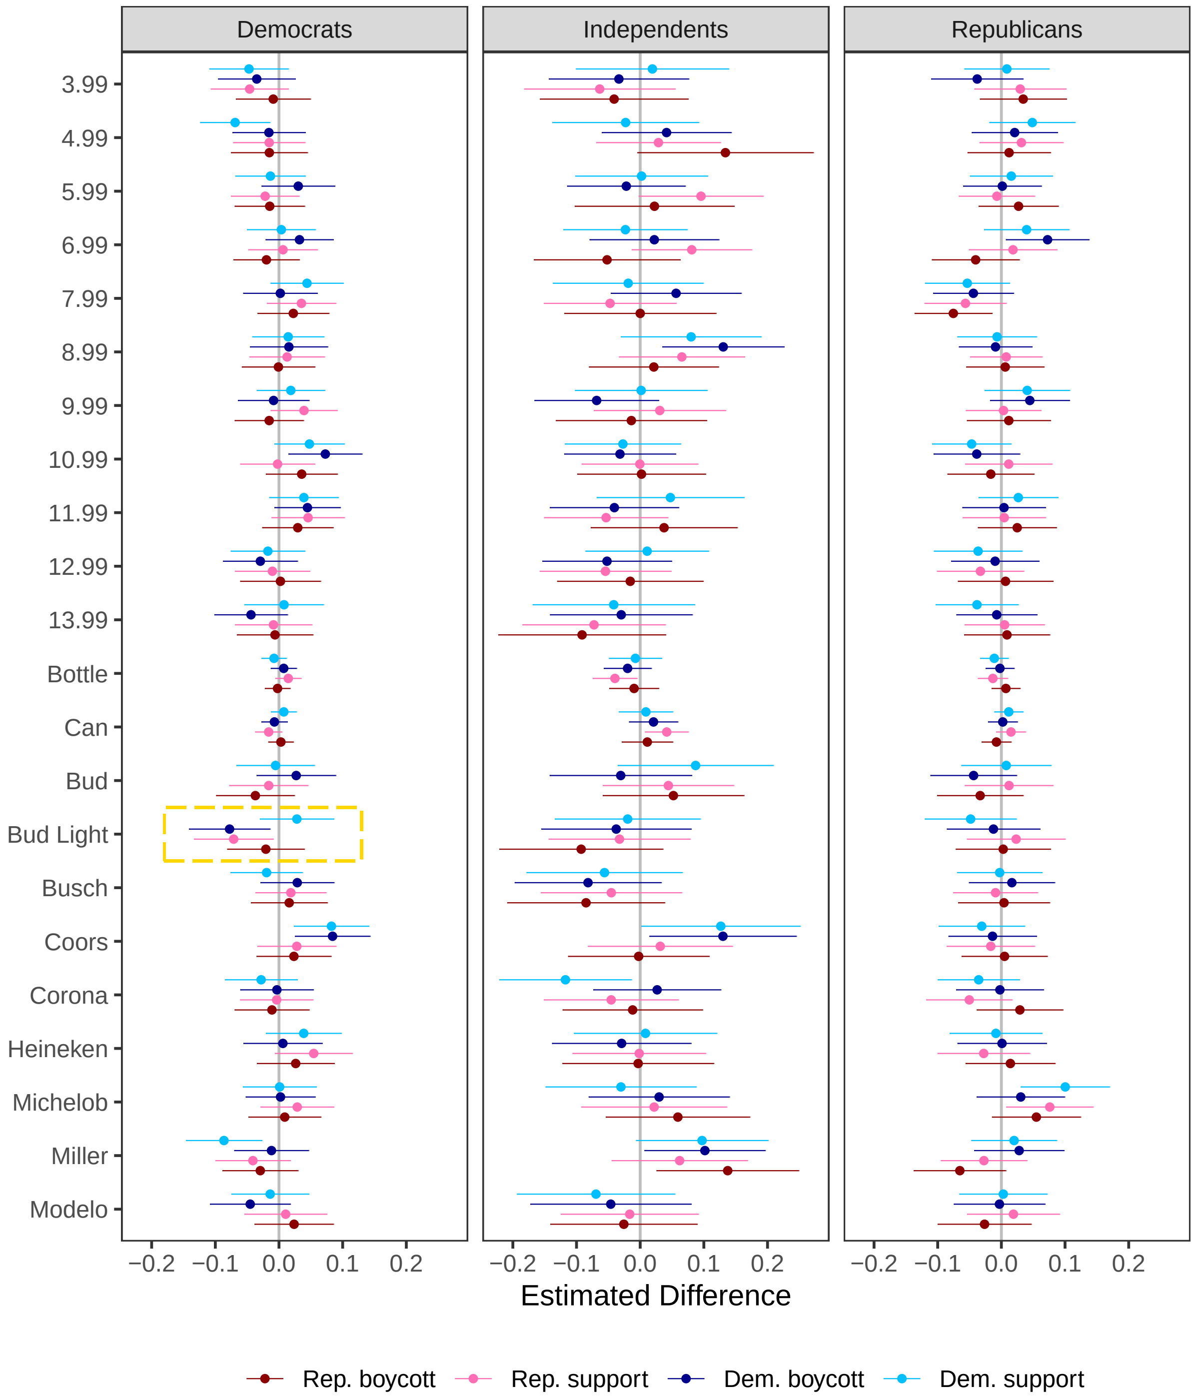

Differences in conjoint marginal means by treatment and party. The baseline is the control treatment. Standard errors are clustered by respondent.

of the question, i.e., "have you personally participated," the question is about respondents' past behavior, so if respondents are truthfully and accurately reporting, there is no reason to expect that treatments should have any effect. Under assignment to this treatment, Democrats increase their own reporting that they are boycotting. Given that this question asks about respondents' past behavior, there is no reason to expect that these should differ by treatment, given true and fully accurate reporting.

Table 3.1: Treatment effects on self-reported boycott participation, by party

	Boycott (self-reported)		
	Democrats	**Independents**	**Republicans**
Republican Boycott	-0.021	-0.102^{**}	0.008
	(0.027)	(0.052)	(0.045)
Republican Support	0.001	0.000	-0.0001
	(0.028)	(0.048)	(0.045)
Democratic Boycott	0.078^{***}	0.055	-0.033
	(0.027)	(0.048)	(0.045)
Democratic Support	-0.017	-0.046	-0.026
	(0.027)	(0.051)	(0.044)
Control	0.119^{***}	0.132^{***}	0.336^{***}
	(0.019)	(0.034)	(0.032)
N	1475	410	1114
R-squared	0.012	0.025	0.001

$^{***}p < .01$; $^{**}p < .05$; $^{*}p < .1$

The likely conclusion is that respondents are changing their reporting of past behavior to comport with a partisan cue. When told that Democrats are boycotting Bud Light, Democratic respondents are more likely to believe that they themselves should also be boycotting to match the group norm, and thus they report that they themselves are boycotting—even if they have not actually been doing so. Lest their boycotting behavior possibly be confused or misinterpreted as favoring Democrats, they decrease their reported probability of boycotting.

Additional evidence can be found by comparing respondents' answers to the boycott question with their earlier (pre-treatment) answers to questions about which beers they had recently consumed. Logically, respondents who are boycotting should (mostly) answer that the had consumed Bud Light in the past, but not in the last 30 days. There are two ways respondents may misreport their own boycotting. First, respondents may falsely overstate their boycott because they claim to be boycotting, but also consumed Bud Light in the last 30 days. This is not necessarily definitive—some respondents may have began a boycott more recently—but a high rate of respondents who claim to be boycotting but also recently consumed Bud Light is at least suspicious. Second, respondents may claim both to be boycotting but also to have *never* consumed

Bud Light. This type of misreporting stems from a different understanding of what it means to be boycotting. As Liaukonytė, Tuchman and Zhu (2022) note, many widely-publicized boycotts may fail to have a significant impact on the targeted brand because many of those who claim to be boycotting were never customers of the brand in the first place. While they may accurately and earnestly believe that they are "boycotting," refusing to buy a product one never had any intention of buying does not have any economic effect. A customer who has never purchased a particular brand may claim that they are boycotting, but the term—as understood by activists and scholars—requires that respondents first be customers before stopping their patronage.

Further evidence of Democratic over-reporting of boycotting can be found in Figure 3.4, which depicts only respondents who claimed to be boycotting. Among independent and Republican boycotters, the breakdown is as one would expect. The vast majority of boycotters used to drink Bud Light (more than 30 days ago), but have not done so within the last 30 days. Among Democrats, however, nearly 2 out of 3 alleged boycotters also report having drank Bud Light within the last 30 days, suggesting that many of the Democrats who claim to be boycotting may have been over-stating their political consumerism.

Table 3.2 further suggests that this is the case. It is a linear probability model where the dependent variable is the likelihood of "false" or inconsistent reporting—i.e., claiming to be boycotting while also stating that one has consumed Bud Light within the last 30 days (or claiming both to be boycotting and to never have consumed Bud Light). The coefficient on Democratic Boycott is positive and statistically significant, suggesting that it is the respondents who were told that other Democrats were boycotting who are most likely to give inconsistent responses indicative of misreporting their own boycotting behavior to fit in with the partisan norm.

There also may be Republican misreporting in the opposite direction. 21.5% of Republicans claim that they never have drank Bud Light. Given that Bud Light was, until recently, the most popular brand in U.S., this claim be overstated. This number is significantly higher than that for Democrats and Republicans; there are no other differences of this magnitude in those claiming to have never consumed other brands. Given that many Republicans have come to view Bud Light negatively, it appears that some may be misremembering or engaging in motivated reasoning, and thus (inaccurately) reporting that they have *never* consumed Bud Light when in fact they have at some point in their life.

Free-response questions to the prompt "In your own words, why are you boycotting Bud Light" further suggest that this is the case amoong Republicans. One respondent wrote "I didn't drink that swill even before all this controversy, so me not buying it has nothing to do with being woke." Another noted "I just hate the taste of the beer," and another said "Been boycotting it way before it was a thing. the beer sucks and has sucked for a long time, not just because it's promoted by a transgender woman." One respondent simply wrote "I never drank Bud Light." While they may feel that they are boycotting, scholars would not describe these individuals as boycotting.

Figure 3.4: Recency of Bud Light consumption among "boycotters," by party

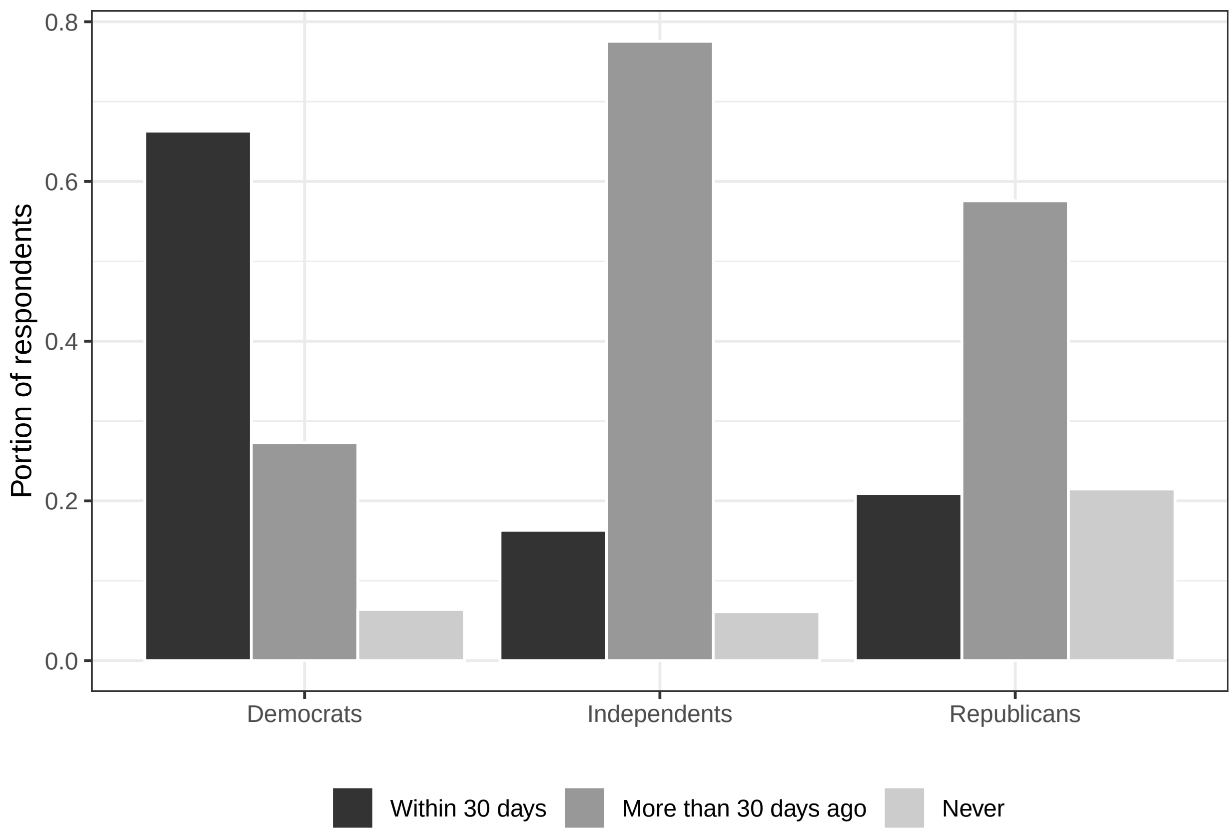

This figure takes only those who stated (post-treatment) that they were boycotting Bud Light and groups by their stated recency of Bud Light consumption, which was measured pre-treatment.

3.5 Conclusion

The conclusions of this paper are twofold. First, when brands are already politically polarized, partisan cues do not appear effective at reducing already-existing negative affect towards the brand. They can, however, further exacerbate negative perceptions of the brand among partisans who do not yet hold unfavorable views towards the brand.

A second conclusion suggests further limitations on the potential for partisan cues to shape political behavior. This study unobtrusively allowed us to compare pre-treatment and post-treatment statements about respondent behavior. When Democratic respondents were given a partisan cue suggesting that members of their own party were boycotting, they were far more likely to report that they too were boycotting—even when their own pre-treatment responses suggested that they had very recently consumed the brand they claimed to be boycotting.

These findings suggest important limitations to existing work on the framing effects of political cues on political consumerism. While political cues can induce political consumerism, this

Table 3.2: Likelihood of inconsistent boycott reporting

		'False' Boycotts	
	Democrats	**Independents**	**Republicans**
	Model 1	**Model 2**	**Model 3**
Republican Boycott	−0.016	−0.029	0.032
	(0.021)	(0.033)	(0.033)
Republican Support	0.006	−0.0005	0.017
	(0.021)	(0.030)	(0.033)
Democratic Boycott	0.043^{**}	0.033	0.007
	(0.021)	(0.030)	(0.033)
Democratic Support	−0.012	−0.015	−0.008
	(0.021)	(0.033)	(0.033)
Control	0.068^{***}	0.044^{**}	0.129^{***}
	(0.015)	(0.021)	(0.023)
N	1475	411	1114
R-squared	0.006	0.010	0.002

$^{***}p < .01$; $^{**}p < .05$; $^{*}p < .1$

may be a one-way street: cues can induce polarization but are not necessarily able to reverse already-existing polarization. They also suggest the need for caution in considering self-report data on political consumption behavior, as the respondents appear likely to misreport their own past behavior to match partisan behavioral norms.

Bibliography

Abadie, Alberto. 2005. "Semiparametric Difference-in-Differences Estimators." *The Review of Economic Studies* 72(1):1–19.

Abadie, Alberto, Alexis Diamond and Jens Hainmueller. 2010. "Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program." *Journal of the American Statistical Association* 105(490):493–505.

Abadie, Alberto, Alexis Diamond and Jens Hainmueller. 2015. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science* 59(2):495–510.

Abraham, Mabel and Vanessa Burbano. 2022. "Congruence Between Leadership Gender and Organizational Claims Affects the Gender Composition of the Applicant Pool: Field Experimental Evidence." *Organization Science* 33(1):393–413.

Agarwal, Sumit, Yupeng Lin, Michael Shen and Sirui Wu. 2023. "Banking Crisis Regulator.".

Ahler, Douglas J. and Gaurav Sood. 2018. "The Parties in Our Heads: Misperceptions about Party Composition and Their Consequences." *The Journal of Politics* 80(3):964–981.

Ahn, Jaehan, Rani Hoitash, Udi Hoitash and Eric Krause. 2023. "The Turnover, Retention, and Career Advancement of Female and Racial Minority Auditors: Evidence from Individual LinkedIn Data.".

Akerlof, George A. and Rachel E. Kranton. 2000. "Economics and Identity." *The Quarterly Journal of Economics* 115(3):715–753.

Allport, Gordon W. 1954. *The Nature of Prejudice.* Addison-wesley ed.

Allport, G.W. 1979. *The Nature of Prejudice.* Basic Books Addison-Wesley Publishing Company.

Ansolabehere, Stephen and Eitan Hersh. 2012. "Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate." *Political Analysis* 20(4):437–459.

Apfeld, Brendan, Emanuel Coman, John Gerring and Stephen Jessee. 2023. "Higher Education and Cultural Liberalism: Regression Discontinuity Evidence from Romania." *The Journal of Politics* 85(1):34–48.

Appels, Moritz. 2023. "CEO Sociopolitical Activism as a Signal of Authentic Leadership to Prospective Employees." *Journal of Management* 49(8):2727–2765.

Arif, Salman, Yeo Sang (Johnny) Yoon and Haiwen (Helen) Zhang. 2022. "The Information Content of Mandatory Human Capital Disclosures - Initial Evidence.".

Associated Press. 2016. "Associated Press Poll: December 2016.".

Avery, Dan. 2020. "LGBTQ Rights Fight Reignited 4 Years after N.C.'s 'bathroom Bill' Controversy." *NBC News* .

Ba, Bocar, Jacob Kaplan, Dean Knox, Mayya Komisarchik, Rachel Mariman, Jonathan Mummolo, Roman Rivera and Michelle Torres. 2022. "Who Are the Police? Descriptive Representation in the Coercive Arm of Government.".

Bai, Jushan. 2009. "Panel Data Models With Interactive Fixed Effects." *Econometrica* 77(4):1229–1279.

Baker, Andrew, David F. Larcker, Charles McClure, Durgesh Saraph and Edward Watts. 2022. "Diversity Washing." *SSRN Electronic Journal* .

Banda, Kevin K., Thomas M. Carsey and Serge Severenchuk. 2020. "Evidence of Conflict Extension in Partisans' Evaluations of People and Inanimate Objects." *American Politics Research* 48(2):275–285.

Bangerter, Adrian, Nicolas Roulin and Cornelius J. König. 2012. "Personnel Selection as a Signaling Game." *The Journal of Applied Psychology* 97(4):719–738.

Bansak, Kirk, Jens Hainmueller, Daniel J. Hopkins and Teppei Yamamoto. 2018. "The Number of Choice Tasks and Survey Satisficing in Conjoint Experiments." *Political Analysis* 26(1):112–119.

Barber, Benjamin and Daniel J. Blake. 2023. "My Kind of People: Political Polarization, Ideology, and Firm Location." *Strategic Management Journal* p. smj.3572.

Barber, Michael and John B. Holbein. 2022. "400 Million Voting Records Show Profound Racial and Geographic Disparities in Voter Turnout in the United States." *PLOS ONE* 17(6):e0268134.

Barbulescu, Roxana and Matthew Bidwell. 2013. "Do Women Choose Different Jobs from Men? Mechanisms of Application Segregation in the Market for Managerial Workers." *Organization Science* 24(3):737–756.

Barzuza, Michal, Quinn Curtis and David H. Webber. 2021. "The Millennial Corporation." *SSRN Electronic Journal* .

Becker, Gary. 1957. *The Economics of Discrimination.* University Of Chicago Press.

Bedendo, Mascia and Linus Siming. 2021. "To Advocate or Not to Advocate: Determinants and Financial Consequences of CEO Activism." *British Journal of Management* 32(4):1062–1081.

Bennett, W. Lance and Robert M. Entman, eds. 2000. *Mediated Politics: Communication in the Future of Democracy*. Communication, Society and Politics Cambridge: Cambridge University Press.

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

Besley, Timothy and Maitreesh Ghatak. 2005. "Competition and Incentives with Motivated Agents." *American Economic Review* 95(3):616–636.

Bhagwat, Yashoda, Nooshin L. Warren, Joshua T. Beck and George F. Watson. 2020. "Corporate Sociopolitical Activism and Firm Value." *Journal of Marketing* 84(5):1–21.

Billings, Stephen B, Noah Braun, Daniel B Jones and Ying Shi. 2022. "Disparate Racial Impacts of Shelby County v. Holder on Voter Turnout.".

Bishop, Bill. 2009. *The Big Sort: Why the Clustering of Like-Minded American Is Tearing Us Apart*. First edition ed. Boston: Mariner Books.

Black, Ines, Sharique Hasan and Rembrand Koning. 2024. "Hunting for Talent: Firm-driven Labor Market Search in the United States." *Strategic Management Journal* 45(3):429–462.

Blau, Francine D. and Lawrence M. Kahn. 2017. "The Gender Wage Gap: Extent, Trends, and Explanations." *Journal of Economic Literature* 55(3):789–865.

Bolton, Patrick, Tao Li, Enrichetta Ravina and Howard Rosenthal. 2020. "Investor Ideology." *Journal of Financial Economics* 137(2):320–352.

Bondi, Tommaso, Vanessa Burbano and Fabrizio Dell'Acqua. 2022. "When to Talk Politics in Business: Theory and Experimental Evidence of Stakeholder Responses to CEO Political Activism.".

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

Bonica, Adam, Adam S. Chilton and Maya Sen. 2016. "The Political Ideologies of American Lawyers." *Journal of Legal Analysis* 8(2):277–335.

Bonica, Adam and Jacob M Grumbach. 2022. "Old Money: Campaign Finance and Gerontocracy in the United States.".

Bonica, Adam, Jacob M. Grumbach, Charlotte Hill and Hakeem Jefferson. 2021. "All-Mail Voting in Colorado Increases Turnout and Reduces Turnout Inequality." *Electoral Studies* 72:102363.

Branicki, Layla, Stephen Brammer, Alison Pullen and Carl Rhodes. 2021. "The Morality of "New" CEO Activism." *Journal of Business Ethics* 170(2):269–285.

Briscoe, Forrest, M. K. Chin and Donald C. Hambrick. 2014. "CEO Ideology as an Element of the Corporate Opportunity Structure for Social Activists." *Academy of Management Journal* 57(6):1786–1809.

Briscoe, Forrest and Sean Safford. 2008. "The Nixon-in-China Effect: Activism, Imitation, and the Institutionalization of Contentious Practices." *Administrative Science Quarterly* 53(3):460–491.

Broockman, David E. and Daniel M. Butler. 2017. "The Causal Effects of Elite Position-Taking on Voter Attitudes: Field Experiments with Elite Communication." *American Journal of Political Science* 61(1):208–221.

Broockman, David, Joshua Kalla and Sean Westwood. 2020. "Does Affective Polarization Undermine Democratic Norms or Accountability? Maybe Not.".

Brown, Jacob R. 2023. "Partisan Conversion Through Neighborhood Influence: How Voters Adopt the Partisanship of Their Neighbors.".

Brown, Jacob R., Enrico Cantoni, Ryan D. Enos, Vincent Pons and Emilie Sartre. 2023. "The Increase in Partisan Segregation in the United States.".

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

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

Burbano, Vanessa C. 2021. "The Demotivating Effects of Communicating a Social-Political Stance: Field Experimental Evidence from an Online Labor Market Platform." *Management Science* 67(2):1004–1025.

Bureau of Labor Statistics. 2023. American Time Use Survey - 2022 Results. Technical Report USDL-23-1364 U.S. Department of Labor Washington, D.C.: .

Cai, Wei, Aiyesha Dey, Jillian Grennan, Joseph Pacelli and Lin Qiu. 2022. "Do Diverse Directors Influence DEI Outcomes?".

Cai, Wei, Yue Chen, Shivaram Rajgopal and Li Yang. 2022. "Diversity Targets.".

Cain, Áine, Marguerite Ward, Hoff Madison and Tanya Dua. 2022. "The Country's Biggest Employers, Including Walmart and Amazon, Should 'say Goodbye to Attracting Top Female Talent' in Abortion 'Trigger Law' States." https://www.businessinsider.com/employers-in-abortion-ban-states-face-retention-crisis-2022-5.

Cappelli, Peter. 2019. "Your Approach to Hiring Is All Wrong." *Harvard Business Review* .

Cappelli, Peter H. 2015. "Skill Gaps, Skill Shortages, and Skill Mismatches: Evidence and Arguments for the United States." *ILR Review* 68(2):251–290.

Carpenter, Jeffrey and Erick Gong. 2015. "Motivating Agents: How Much Does the Mission Matter?" *Journal of Labor Economics* 34(1):211–236.

Chatman, Jennifera. 1989. "Improving Interactional Organizational Research: A Model of Person-OrganizaFit Tion." *Academy of Management Review* 14(3):333–349.

Chatterji, Aaron K. and Michael W. Toffel. 2018. "The New CEO Activists." *Harvard Business Review* .

Chatterji, Aaron K. and Michael W. Toffel. 2019. "Assessing the Impact of CEO Activism." *Organization & Environment* 32(2):159–185.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, Emmanuel Saez and Nicholas Turner. 2014. "Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility." *The American Economic Review* 104(5):141–147.

Coff, Russell and David Kryscynski. 2011. "Invited Editorial: Drilling for Micro-Foundations of Human Capital–Based Competitive Advantages." *Journal of Management* 37(5):1429–1443.

Colonnelli, Emanuele, Valdemar Pinho Neto and Edoardo Teso. 2022. "Politics At Work.".

Connelly, Brian L., S. Trevis Certo, R. Duane Ireland and Christopher R. Reutzel. 2011. "Signaling Theory: A Review and Assessment." *Journal of Management* 37(1):39–67.

Corno, Lucia, Eliana La Ferrara and Justine Burns. 2022. "Interaction, Stereotypes, and Performance: Evidence from South Africa." *American Economic Review* 112(12):3848–3875.

Cowgill, Bo and Patryk Perkowski. 2020. "Delegation in Hiring: Evidence from a Two-Sided Audit.".

Culpepper, Pepper D. 2010. *Quiet Politics and Business Power: Corporate Control in Europe and Japan.* Cambridge Studies in Comparative Politics Cambridge: Cambridge University Press.

Dalesio, Emery P. and Jonathan Drew. 2017. "AP Exclusive: 'Bathroom Bill' to Cost North Carolina $3.76B." *AP News* .

de Chaisemartin, Clément and Xavier D'Haultfœuille. 2023. "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey." *The Econometrics Journal* 26(3):C1–C30.

Dodd, Melissa D and Dustin W. Supa. 2014. "Conceptualizing and Measuring "Corporate Social Advocacy" Communication: Examining the Impact on Corporate Financial." *Public Relations Journal* 8(3):23.

Downey, Mitch and Jinci Liu. 2023. "Political Preferences and Migration of College Educated Workers.".

Druckman, James N and Matthew S Levendusky. 2019. "What Do We Measure When We Measure Affective Polarization?" *Public Opinion Quarterly* 83(1):114–122.

Egami, Naoki and Soichiro Yamauchi. 2023. "Using Multiple Pretreatment Periods to Improve Difference-in-Differences and Staggered Adoption Designs." *Political Analysis* 31(2):195–212.

Elder, Elizabeth Mitchell and Neil A. O'Brian. 2022. "Social Groups as the Source of Political Belief Systems: Fresh Evidence on an Old Theory." *American Political Science Review* pp. 1–18.

Enamorado, Ted. 2021. A Primer on Probabilistic Record Linkage. In *Handbook of Computational Social Science, Volume 2.* Routledge.

Endres, Kyle and Costas Panagopoulos. 2017. "Boycotts, Buycotts, and Political Consumerism in America." *Research & Politics* 4(4):2053168017738632.

Endres, Kyle, Costas Panagopoulos and Donald P. Green. 2021. "Elite Messaging and Partisan Consumerism: An Evaluation of President Trump's Tweets and Polarization of Corporate Brand Images." *Political Research Quarterly* 74(4):834–851.

Fadhel, Adel, Katie Panella, Ethan Rouen and George Serafeim. 2021. "Accounting for Employment Impact at Scale.".

Fausset, Richard. 2014. "North Carolina, in Political Flux, Battles for Its Identity." *The New York Times* .

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

Florida, Richard. 2014. *The Rise of the Creative Class–Revisited: Revised and Expanded.* Basic Books.

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

Freeman, R. Edward. 1984. *Strategic Management: A Stakeholder Approach.* Pitman.

Gao, Janet, Yufeng Wu and Wang Zhang. 2023. "Decentralized Banking in Mortgage Market: Evidence from Branch Manager's Past Experience.".

Gelles, David. 2021. "Delta and Coca-Cola Reverse Course on Georgia Voting Law, Stating 'Crystal Clear' Opposition." *The New York Times* .

Gift, Karen and Thomas Gift. 2015. "Does Politics Influence Hiring? Evidence from a Randomized Experiment." *Political Behavior* 37(3):653–675.

Gobillon, Laurent and Thierry Magnac. 2016. "Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls." *The Review of Economics and Statistics* 98(3):535–551.

Goldberg, Abbie E. 2023. Impact of HB 1557 (Florida's Don't Say Gay Bill) on LGBTQ+ Parents in Florida. Technical report UCLA School of Law Williams Institute.

Goldin, Claudia Dale. 2021. *Career and Family: Women's Century-Long Journey toward Equity.* Princeton, New Jersey: Princeton University Press.

Goldin, Claudia, Sari Pekkala Kerr, Claudia Olivetti and Erling Barth. 2017. "The Expanding Gender Earnings Gap: Evidence from the LEHD-2000 Census." *American Economic Review* 107(5):110–114.

Goldman, Eitan, Nandini Gupta and Ryan D. Israelsen. 2021. "Political Polarization in Financial News.".

Grumbach, Jacob M. and Alexander Sahn. 2020. "Race and Representation in Campaign Finance." *American Political Science Review* 114(1):206–221.

Gubler, Joshua R. and Joel Sawat Selway. 2012. "Horizontal Inequality, Crosscutting Cleavages, and Civil War." *Journal of Conflict Resolution* 56(2):206–232.

Gupta, Abhinav and Forrest Briscoe. 2020. "Organizational Political Ideology and Corporate Openness to Social Activism." *Administrative Science Quarterly* 65(2):524–563.

Hainmueller, Jens, Daniel J. Hopkins and Teppei Yamamoto. 2014. "Causal Inference in Conjoint Analysis: Understanding Multidimensional Choices via Stated Preference Experiments." *Political Analysis* 22(1):1–30.

Hambrick, Donald C. and Adam J. Wowak. 2021. "CEO Sociopolitical Activism: A Stakeholder Alignment Model." *Academy of Management Review* 46(1):33–59.

Hersh, Eitan D. 2015. *Hacking the Electorate: How Campaigns Perceive Voters.* New York, NY: Cambridge University Press.

Hersh, Eitan D. and Matthew N. Goldenberg. 2016. "Democratic and Republican Physicians Provide Different Care on Politicized Health Issues." *Proceedings of the National Academy of Sciences* 113(42):11811–11816.

Hersh, Eitan and Sarang Shah. 2023. "The Partisan Realignment of American Business.".

Hersh, Eitan and Yair Ghitza. 2018. "Mixed Partisan Households and Electoral Participation in the United States." *PLOS ONE* 13(10):e0203997.

Hertel-Fernandez, Alexander. 2018. *Politics at Work: How Companies Turn Their Workers into Lobbyists.* 1st edition ed. New York, NY, United States of America: Oxford University Press.

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

Hou, Young and Christopher W. Poliquin. 2023. "The Effects of CEO Activism: Partisan Consumer Behavior and Its Duration." *Strategic Management Journal* 44(3):672–703.

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

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

Human Rights Campaign. 2016. "More Than 100 Major CEOs & Business Leaders Urge North Carolina To Repeal Anti-LGBT Law." https://www.hrc.org/press-releases/more-than-100-major-ceos-business-leaders-demand-north-carolina-repeal-radi.

Hurst, Reuben. 2023. "Countervailing Claims: Pro-Diversity Responses to Stigma by Association Following the Unite the Right Rally." *Administrative Science Quarterly* 68(4):1094–1132.

Hurst, Reuben, Saerom (Ronnie) Lee and Justin Frake. 2024. "The Effect of Flatter Hierarchy on Applicant Pool Gender Diversity: Evidence from Experiments." *Strategic Management Journal* n/a(n/a).

Hyatt, Henry, Erika McEntarfer, Kevin McKinney, Stephen Tibbets, Lars Vilhuber, Douglas Walton, Joyce Key Hahn and Hubert Janicki. 2017. "Job-to-Job Flows: New Statistics on Worker Reallocation and Job Turnover." *US Census Bureau* .

Igielnik, Ruth, Scott Keeter, Courtney Kennedy and Bradley Spahn. 2018. "Commercial Voter Files and the Study of U.S. Politics.".

Imai, Kosuke, Santiago Olivella and Evan T. R. Rosenman. 2022. "Addressing Census Data Problems in Race Imputation via Fully Bayesian Improved Surname Geocoding and Name Supplements." *Science Advances* 8(49):eadc9824.

Iyengar, Shanto and Sean J. Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." *American Journal of Political Science* 59(3):690–707.

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

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22(1):129–146.

Jensen, Nathan M. 2017*a*. "The Effect of Economic Development Incentives and Clawback Provisions on Job Creation: A Pre-Registered Evaluation of Maryland and Virginia Programs." *Research & Politics* 4(2):2053168017713646.

Jensen, Nathan M. 2017*b*. "Job Creation and Firm-Specific Location Incentives." *Journal of Public Policy* 37(1):85–112.

Jensen, Nathan M. and Calvin Thrall. 2021. "Who's Afraid of Sunlight? Explaining Opposition to Transparency in Economic Development." *Business and Politics* 23(4):474–491.

Jensen, Nathan M. and Edmund J. Malesky. 2018. *Incentives to Pander: How Politicians Use Corporate Welfare for Political Gain.* Cambridge, United Kingdom ; New York, NY: Cambridge University Press.

Jones, David A. 2013. "The Polarizing Effect of a Partisan Workplace." *PS: Political Science & Politics* 46(1):67–73.

Kahn-Lang, Ariella and Kevin Lang. 2020. "The Promise and Pitfalls of Differences-in-Differences: Reflections on *16 and Pregnant* and Other Applications." *Journal of Business & Economic Statistics* 38(3):613–620.

Kam, Cindy D. and Maggie Deichert. 2020. "Boycotting, Buycotting, and the Psychology of Political Consumerism." *The Journal of Politics* 82(1):72–88.

King, Brayden. 2011. "The Tactical Disruptiveness of Social Movements: Sources of Market and Mediated Disruption in Corporate Boycotts." *Social Problems* 58(4):491–517.

Klofstad, Casey A., Rose McDermott and Peter K. Hatemi. 2013. "The Dating Preferences of Liberals and Conservatives." *Political Behavior* 35(3):519–538.

Knol, Mirjam J., Wiebe R. Pestman and Diederick E. Grobbee. 2011. "The (Mis)Use of Overlap of Confidence Intervals to Assess Effect Modification." *European Journal of Epidemiology* 26(4):253–254.

Korschun, Daniel, Kelly D. Martin and Gautham Vadakkepatt. 2020. "Marketing's Role in Understanding Political Activity." *Journal of Public Policy & Marketing* 39(4):378–387.

Ladika, Susan. 2023. "Many Workers Leaving States That Pass Anti-LGBTQ Laws." https://www.shrm.org/topics-tools/news/inclusion-equity-diversity/many-workers-leaving-states-pass-anti-lgbtq-laws.

Larcker, David F, Stephen A Miles and Brian Tayan. 2018. "The Double-Edged Sword of CEO Activism.".

Leeper, Thomas J., Sara B. Hobolt and James Tilley. 2020. "Measuring Subgroup Preferences in Conjoint Experiments." *Political Analysis* 28(2):207–221.

Lelkes, Yphtach and Sean J. Westwood. 2017. "The Limits of Partisan Prejudice." *The Journal of Politics* 79(2):485–501.

Lenz, Gabriel S. 2009. "Learning and Opinion Change, Not Priming: Reconsidering the Priming Hypothesis." *American Journal of Political Science* 53(4):821–837.

Lenz, Gabriel S. 2012. *Follow the Leader?: How Voters Respond to Politicians' Policies and Performance.* Chicago Studies in American Politics Chicago, IL: University of Chicago Press.

Leung, Ming, Chuchu Liang, Ben Lourie and Chenqi Zhu. 2023. "Effect of Corporate Environment Social and Governance Reputation on Employee Turnover." *SSRN Electronic Journal* .

Li, Qin, Ben Lourie, Alexander Nekrasov and Terry Shevlin. 2022. "Employee Turnover and Firm Performance: Large-Sample Archival Evidence." *Management Science* 68(8):5667–5683.

Li, Zhao. 2018. "How Internal Constraints Shape Interest Group Activities: Evidence from Access-Seeking PACs." *American Political Science Review* 112(4):792–808.

Li, Zhao and Richard W. Disalvo. 2022. "Can Stakeholders Mobilize Businesses for the Protection of Democracy? Evidence from the U.S. Capitol Insurrection." *American Political Science Review* pp. 1–7.

Liaukonytė, Jūra, Anna Tuchman and Xinrong Zhu. 2022. "Frontiers: Spilling the Beans on Political Consumerism: Do Social Media Boycotts and Buycotts Translate to Real Sales Impact?" *Marketing Science* p. mksc.2022.1386.

Liaukonytė, Jūra, Anna Tuchman and Xinrong Zhu. 2023. "Rejoinder: Spilling More Beans on Political Consumerism: It's More of the Same Tune." *Marketing Science* 42(1):32–36.

Lindblom, Charles E. 1982. "The Market as Prison." *The Journal of Politics* 44(2):324–336.

Lipset, Seymour Martin. 1981. *Political Man: The Social Bases of Politics.* Expanded ed ed. Baltimore: Johns Hopkins University Press.

Maks-Solomon, Cory. 2020. ""Liberal Elite" CEOs and Corporate Social Activism.".

Maks-Solomon, Cory and Josiah Mark Drewry. 2021. "Why Do Corporations Engage in LGBT Rights Activism? LGBT Employee Groups as Internal Pressure Groups." *Business and Politics* 23(1):124–152.

Malina, Gabrielle and Eitan Hersh. 2021. "The Politics of 130,000 American Religious Leaders: A New Methodological Approach." *Journal for the Scientific Study of Religion* 60(4):709–725.

Mallory, Christy and Brad Sears. 2016. The Fiscal Impact of North Carolina's HB2. Technical report UCLA School of Law Williams Institute.

Margolis, Michele F. 2018. *From Politics to the Pews: How Partisanship and the Political Environment Shape Religious Identity.* University of Chicago Press.

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

Mason, Lilliana. 2015. ""I Disrespectfully Agree": The Differential Effects of Partisan Sorting on Social and Issue Polarization." *American Journal of Political Science* 59(1):128–145.

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

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

Melloni, Gaia, Andrea Patacconi and Nick Vikander. 2023. "Cashing in on the Culture Wars? CEO Activism, Wokewashing, and Firm Value." *Strategic Management Journal* 44(13):3098–3121.

Micheletti, M., ed. 2003. *Political Virtue and Shopping: Individuals, Consumerism, and Collective Action.* 2003rd edition ed. New York: Palgrave Macmillan.

Mohliver, Aharon, Donal Crilly and Aseem Kaul. 2023. "Corporate Social Counterpositioning: How Attributes of Social Issues Influence Competitive Response." *Strategic Management Journal* 44(5):1199–1217.

Morrill, Calvin, Mayer N. Zald and Hayagreeva Rao. 2003. "Covert Political Conflict in Organizations: Challenges from Below." *Annual Review of Sociology* 29(Volume 29, 2003):391–415.

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

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

Nalick, Michael, Matthew Josefy, Asghar Zardkoohi and Leonard Bierman. 2016. "Corporate Sociopolitical Involvement: A Reflection of Whose Preferences?" *Academy of Management Perspectives* 30(4):384–403.

Newman, Benjamin J. and Brandon L. Bartels. 2011. "Politics at the Checkout Line: Explaining Political Consumerism in the United States." *Political Research Quarterly* 64(4):803–817.

Noah, Timothy. 2023. "The Red State Brain Drain Isn't Coming. It's Happening Right Now." *The New Republic* .

Nordlinger, Eric A. 1972. *Conflict Regulation in Divided Societies.* Cambridge, Mass.: Center for International Affairs, Harvard University.

Olney, William W. and Owen Thompson. 2024. "The Determinants of Declining Internal Migration.".

Orgera, Kendal, Hasan Mahmood and Atul Grover. 2023. "Training Location Preferences of U.S. Medical School Graduates Post Dobbs v. Jackson Women's Health." https://www.aamcresearchinstitute.org/our-work/data-snapshot/training-location-preferences-us-medical-school-graduates-post-dobbs-v-jackson-women-s-health.

Pacelli, Joseph, Tianshuo (Terrence) Shi and Yuan Zou. 2023. "Communicating Corporate Culture in Labor Markets: Evidence from Job Postings.".

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

Parker, Kim, Juliana Horowitz, Anna Browns, Richard Fry, D'Vera Cohn and Ruth Igielnik. 2018. "What Unites and Divides Urban, Suburban, and Rural Communities.".

Pasirayi, Simbarashe, Patrick B. Fennell and Kayla B. Follmer. 2023. "#Activism: Investor Reactions to Corporate Sociopolitical Activism." *Business & Society* 62(4):704–744.

Raeburn, Nicole C. 2004. *Changing Corporate America from inside out: Lesbian and Gay Workplace Rights.* Number v. 20 *in* "Social Movements, Protest, and Contention" Minneapolis: University of Minnesota Press.

Raffiee, Joseph, Florenta Teodoridis and Daniel Fehder. 2023. "Partisan Patent Examiners? Exploring the Link between the Political Ideology of Patent Examiners and Patent Office Outcomes." *Research Policy* 52(9):104853.

Rho, Eunice Hyon Min. 2017. "After the Fake Repeal of HB2, We Seek Transparency in North Carolina | ACLU.".

Rosenman, Evan T. R., Santiago Olivella and Kosuke Imai. 2023. "Race and Ethnicity Data for First, Middle, and Surnames." *Scientific Data* 10(1):299.

Roth, Jonathan and Pedro H. C. Sant'Anna. 2023. "When Is Parallel Trends Sensitive to Functional Form?" *Econometrica* 91(2):737–747.

Roth, Jonathan, Pedro H.C. Sant'Anna, Alyssa Bilinski and John Poe. 2023. "What's Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature." *Journal of Econometrics* 235(2):2218–2244.

Rothacker, Rick, Ely Portillo and Katherine Peralta. 2016. "PayPal Withdraws Plans for Charlotte Expansion over HB2." *Charlotte Observer* .

Roy, A. D. 1951. "Some Thoughts on the Distribution of Earnings." *Oxford Economic Papers* 3(2):135–146.

Sax, Christoph and Dirk Eddelbuettel. 2018. "Seasonal Adjustment by X-13ARIMA-SEATS in R." *Journal of Statistical Software* 87(11).

Sax, Christoph and Dirk Eddelbuettel. 2022. "Seasonal: R Interface to X-13-ARIMA-SEATS.".

Schaffner, Brian, Stephen Ansolabehere and Marissa Shih. 2023. "Cooperative Election Study Common Content, 2022.".

Schneider, Benjamin, Harold W. Goldstiein and D. Brent Smith. 1995. "THE ASA FRAMEWORK: AN UPDATE." *Personnel Psychology* 48(4):747–773.

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

Shanor, Amanda, Mary-Hunter McDonnell and Timothy Werner. 2022. "Corporate Political Power: The Politics of Reputation & Traceability." *Emory Law Journal* 71(2):65.

Smith, Adam. 1776. *An Inquiry into the Nature and Causes of the Wealth of Nations.*

Smith, Mark A. 2000. *American Business and Political Power: Public Opinion, Elections, and Democracy.* Chicago: University of Chicago Press.

Soule, Sarah Anne. 2009. *Contention and Corporate Social Responsibility.* Cambridge Studies in Contentious Politics New York: Cambridge University Press.

Spenkuch, Jörg L., Edoardo Teso and Guo Xu. 2023. "Ideology and Performance in Public Organizations." *Econometrica* 91(4):1171–1203.

Stolle, Dietlind and Michele Micheletti. 2013. *Political Consumerism: Global Responsibility in Action.* Cambridge: Cambridge University Press.

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

Tajfel, Henri and John Turner. 1979. *An Integrative Theroy of Intergroup Conflict.* Monterey, CA: Brooks Cole pp. 33–47.

The National Resident Matching Program. 2023. "NRMP® Releases the 2023 Main Residency Match Results and Data Report, the Most Trusted Data Resource for the Main Residency Match®." https://www.nrmp.org/about/news/2023/05/nrmp-releases-the-2023-main-residency-match-results-and-data-report-the-most-trusted-data-resource-for-the-main-residency-match/.

Tiebout, Charles M. 1956. "A Pure Theory of Local Expenditures." *Journal of Political Economy* 64(5):416–424.

US Census Bureau. 2017. X13-ARIMA-SEATS Reference Manual. Technical report Time Series Research Staff, Center for Statistical Research and Methodology, US Census Bureau.

Vogel, David. 1978. "Why Businessmen Distrust Their State: The Political Consciousness of American Corporate Executives." *British Journal of Political Science* 8(1):45–78.

Vogel, David. 2006. *The Market for Virtue: The Potential and Limits of Corporate Social Responsibility.* 2nd edition ed. Washington, DC: Brookings Institution Press.

Werner, Timothy. 2017. "Investor Reaction to Covert Corporate Political Activity." *Strategic Management Journal* 38(12):2424–2443.

Werner, Timothy. 2018. *Public Forces and Private Politics in American Big Business.*

Wilson, Valerie, Ethan Miller and Melat Kassa. 2021. Racial Representation in Professional Occupations: By the Numbers. Technical report Economic Policy Institute.

Wojcieszak, Magdalena and Benjamin R. Warner. 2020. "Can Interparty Contact Reduce Affective Polarization? A Systematic Test of Different Forms of Intergroup Contact." *Political Communication* 37(6):789–811.

Wowak, Adam J., John R. Busenbark and Donald C. Hambrick. 2022. "How Do Employees React When Their CEO Speaks Out? Intra- and Extra-Firm Implications of CEO Sociopolitical Activism." *Administrative Science Quarterly* 67(2):553–593.

Xu, Yiqing. 2017. "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models." *Political Analysis* 25(1):57–76.

Yang, Bo and Nan Jia. 2023. "Controversial Corporate Political Activities and Shareholder Activism: A Natural Experiment of the U.S. Capitol Riot.".

Zaller, John R. 1992. *The Nature and Origins of Mass Opinion.* Cambridge Studies in Public Opinion and Political Psychology Cambridge: Cambridge University Press.

Zingher, Joshua N. 2022. "Diploma Divide: Educational Attainment and the Realignment of the American Electorate." *Political Research Quarterly* .

Appendix A

Appendices for Chapter 1

Figure A.1: Voter file match rate, by age

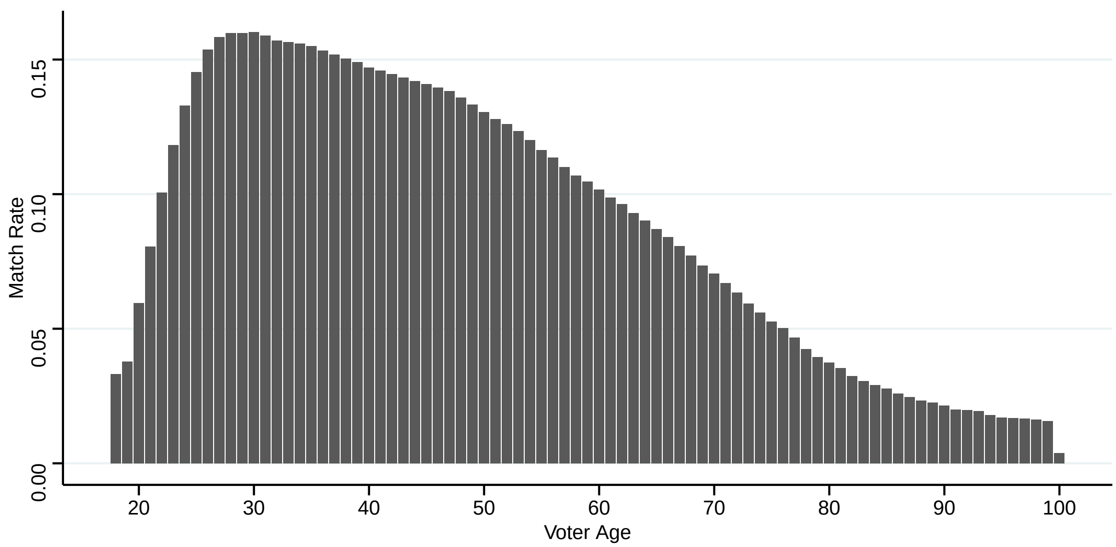

This figure depicts the portion of voter file records which we are able to match using the exact matching procedure. Among registered voters in their 30s, we are able to match about 15% with an corresponding employment record from Revelio. This is unsurprising given that these voters are most likely to be in employment. They may also be more likely to appear on online job postings and social media, such as LinkedIn. We are less able to match the very young, who are less likely to be registered to vote, and the very old, who are less likely to be in employment.

Figure A.2: Number of coworkers per position

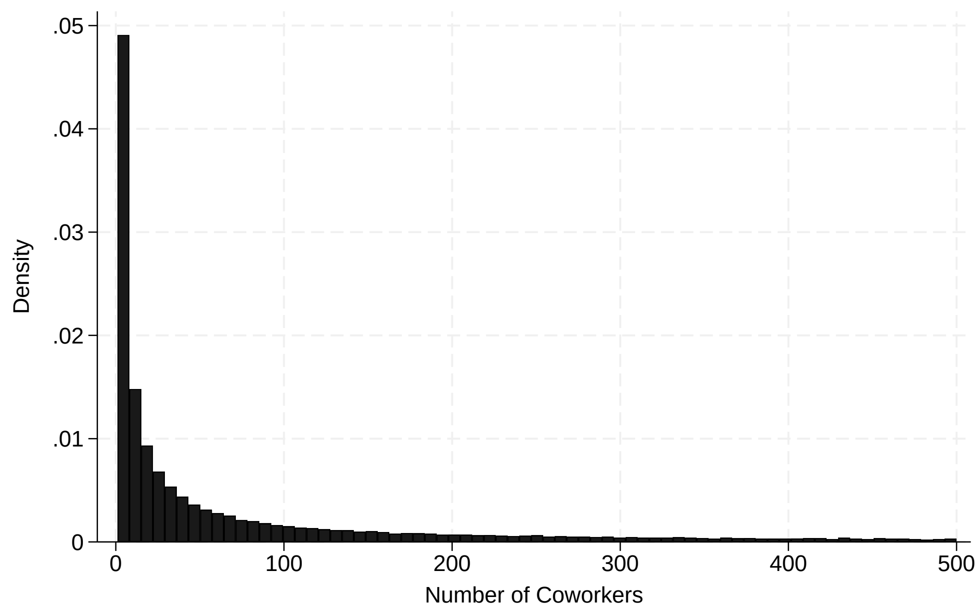

For each employment position in the dataset, this figure depicts the number of co-workers at the same workplace (i.e., at the same employer and MSA). Note that positions are not necessarily individual workers; a worker may have multiple positions as her or she changes jobs or concurrently if he or she works multiple jobs. This figure excludes positions where the number of co-workers is greater than 500 (18% of all positions).

Figure A.3: Number of positions per workplace

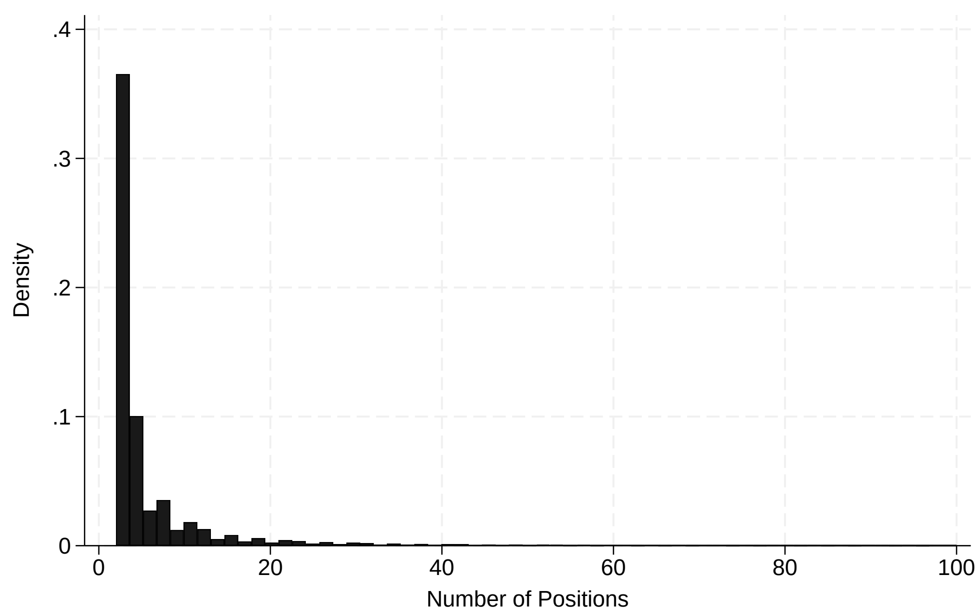

Units of analysis are workplaces (employer-MSA dyads). This figure excludes employer-MSAs for which the number of positions is greater than 100, which constitute 1% of employer-MSAs.

Figure A.4: Number of positions per year

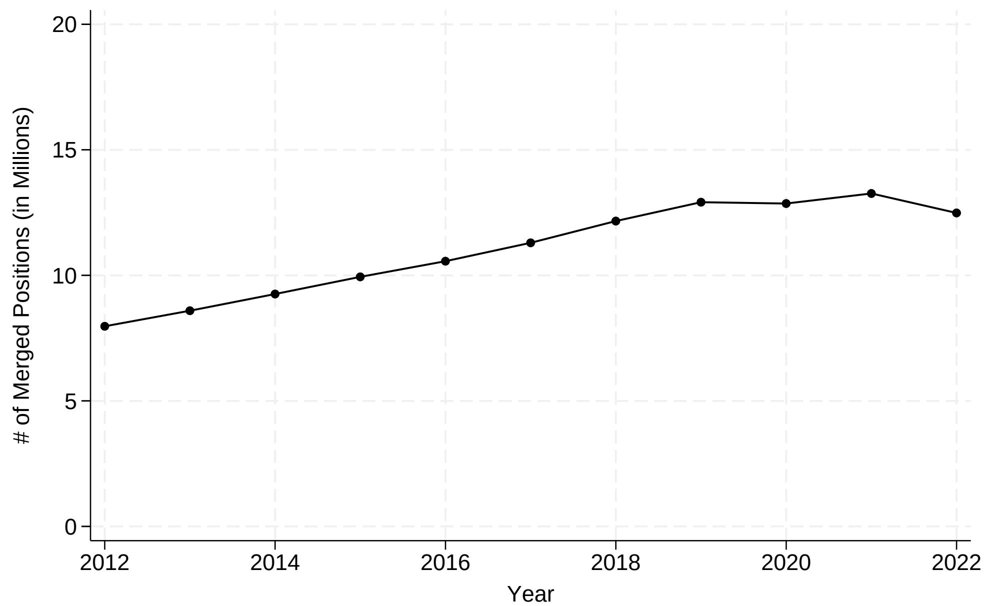

Each point indicates the number of positions for a given year. The slight dip in 2022 reflects the fact that our data end in August 2022.

Figure A.5: Nationwide distribution of coworkers that are Democrats and Republicans, overall and by focal worker partisanship.

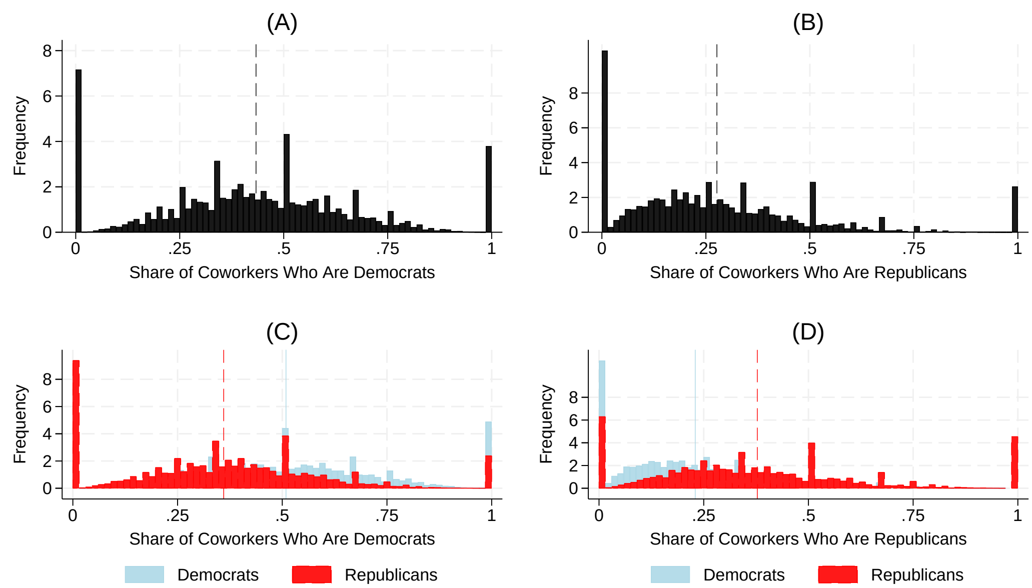

Sub-figures (A) and (B) show the overall distribution in the share of coworkers, by position, that are Democrats and Republicans, respectively. Sub-figures (C) and (D) present these distributions separately for Democrats and Republicans. All plots are histograms. Vertical lines are averages. Figure 1.1 in the main text presents these patterns as kernel density plots rather than histograms.

Figure A.6: Most Democratic and Republican occupations and industries

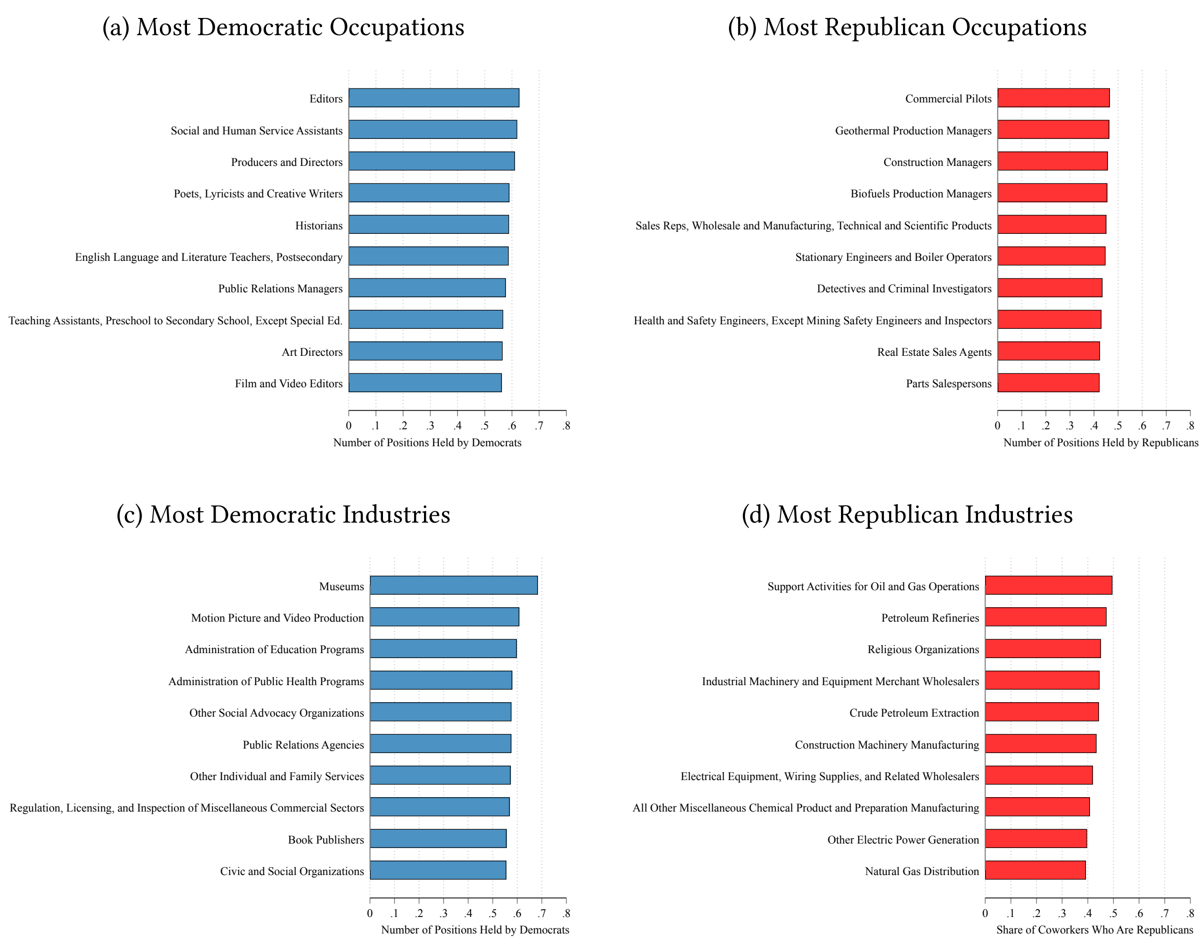

Figure A.7: Gender sorting in the labor market

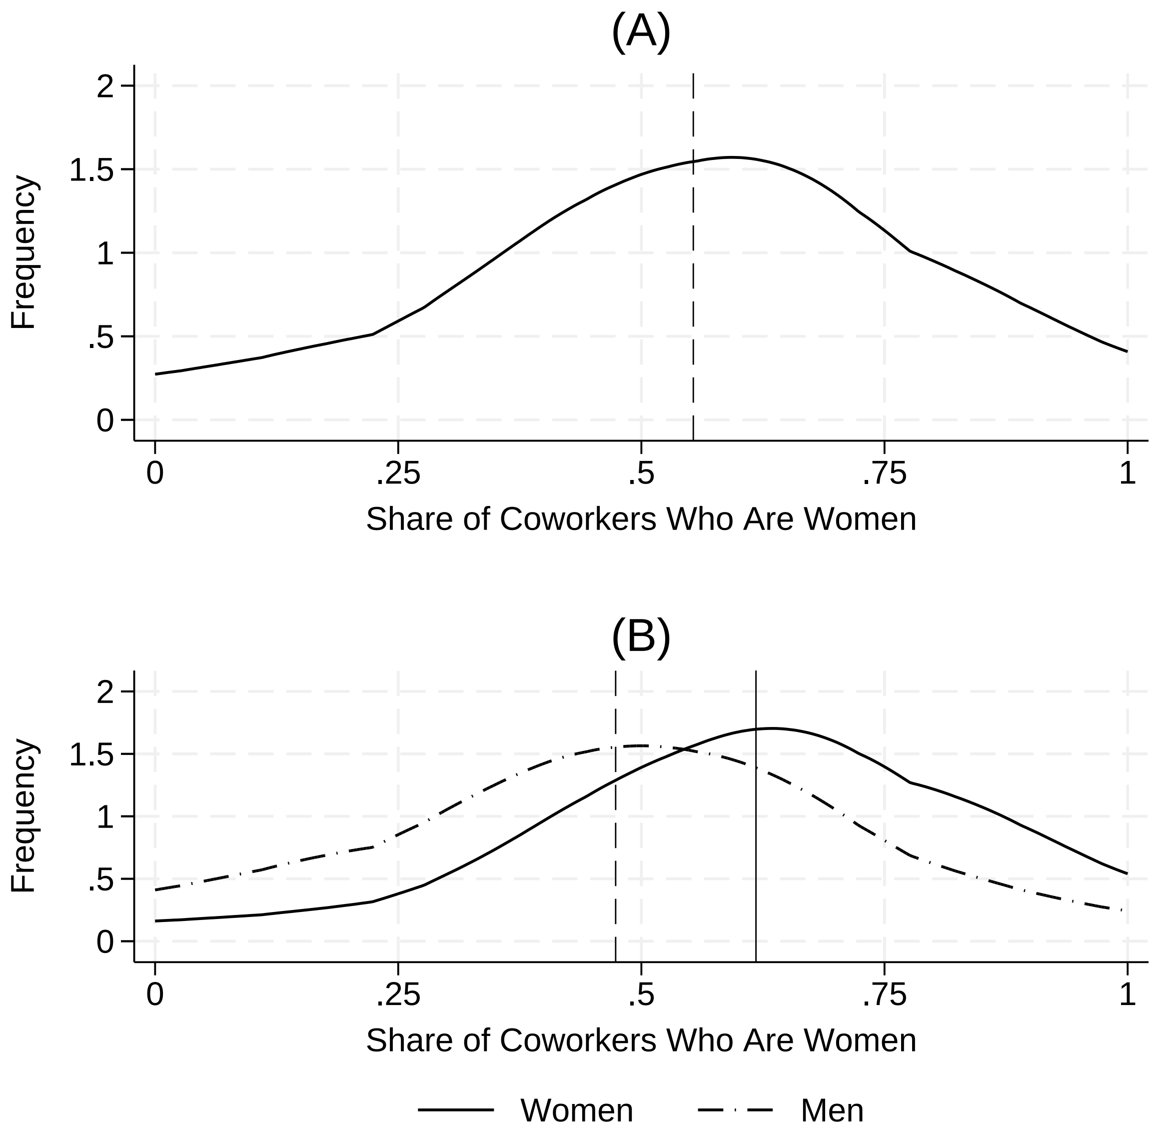

Sub-figure (A) shows the overall distribution in the share of coworkers, by position, that are women. Sub-figures (B) present this distribution separately for women and men. Compared to men, women tend to work in employer-MSAs in which a greater share of co-workers are women. Both plots are kernel density plots with a bandwidth of 0.10. Vertical lines are averages.

Figure A.8: Racial sorting in the labor market

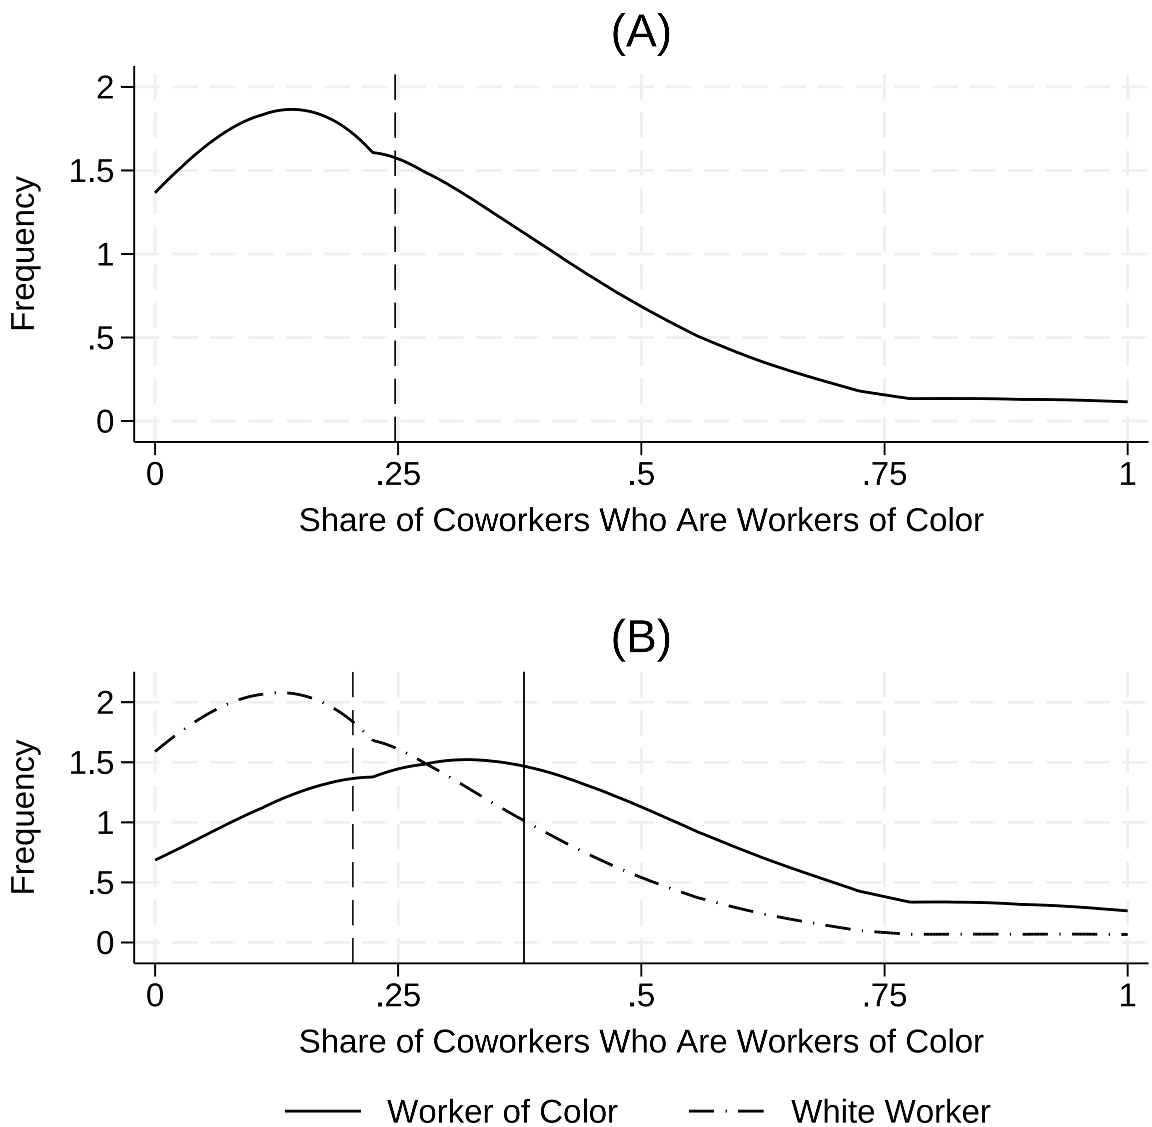

Sub-figure (A) shows the overall distribution in the share of coworkers, by position, that are workers of color. Sub-figures (B) present this distribution separately for workers of color and white workers. Compared to white workers, workers of color tend to work in employer-MSAs in which a greater share of co-workers are workers of color. Both plots are kernel density plots with a bandwidth of 0.10. Vertical lines are averages.

Figure A.9: Specification curve without fixed effects and Democrat as reference partisanship

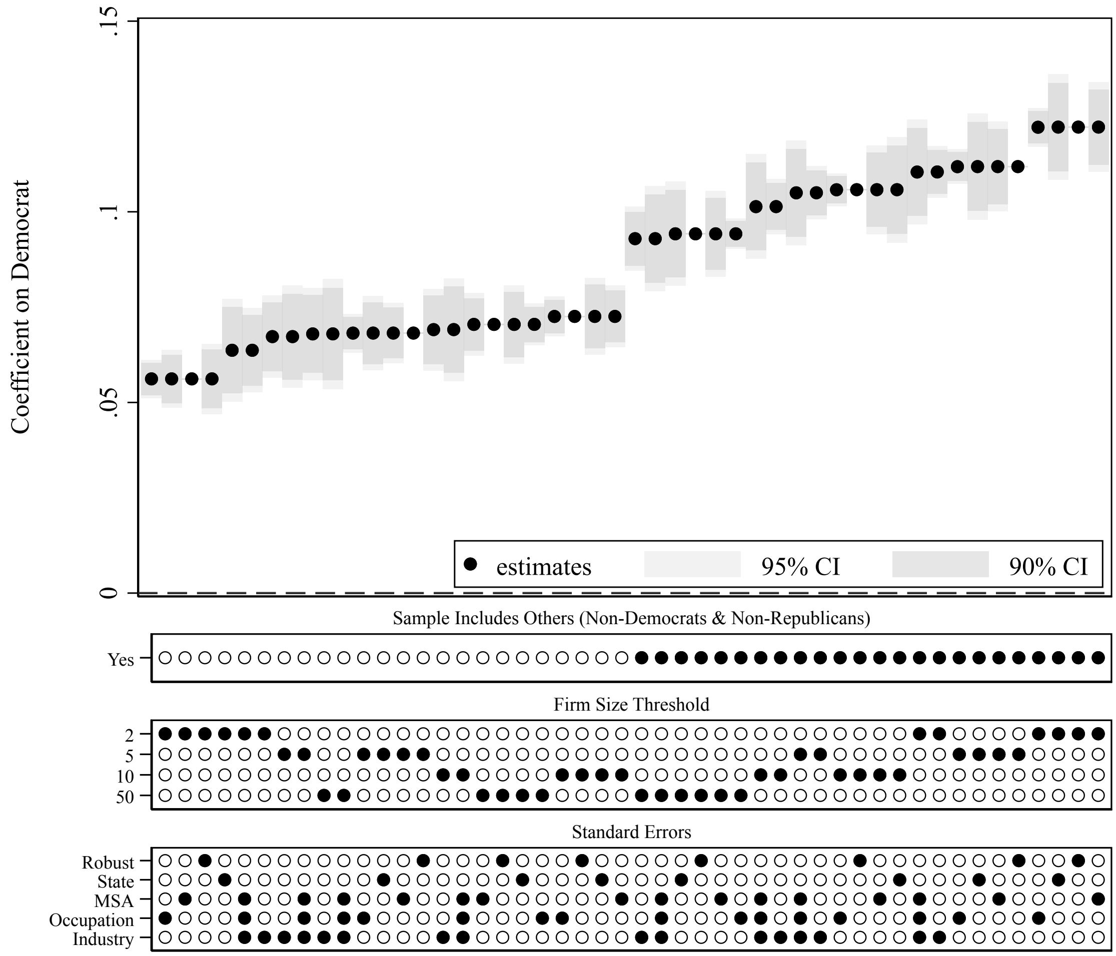

Points are from estimates of different versions of Equation 1.1 and vary in terms of whether they include workers who do not identify as either Democrat or Republican, the inclusion criteria in terms of number of co-workers, and the level at which we cluster standard errors.

Figure A.10: Specification curve without fixed effects and Republican as reference partisanship

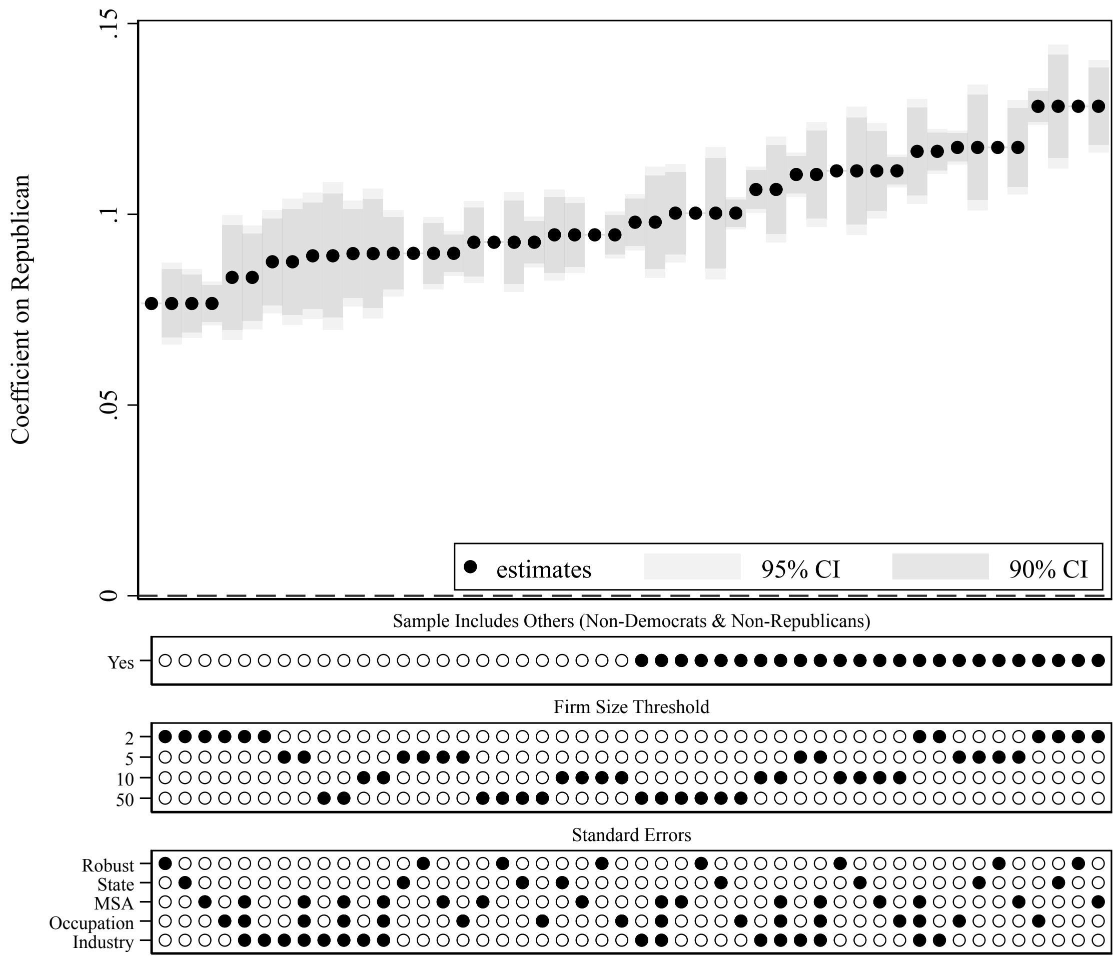

Points are from estimates of different versions of Equation 1.1 except in all cases, the outcome variable is the share of Republican coworkers and the explanatory variable is whether or not the focal worker is a Republican. The specifications vary in terms of whether they include workers who do not identify as either Democrat or Republican, the inclusion criteria in terms of number of co-workers, and the level at which we cluster standard errors.

Figure A.11: Specification curve with full fixed effects and Democrat as reference partisanship

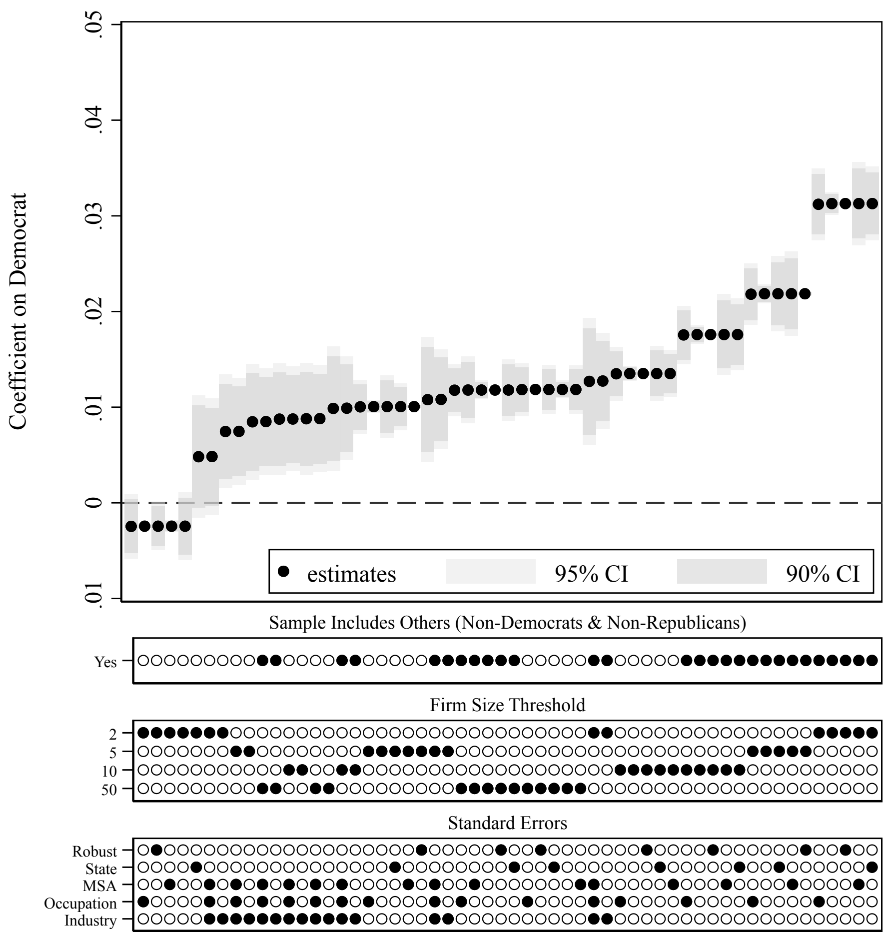

Points are from various specifications of the version of Equation 1.2 that includes MSA-occupation-industry fixed effects. The specifications vary in terms of whether they include workers who do not identify as either Democrat or Republican, the inclusion criteria in terms of number of co-workers, and the level at which we cluster standard errors.

Figure A.12: Specification curve with full fixed effects and Republican as reference partisanship

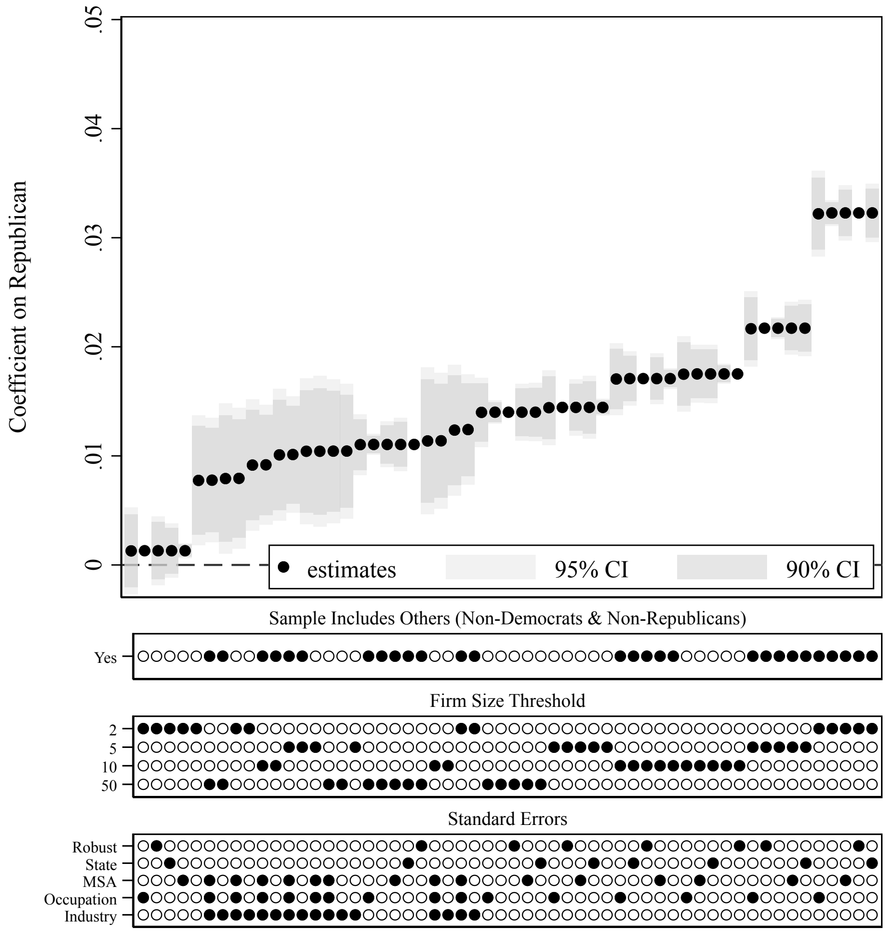

Points are from various specifications of the version of Equation 1.2 that includes MSA-occupation-industry fixed effects. Here, the outcome variable is the share of Republican coworkers and the explanatory variable is whether or not the focal worker is a Republican. The specifications vary in terms of whether they include workers who do not identify as either Democrat or Republican, the inclusion criteria in terms of number of co-workers, and the level at which we cluster standard errors.

Figure A.13: Distribution of in-party and cross-party exposure, by political affiliation excluding non-partisans

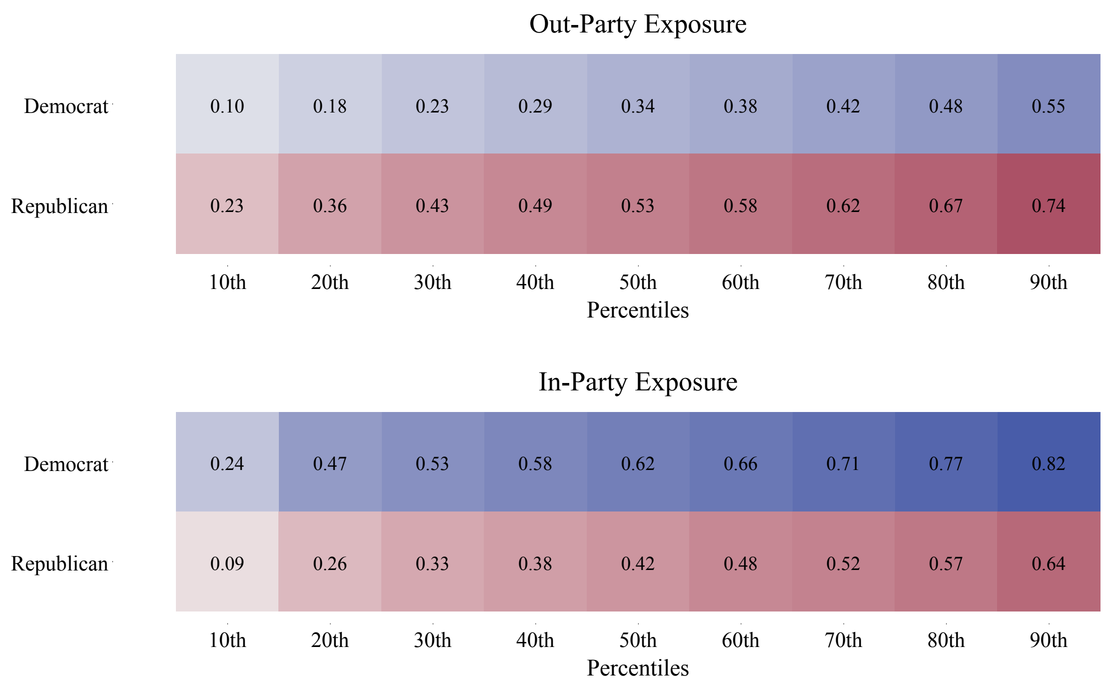

Table A.1: MSAs with highest match rates

	L2 Records	Matched	Rate
Columbia MO MSA	126,830	23,701	18.7
Salt Lake City UT MSA	591,290	110,246	18.6
Austin-Round Rock TX MSA	1,468,106	270,278	18.4
Des Moines-West Des Moines IA MSA	462,846	84,278	18.2
Fargo ND-MN MSA	136,905	24,701	18.0
Lincoln NE MSA	214,045	38,395	17.9
Boise City-Nampa ID MSA	412,643	73,704	17.9
Seattle-Tacoma-Bellevue WA MSA	2,545,161	450,719	17.7
Cheyenne WY MSA	44,519	7,756	17.4
Denver-Aurora-Lakewood CO MSA	2,030,423	353,476	17.4
Omaha-Council Bluffs NE-IA MSA	617,117	106,332	17.2
Provo-Orem UT MSA	329,921	56,727	17.2
Burlington-South Burlington VT MSA	167,381	28,423	17.0
Raleigh-Cary NC MSA	1,023,192	173,664	17.0
Ames IA MSA	75,949	12,730	16.8
Colorado Springs CO MSA	487,450	79,248	16.3
Missoula MT MSA	82,972	13,453	16.2
Wichita KS MSA	394,898	63,772	16.1
Ithaca NY MSA	80,512	12,997	16.1
State College PA MSA	106,014	17,095	16.1

This table depicts the MSAs where we were able to attain the highest match rates using the exact name matching method. Many of these MSAs are college towns, which could be associated with relatively high use of online professional social media (e.g., LinkedIn), facilitating higher match rates.

Table A.2: MSAs with lowest match rates

	L2 Records	Matched	Rate
Albany GA MSA	123,145	808	0.7
Gainesville GA MSA	269,284	2,006	0.7
Rome GA MSA	118,469	935	0.8
Brunswick GA MSA	94,010	1,009	1.1
Naples-Marco Island FL MSA	268,556	3,294	1.2
Mansfield OH MSA	124,688	1,725	1.4
Dalton GA MSA	113,003	1,564	1.4
Warner Robins GA MSA	134,778	2,098	1.6
Athens-Clarke County GA MSA	190,975	2,986	1.6
Savannah GA MSA	277,471	4,388	1.6
Jackson TN MSA	97,833	1,548	1.6
Columbus GA-AL MSA	248,896	4,041	1.6
Valdosta GA MSA	101,366	1,743	1.7
Hinesville-Fort Stewart GA MSA	49,117	860	1.8
Columbus IN MSA	74,374	1,465	2.0
St. Joseph MO-KS MSA	91,935	1,865	2.0
Fond du Lac WI MSA	97,654	2,300	2.4
Lima OH MSA	81,498	2,039	2.5
Macon GA MSA	247,794	6,310	2.5
Cleveland TN MSA	85,193	2,234	2.6

NOTE: The lowest match rates are found in less-populated MSAs, mostly located in the South, which may have lower rates of online job profile usage. Others (such as Brunswick, GA and Naples-Marco Island, FL) are coastal retirement destinations where many residents may not be in employment.

Table A.3: Position-level summary statistics for full and partisan samples for the 2012-2022 panel

	(1) Mean	(2) Standard Deivation	(3) Min	(4) Max	(5) Obs
For Full Sample					
Democrat	0.43	0.50	0.00	1.00	121,317,174
Republican	0.28	0.45	0.00	1.00	121,317,174
Other Party or Independent	0.29	0.45	0.00	1.00	121,317,174
Woman	0.54	0.50	0.00	1.00	121,317,174
Asian	0.06	0.23	0.00	1.00	99,253,667
Black	0.03	0.17	0.00	1.00	99,253,667
Hispanic	0.10	0.30	0.00	1.00	99,253,667
Other	0.05	0.22	0.00	1.00	99,253,667
White	0.77	0.42	0.00	1.00	99,253,667
Holds Bachelor’s	0.89	0.31	0.00	1.00	99,253,667
Holds Graduate Degree	0.36	0.48	0.00	1.00	99,253,667
Percent Dem Coworkers	0.43	0.25	0.00	1.00	121,317,174
Percent Rep Coworkers	0.28	0.24	0.00	1.00	121,317,174
Percent Other Coworkers	0.29	0.22	0.00	1.00	121,317,174
Unique Workers					14,824,737
Unique Positions					32,259,929
For Democrats					
Woman	0.59	0.49	0.00	1.00	52,275,116
Asian	0.06	0.25	0.00	1.00	40,499,849
Black	0.06	0.23	0.00	1.00	40,499,849
Hispanic	0.13	0.34	0.00	1.00	40,499,849
Other	0.06	0.23	0.00	1.00	40,499,849
White	0.69	0.46	0.00	1.00	40,499,849
Holds Bachelor’s	0.90	0.29	0.00	1.00	37,615,213
Holds Graduate Degree	0.39	0.49	0.00	1.00	37,615,213
Percent Dem Coworkers	0.51	0.25	0.00	1.00	52,275,116
Percent Rep Coworkers	0.23	0.21	0.00	1.00	52,275,116
Percent Other Coworkers	0.26	0.20	0.00	1.00	52,275,116
Unique Workers					6,395,388
Unique Positions					14,380,215
For Republicans					
Woman	0.50	0.50	0.00	1.00	34,163,478
Asian	0.03	0.17	0.00	1.00	29,511,063
Black	0.00	0.06	0.00	1.00	29,511,063
Hispanic	0.06	0.23	0.00	1.00	29,511,063
Other	0.03	0.18	0.00	1.00	29,511,063
White	0.88	0.33	0.00	1.00	29,511,063
Holds Bachelor’s	0.88	0.32	0.00	1.00	22,003,698
Holds Graduate Degree	0.33	0.47	0.00	1.00	22,003,698
Percent Dem Coworkers	0.35	0.24	0.00	1.00	34,163,478
Percent Rep Coworkers	0.39	0.26	0.00	1.00	34,163,478
Percent Other Coworkers	0.26	0.21	0.00	1.00	34,163,478
Unique Workers					4,078,568
Unique Positions					8,109,288

Units of analysis are position-years. Sample includes jobs active at any point between 2012 and August 2022. Differences in number of observations across variables within the same subsample reflect missing data.

Table A.4: For panel of workers employer, occupation, and industry-level summary statistics for the 2012-2022 panel

	(1) Full Sample	(2) Democrats	(3) Republicans
Employer Level			
Position-Years/Employer	104.31	57.74	45.72
	(1,624.76)	(787.53)	(582.82)
Unique Employers	1,163,065	905,337	747,309
Position-Years/Employer-MSA	54.01	31.070	24.50
	(438.77)	(261.88)	(145.65)
Unique Employer-MSAs	2,246,349	1,682,486	1,394,260
Occupation Level			
Position-Years/Occupation	294,147.6	126,908.3	82,231.65
	(397,790.6)	(185,204.5)	(114,792.9)
Unique Occupations	384	384	384
Industry Level			
Position-Years/Industry	86,249.05	37,115.67	24,161.94
	(37,1564.9)	(18,4945.1)	(89,577.77)
Unique Industries	1,052	1,051	1,050

Column 1 includes our full sample. Columns 2 and 3 include only Democrats and Republicans, respectively. Units of analysis are position-years. Sample includes jobs active at any point between 2012 and August 2022.

Appendix B

Appendices for Chapter 2

Figure B.1: Quarterly employee interstate migration figures (number of migrants), 2011–2019

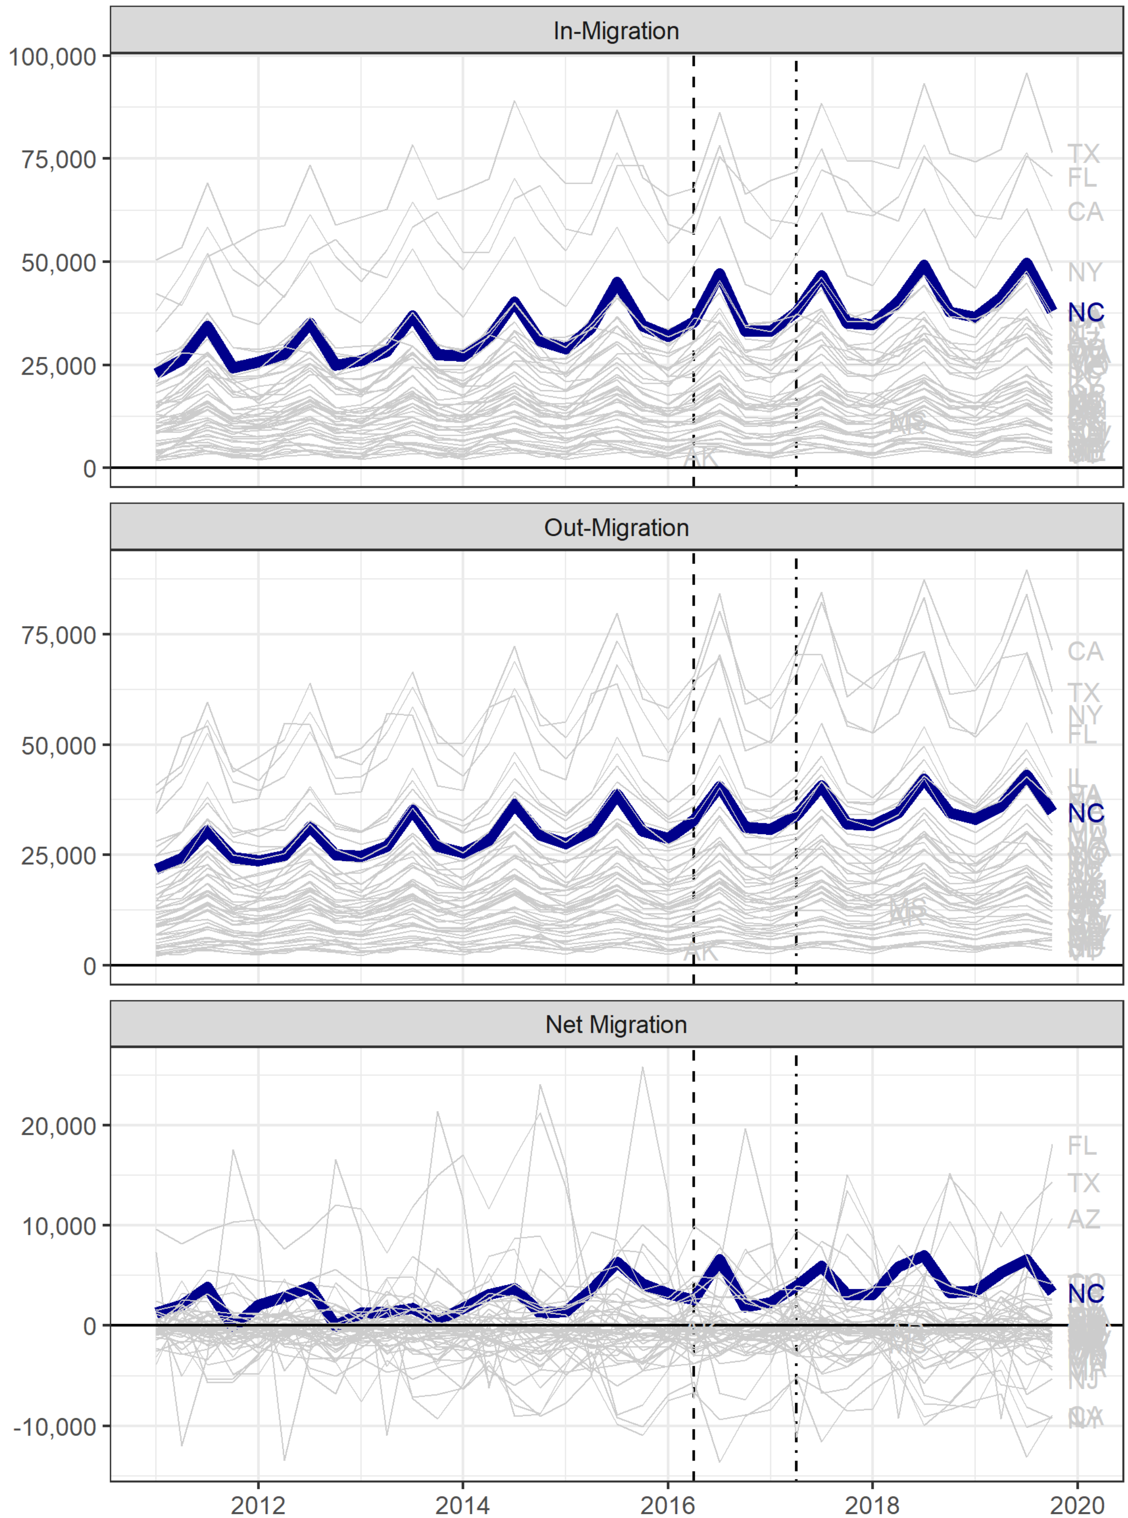

Includes both employer-employer (EE) and adjacent-quarter (AQHires). Dashed vertical dashed line is at 2016 Q2, when HB2 went into effect in North Carolina. Dot-dash vertical line is at 2017 Q2, when HB2 was partially repealed (HB142). Data for AK, AR, and MS are incomplete in LEHD J2J data. All data are highly seasonal, with interstate moves being most common during summer (Q3). In-migration and out-migration rates generally move in tandem, except for natural-resource dependent states with small populations (e.g., AK, ND, WY).

Figure B.2: Quarterly employee interstate migration rate, 2011–2019, by state

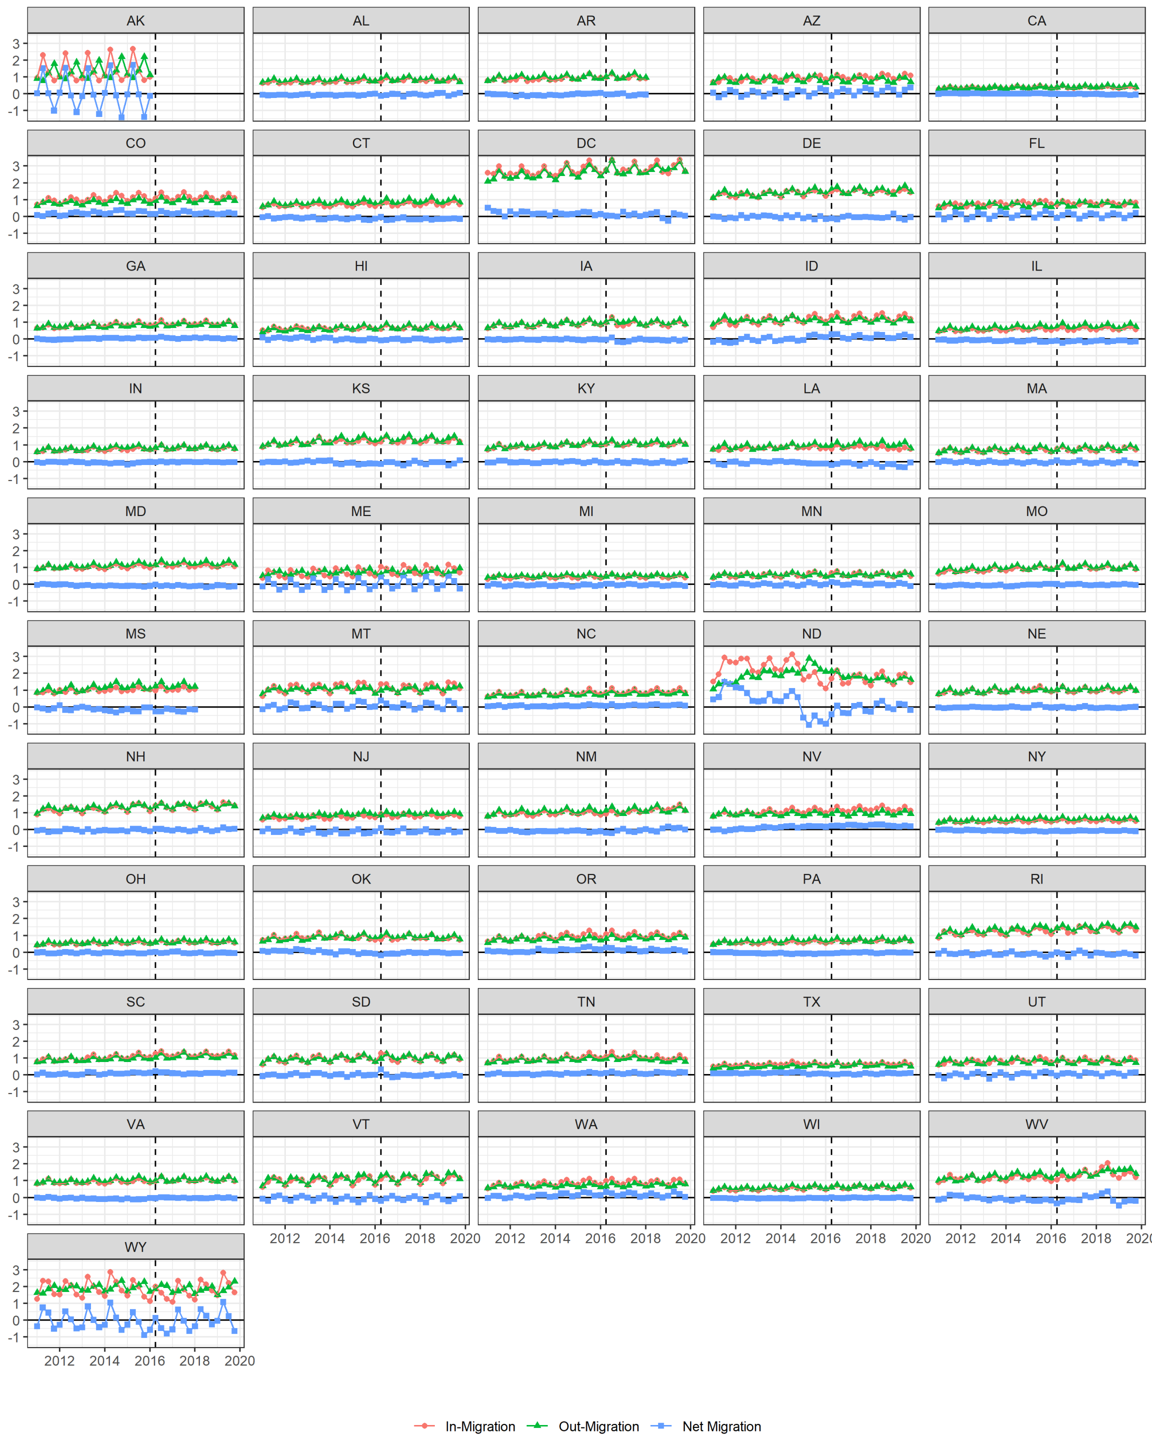

Includes both employer-employer (EE) and adjacent-quarter (AQHires). Vertical dashed line is at 2016 Q2, when HB2 went into effect in North Carolina. Data for AK, AR, and MS are incomplete in LEHD J2J data. All data are highly seasonal, with interstate moves being most common during summer (Q3). In-migration and out-migration rates generally move in tandem, except for natural-resource dependent states with small populations (e.g., AK, ND, WY).

Figure B.3: Quarterly employee interstate migration rate, 2011–2019, by state, seasonally adjusted

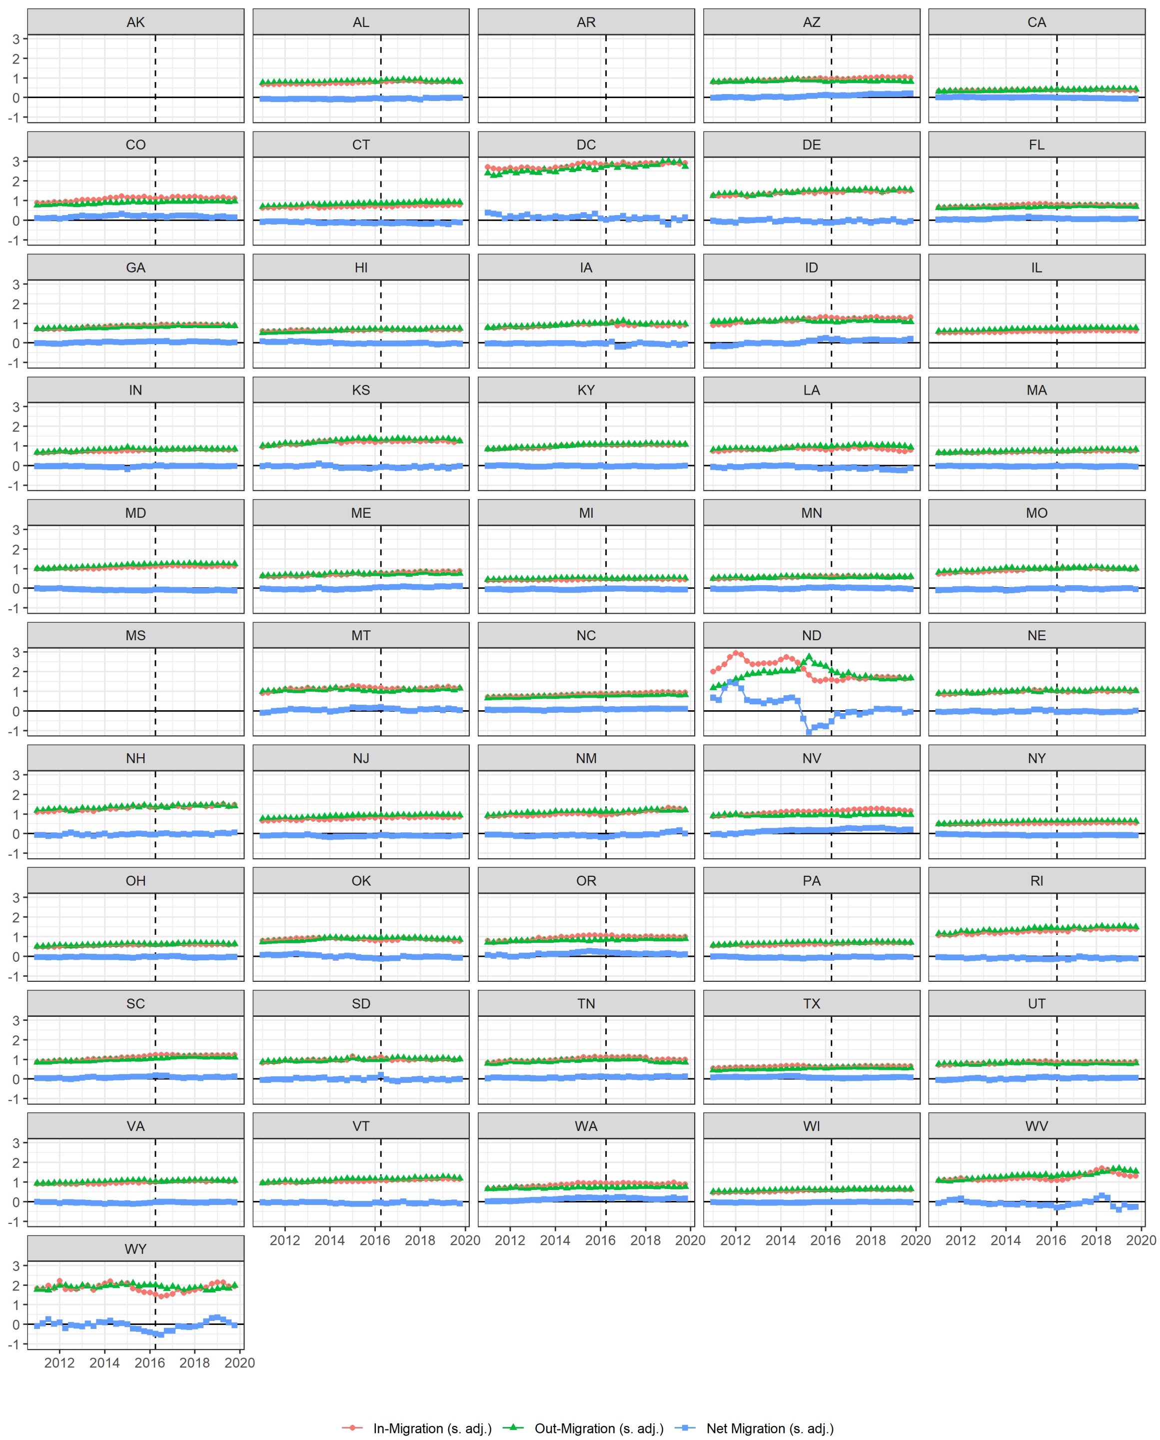

Includes both employer-employer (EE) and adjacent-quarter (AQHires). Vertical dashed line is at 2016 Q2, when HB2 went into effect in North Carolina. I do not calculate seasonal adjustments for AK, AR, and MS as they are incomplete in the LEHD J2J data. Data are seasonally adjusted using the `seasonal` package, which provides an R interface to the `X13ARIMA-SEATS` adjustment program developed by the US Census Bureau (Sax and Eddelbuettel 2018, 2022). In-migration and out-migration rates generally move in tandem, except for natural-resource dependent states with small populations (e.g., AK, ND, WY).

Figure B.4: Quarterly employee interstate migration figures (number of migrants), 2011–2019, by state

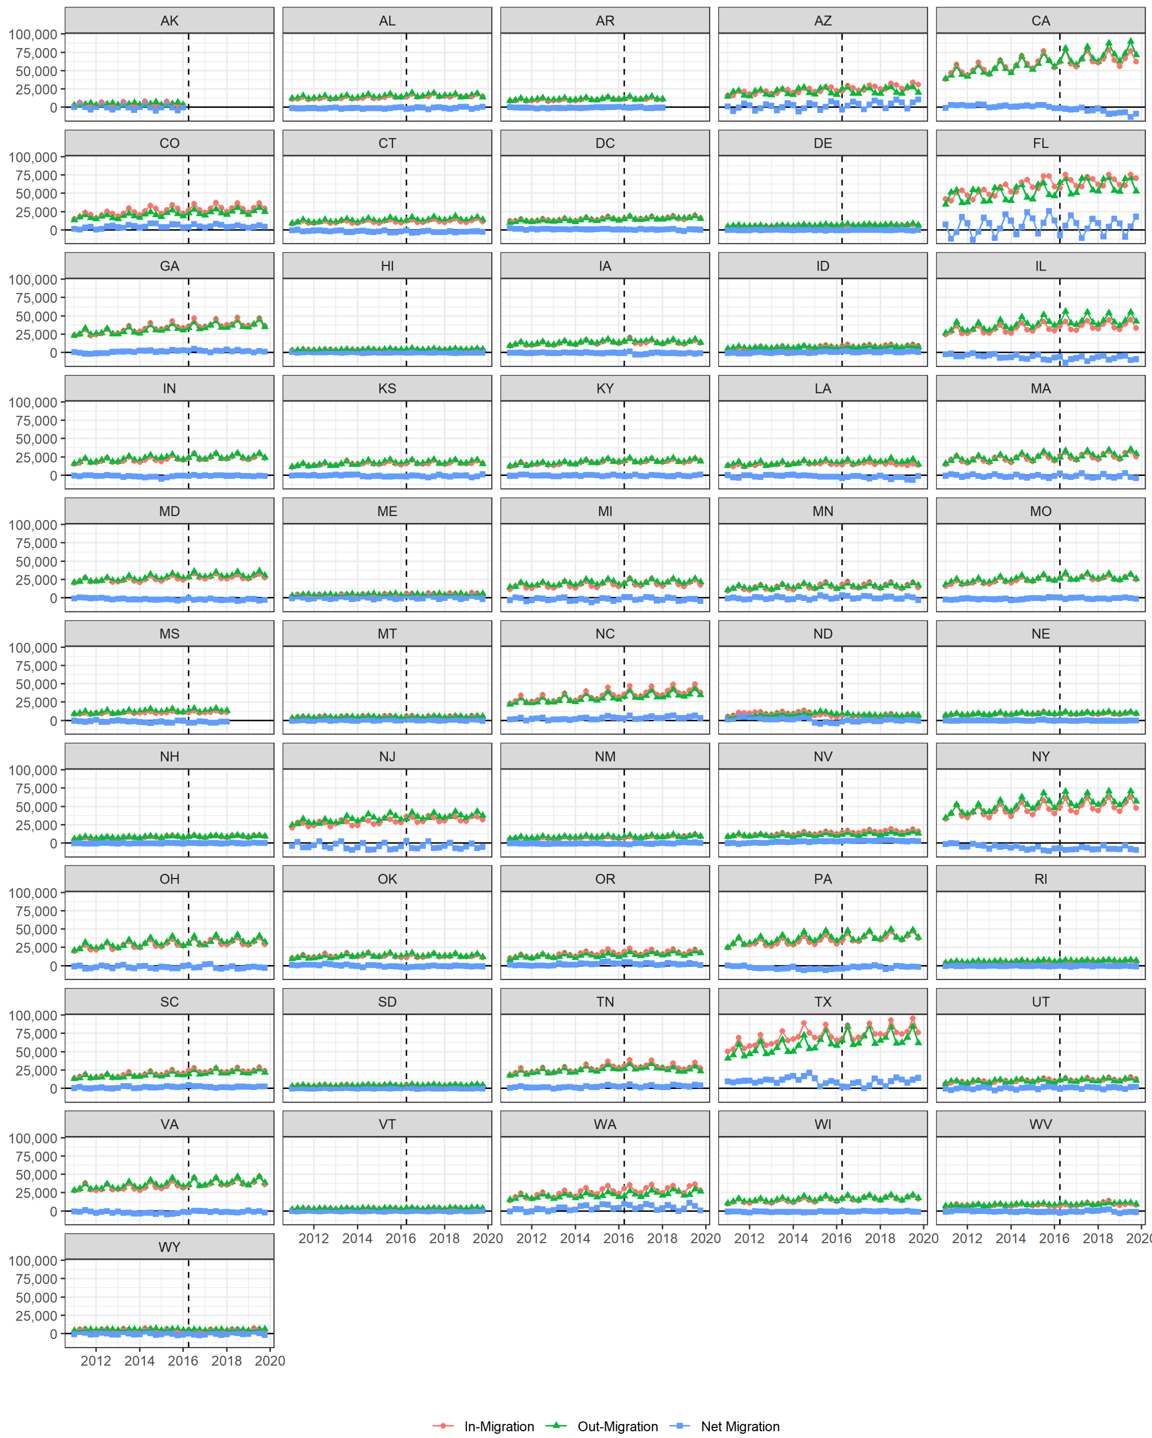

Includes both employer-employer (EE) and adjacent-quarter (AQHires). Vertical dashed line is at 2016 Q2, when HB2 went into effect in North Carolina. Data for AK, AR, and MS are incomplete in LEHD J2J data. All data are highly seasonal, with interstate moves being most common during summer (Q3). In-migration and out-migration rates generally move in tandem, except for natural-resource dependent states with small populations (e.g., AK, ND, WY).

Figure B.5: Quarterly employee interstate migration figures (number of migrants), 2011–2019, by state, seasonally adjusted

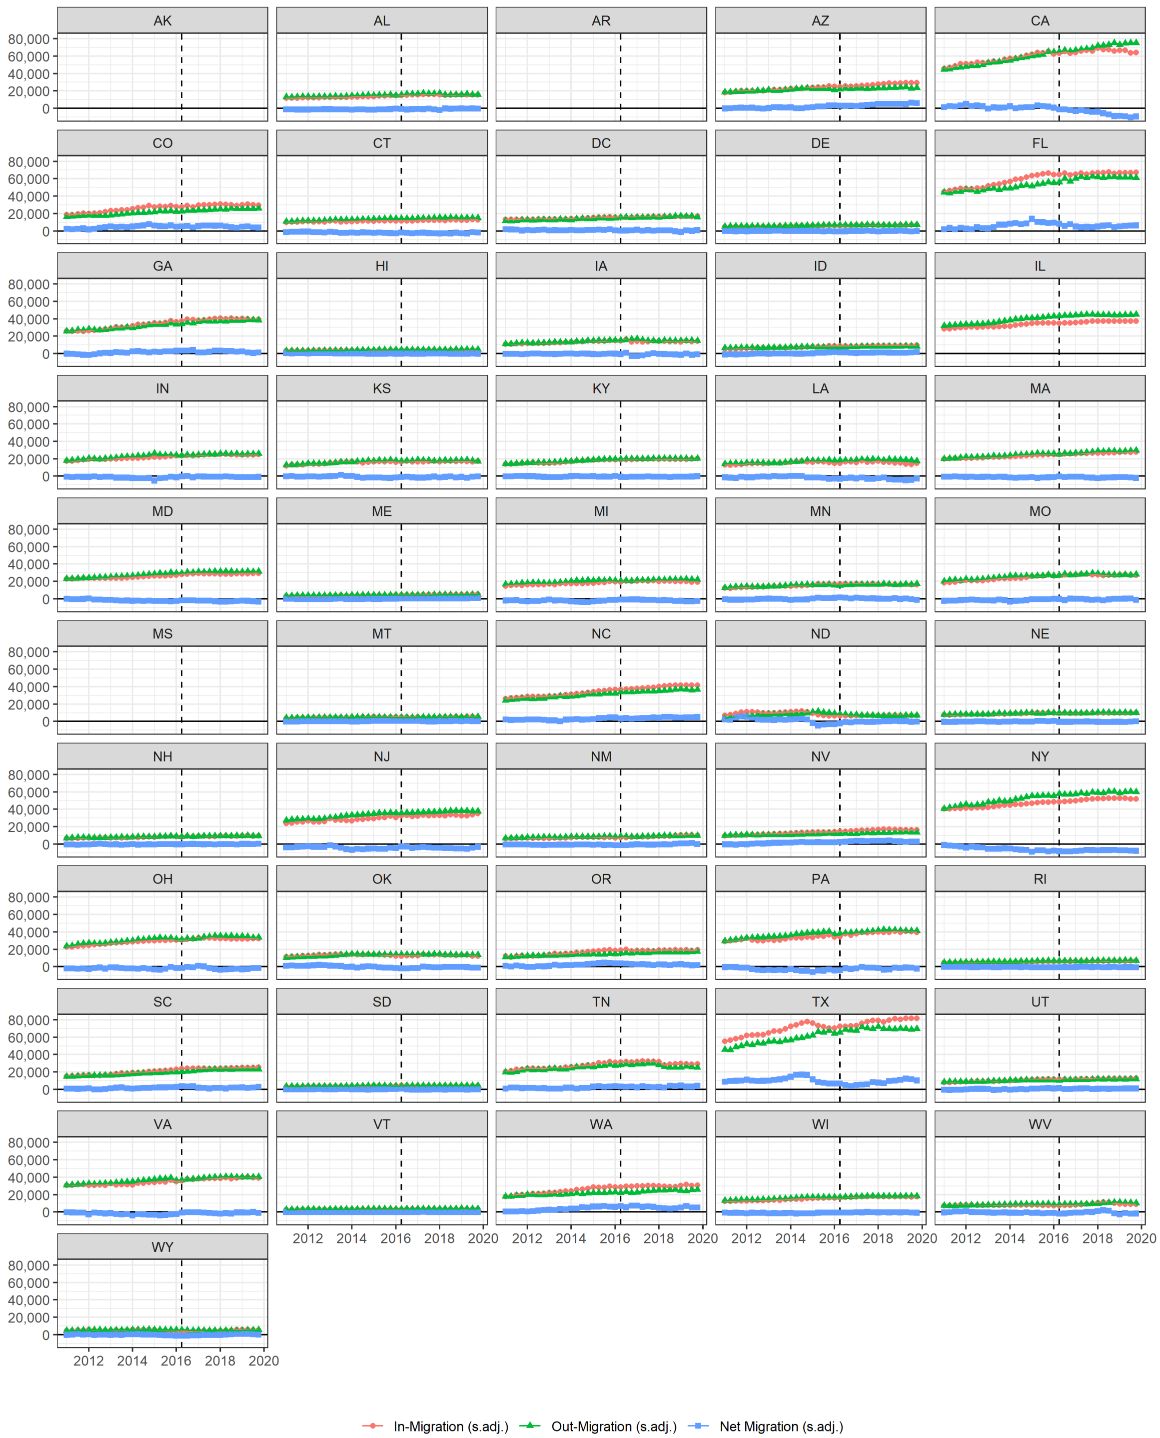

Includes both employer-employer (EE) and adjacent-quarter (AQHires). Vertical dashed line is at 2016 Q2, when HB2 went into effect in North Carolina. I do not calculate seasonal adjustments for AK, AR, and MS as they are incomplete in the LEHD J2J data. Data are seasonally adjusted using the `seasonal` package, which provides an R interface to the `X13ARIMA-SEATS` adjustment program developed by the US Census Bureau (Sax and Eddelbuettel 2018, 2022). In-migration and out-migration rates generally move in tandem, except for natural-resource dependent states with small populations (e.g., AK, ND, WY).

Overall In-Migration

Figure B.6: Estimated effect on in-migration to North Carolina (excluding DC, ND, and WY)

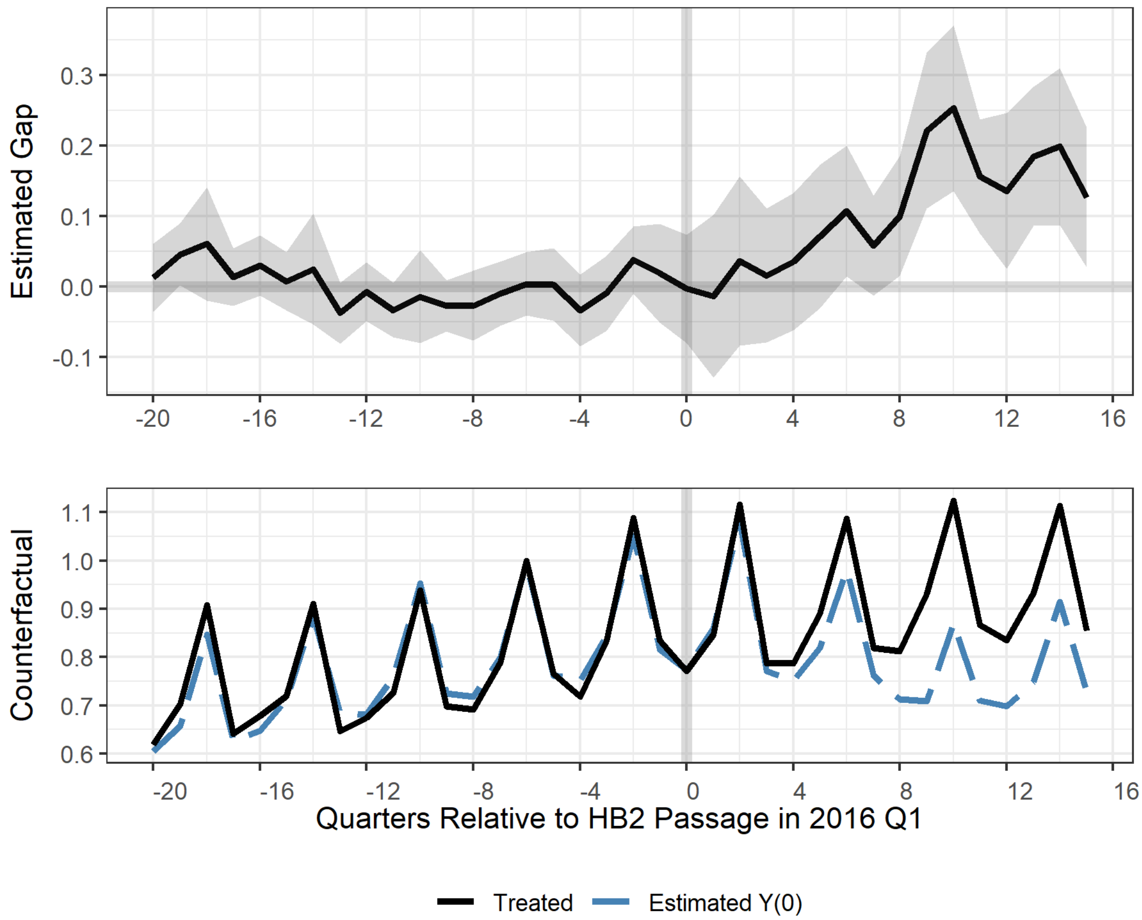

Generalized synthetic control (GSC) estimates of the effect of HB2 on in-migration to North Carolina. Top panel is a "gap" plot and bottom panel depicts the average among treated dyads (dyads with North Carolina as the destination) versus the synthetic estimate for treated dyads. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure B.7: Implied weights for estimated effect of in-migration to North Carolina (excluding DC, ND, and WY)

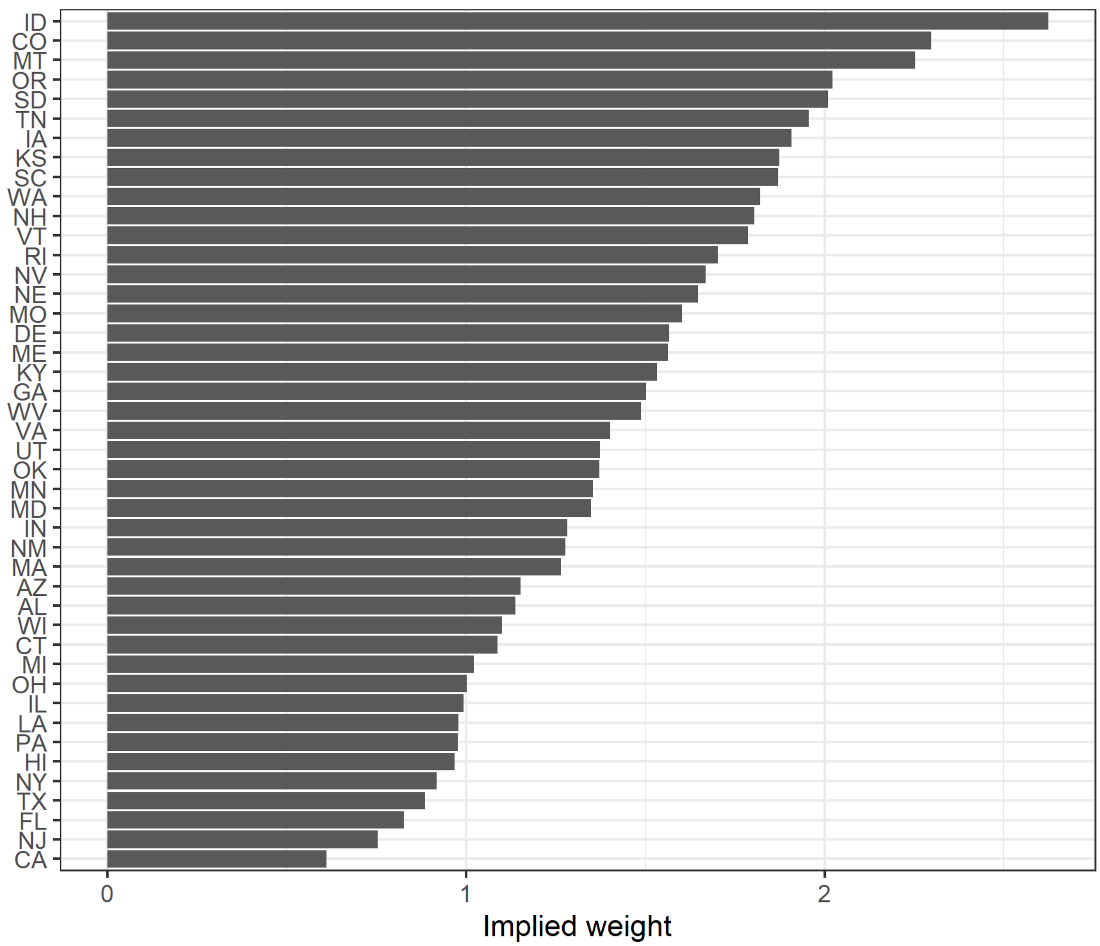

Implied generalized synthetic control weights for estimate.

Overall Out-Migration

Figure B.8: Implied weights for estimated effect of HB2 on out-migration from North Carolina

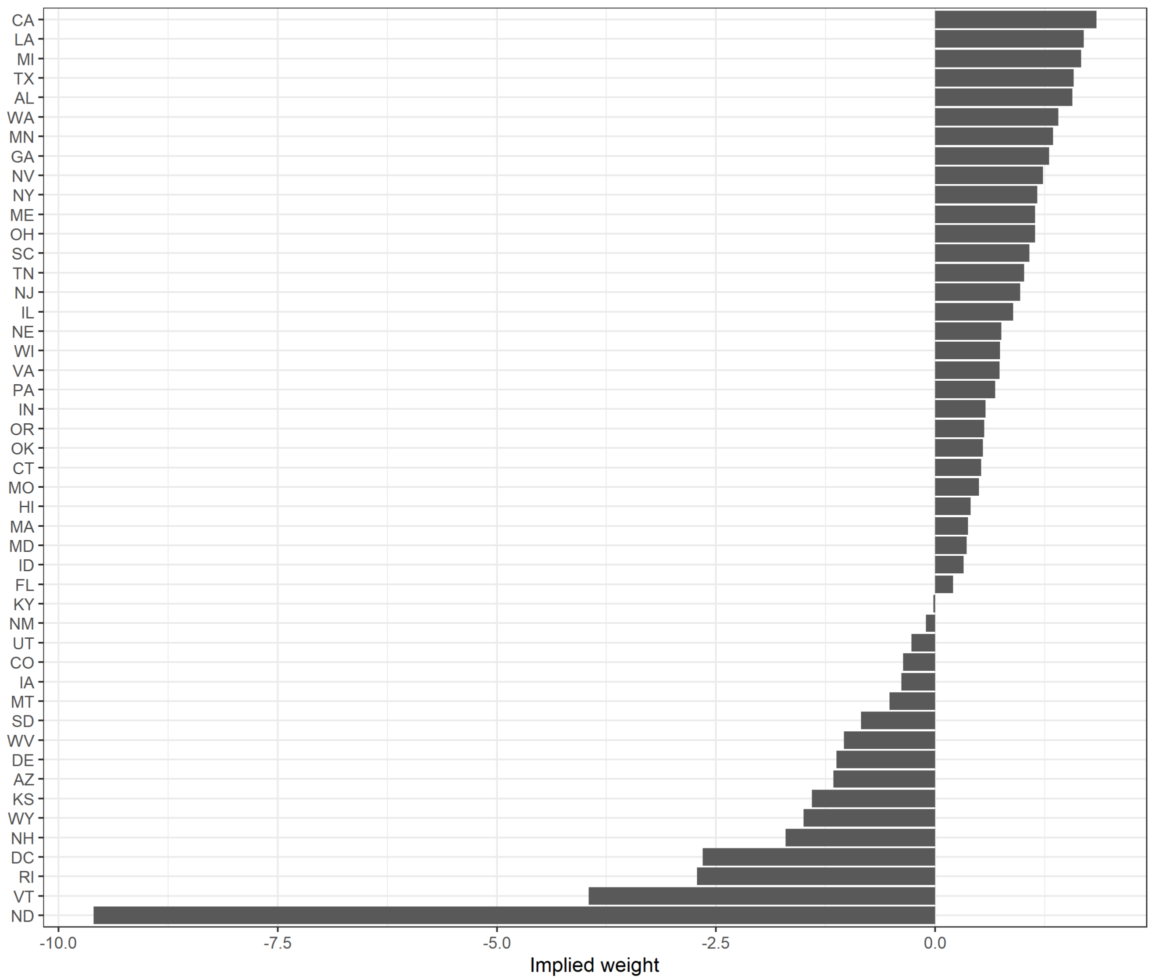

Implied weights for generalized synthetic model fit in Figure B.9.

Figure B.9: Estimated effect of HB2 on out-migration from North Carolina

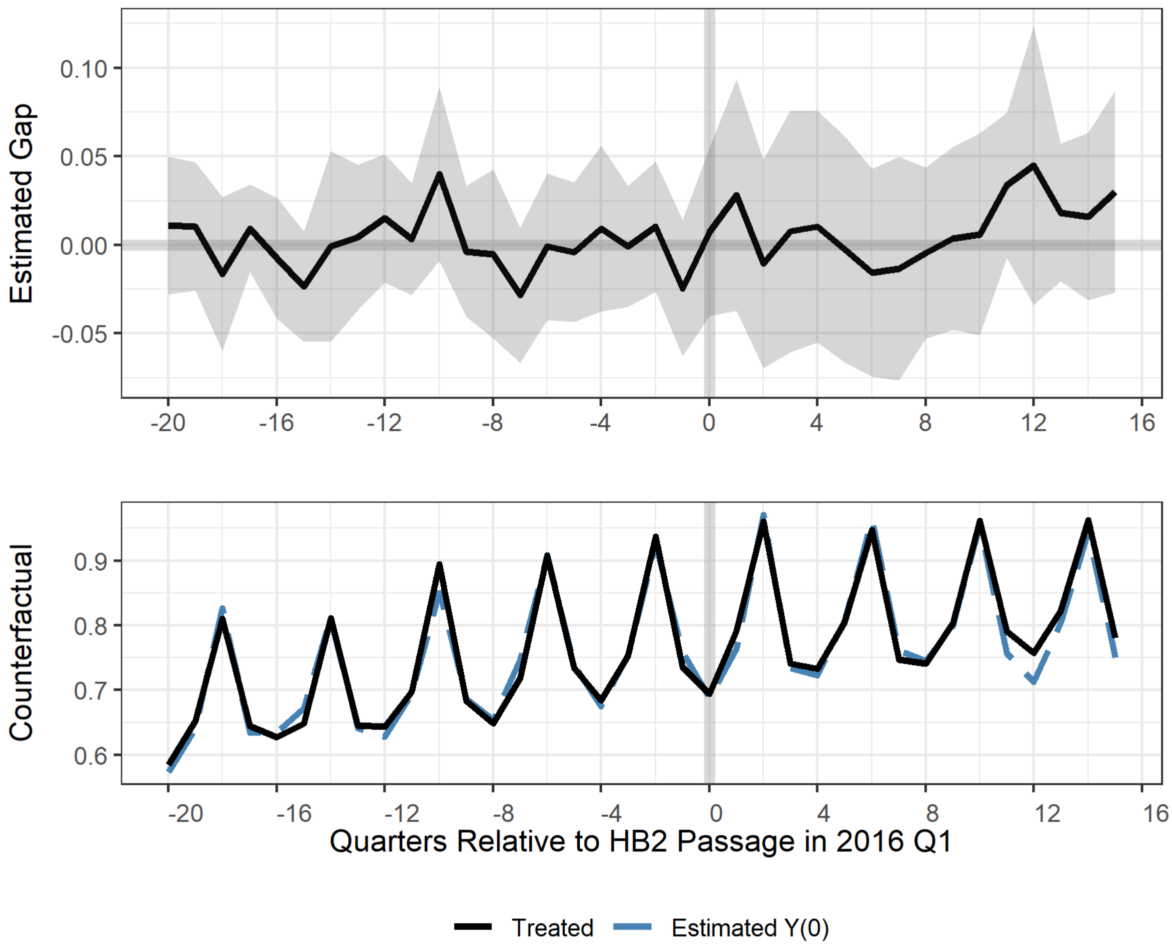

Top panel is a "gap" plot and bottom panel depicts the actual observed values for North Carolina versus the synthetic estimate of North Carolina under control. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure B.10: Estimated effect of HB2 on out-migration from North Carolina, excluding DC, ND, and WY

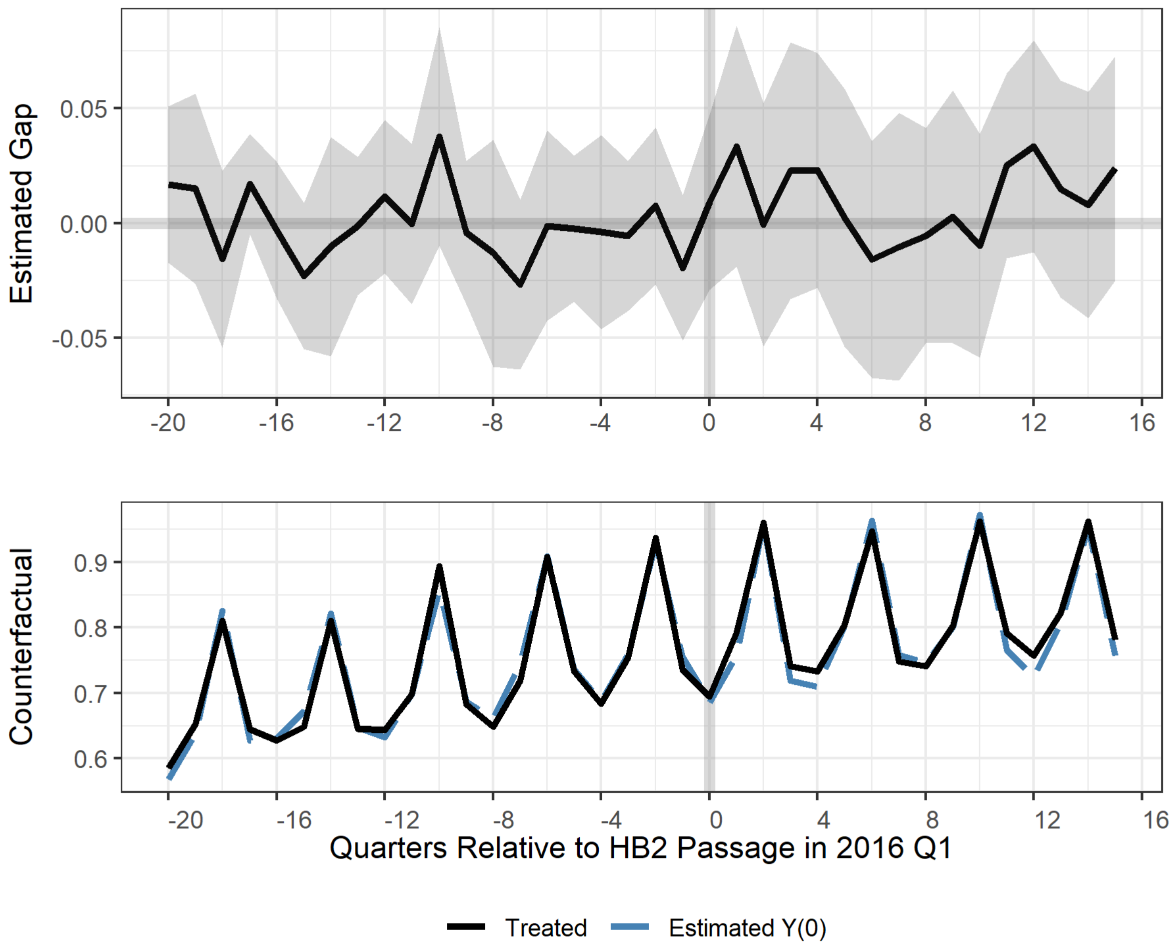

Top panel is a "gap" plot and bottom panel depicts the average among control states versus North Carolina. Analysis makes use of the Expectation Maximization algorithm (Gobillon and Magnac 2016). Inference is parametric; standard errors are estimated based upon 1000 bootstraps.

Figure B.11: Implied weights for estimated effect of HB2 on out-migration from North Carolina

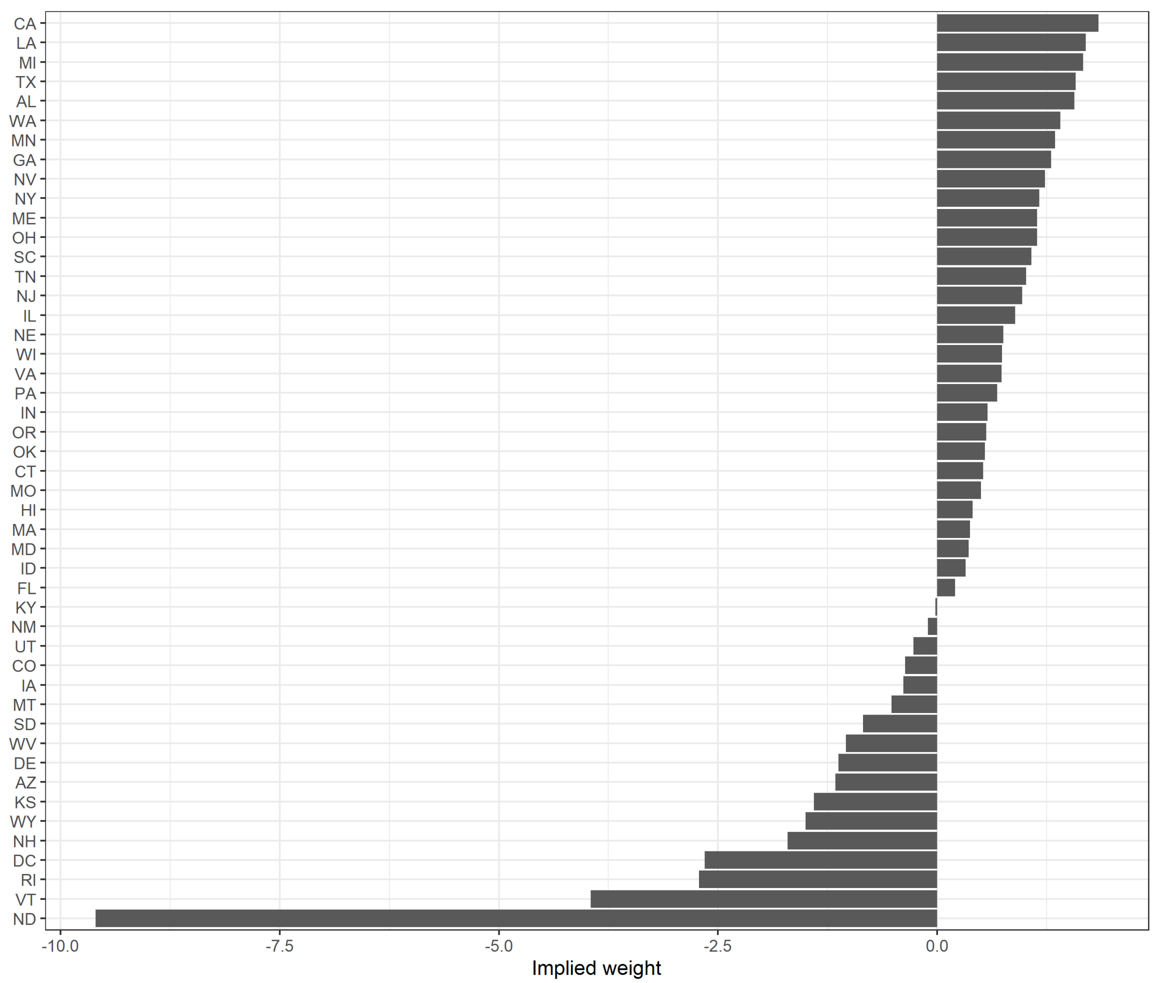

Implied weights for estimate.

Figure B.12: Implied weights for estimated effect of HB2 on out-migration from North Carolina, excluding DC, ND, and WY

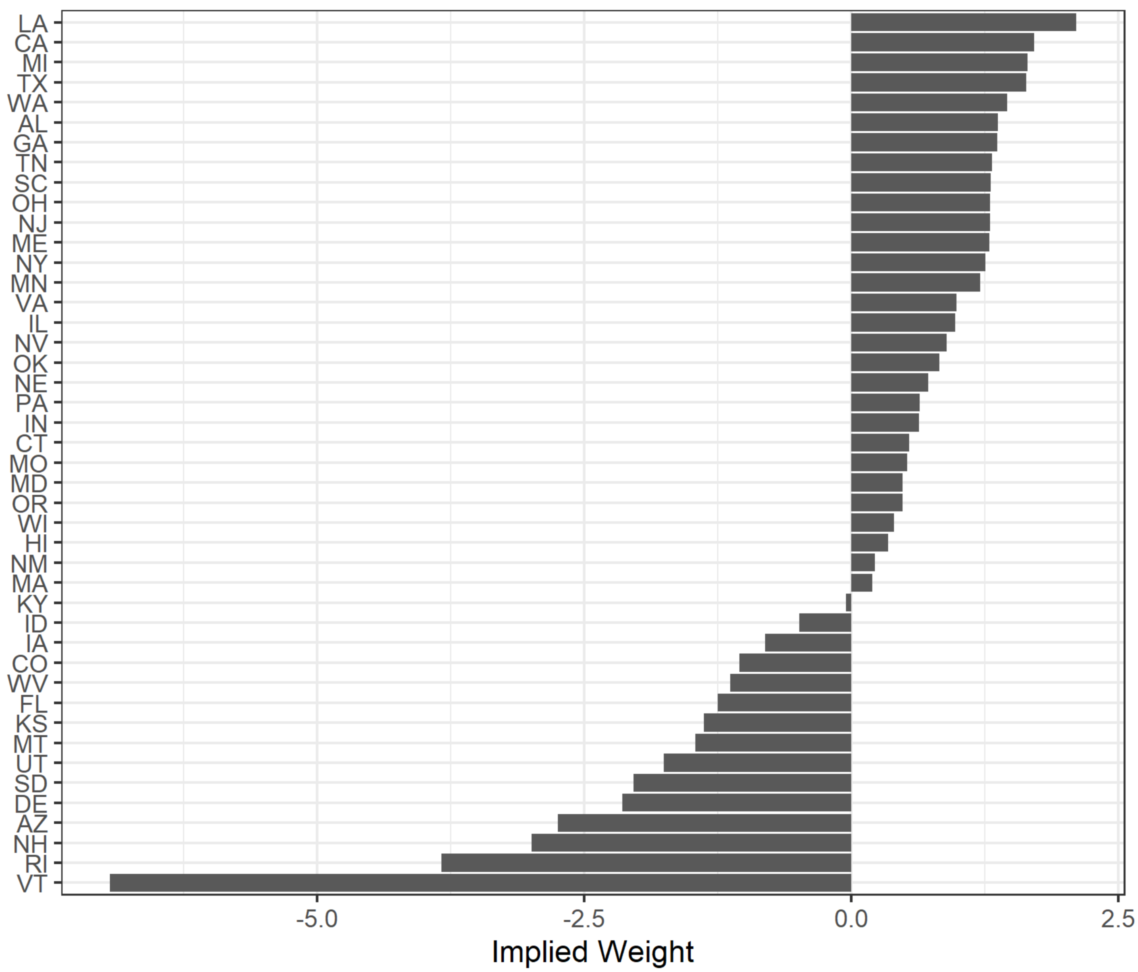

Implied weights for estimate.

Wage differentials

Table B.1: Difference-in-difference of change in quarterly earnings for movers From North Carolina, by education category

	Change in Quarterly Earnings for Movers (Dollars)				
Education	E1	E2	E3	E4	E5
	(1)	(2)	(3)	(4)	(5)
NC Origin × Post HB2	16.34	79.66	123.7	312.5*	-19.39
	(144.4)	(131.9)	(138.3)	(182.9)	(87.77)
Observations	70,765	75,041	75,950	75,776	68,640
Origin-Destination FE	✓	✓	✓	✓	✓
Year-Quarter FE	✓	✓	✓	✓	✓

Difference-in-difference compares those moving *from* North Carolina to other states versus those moving out of other states. If HB2 constituted a significant negative labor supply shock to North Carolina, we would expect a relatively higher compensating wage differential for North Carolina. Similarly, we would also expect a negative wage differential for those leaving North Carolina, as they would be prepared to accept lower wages in order to leave the state. We do not find this. See also event study plot.

Figure B.13: Event plot of DiD estimates for salary differences for movers from NC

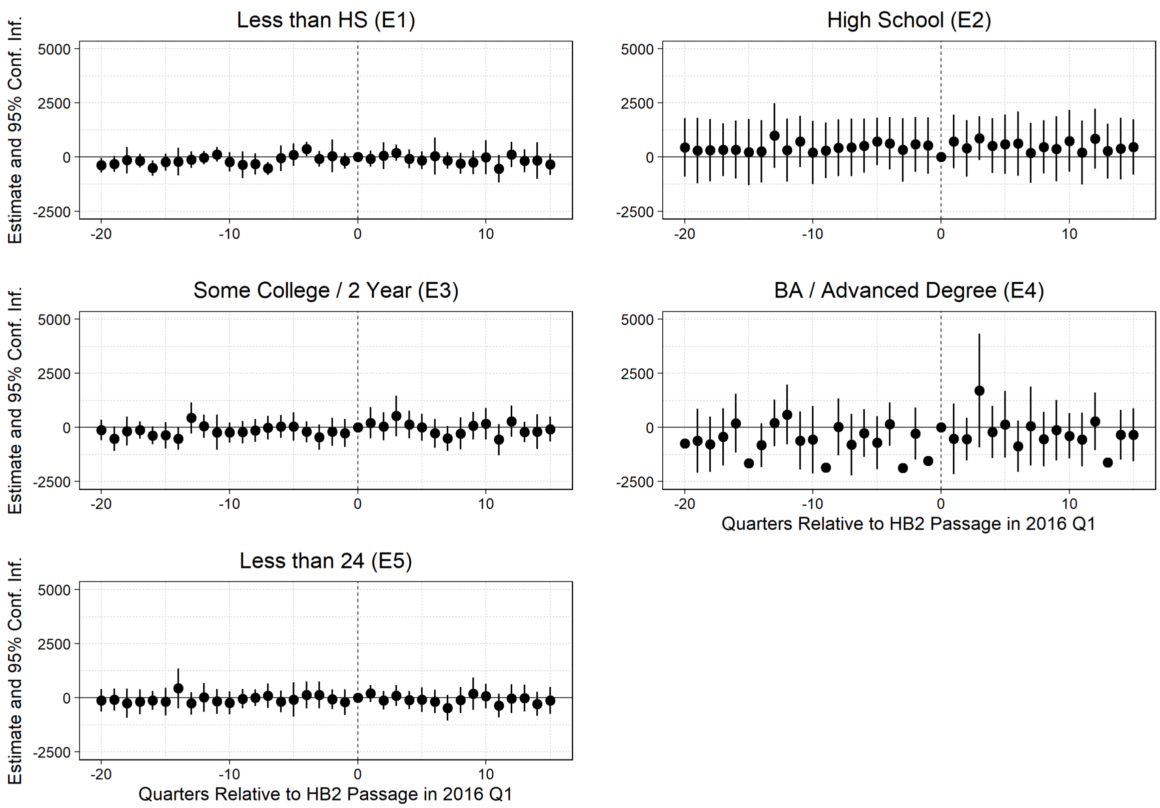

Difference-in-difference (DiD) analysis where the dependent variable is the change in quarterly wages for employer-employer movers (i.e., those who move directly from one job to another without an quarter of labor market inactivity) between their old job (in quarter $t-1$ and their new job in quarter t. Standard errors are clustered by origin-destination pair and regression is weighted by the number of individual workers used to compute the average wages in each cell defined by origin, destination, quarter, and subpopulation (education category). Observations are at the origin-destination-quarter level; missing observations differ across categories due Census Bureau data depression rules for cells with small counts. Each educational category is modeled separately: E1 (less than high school), E2 (high school or equivalent, no college), E3 (some college or Associate's degree), E4 (Bachelor's or advanced degree), and E5 (workers 24 years of age or younger, for whom educational attainment is not available). Time period begins in 2011 (20 pre-treatment periods), treatment begins in 2016 Q2, and time period ends in 2019 Q4 (15 post-treatment periods).

Appendix C

Appendices for Chapter 3

C.1 Additional Details on Experimental Setup

Brands Used in the Survey

The survey asks about nine of the most popular beer brands in America. Because of considerable consolidation within the beer industry, many of the most popular U.S. brands are owned by the same conglomerates. While in some cases, this common ownership is obvious (e.g., Bud Light and Budweiser are both owned by Anheuser-Busch), in other cases, it is less clear (e.g., Michelob Ultra is also owned by Anheuser-Busch). By testing multiple brands with different corporate relationships to the main focal brand, I can better understand the motivations for political consumerism in this context. If respondents are motivated by instrumental reasons, their attitude towards Bud Light should be mirrored by other brands owned by Anheuser-Busch. If, however, respondents are primarily engaged in expressive political consumerism, their attitude towards Bud Light is mainly a signal to other partisans, and there is no reason why other Anheuser-Busch brands should be affected.

The brands used in the study were as follows:

- Anheuser-Busch brands
  - Bud Light (focal brand)
  - Budweiser
  - Busch Light
  - Michelob Ultra
- Other brands[1]
  - Corona (Constellation Brands)

[1]The Corona and Modelo Especial brands are owned by AB Inbev (the parent of Anheuser-Busch) outside of the United States. Due to antitrust concerns, the company divested the rights to these brands in the United States.

  - Coors Light (Miller Coors)
  - Heineken (Heineken)
  - Modelo Especial (Constellation Brands)
  - Miller Light (Miller Coors)

Simulated News Stories Used in Survey Experiment

1. **Republican boycott treatment**

   Republicans Boycott Bud Light After Controversy

   After a controversial social media campaign featuring a transgender influencer, conservatives are calling for a boycott of Bud Light.

   The controversy began when Dylan Mulvaney, a transgender actor and activist, posted a sponsored video celebrating the one-year anniversary of her first identifying as a woman. In the video, Mulvaney poses with a specially-wrapped can of Bud Light featuring her face on the label.

   Since the initial post, Bud Light has faced growing criticism from conservatives, with many pledging to boycott the brand. Prominent figures within the Republican Party have expressed support for the boycott in a variety of ways. Multiple Republican Congressmen posted videos showing themselves purchasing alternative brands or showing off a well-stocked fridge—without any Bud Light. Florida Governor Ron DeSantis reinforced calls for a boycott, stating that he did not want to support "woke companies" and asking "why would you want to drink Bud Light?"

   Many rank-and-file Republicans appear to have taken up the call for a boycott as a way of pushing back against "woke" values. Prominent conservative celebrities have posted themselves destroying Bud Light cans and advertising materials. In one notable video, 1990s singer-songwriter Kid Rock destroyed several cases of Bud Light with an automatic rifle. In others, everyday citizens run over Bud Light with trucks or pour beer down toilets.

   While it remains to be seen how much these events will affect the company, it is clear that many Republicans will not be drinking Bud Light any time soon.

2. **Republican Support treatment**

   Republicans Show Support for Bud Light After Controversy

   After Bud Light apologized for a controversial social media campaign featuring a transgender influencer, some conservatives are standing up to defend the brand.

   The controversy began when Dylan Mulvaney, a transgender actor and activist, posted a sponsored video celebrating the one-year anniversary of her first identifying as a woman. In the video, Mulvaney poses with a specially-wrapped can of Bud Light featuring her face on the label.

After the brand faced backlash on social media, the company sought to distance itself from the controversy in a number of ways. Most notably, the CEO of the company issued a public apology for being ”part of a discussion that divides people.” The company also announced that the marketing executive responsible for approving the campaign was placed on leave. Finally, the brand announced that new ad campaigns would feature patriotic themes and emphasize the brand's support for veterans.

Given the brand's apology, many Republicans and conservatives appear prepared to forgive and move on. For one, President Trump appears to be standing behind the brand: Bud Light remains on the menu at Trump hotels and resorts in Chicago, Florida, and Las Vegas. The President's son, Donald Trump, Jr., also defended the company. On his podcast, Trump, Jr. argued that despite their recent campaign, the company that makes Bud Light has generally been much less ”woke” than other conglomerates. Others also pointed out that Bud Light's parent company has long been a financial supporter of the Republican party.

While it remains to be seen how much these events will affect the company, it is clear that many Republicans are prepared to move on and share a Bud Light with their fellow conservatives.

3. **Democratic Support treatment**

   Democrats Boycott Bud Light After Controversy

   After Bud Light apologized for a controversial social media campaign featuring a transgender influencer, some liberals are calling for a boycott.

   The controversy began when Dylan Mulvaney, a transgender actor and activist, posted a sponsored video celebrating the one-year anniversary of her first identifying as a woman. In the video, Mulvaney poses with a specially-wrapped can of Bud Light featuring her face on the label.

   After the brand faced backlash on social media, the company sought to distance itself from the controversy in a number of ways. Most notably, the CEO of the company issued a public apology for being “part of a discussion that divides people.”

   The company also announced that the marketing executive responsible for approving the campaign was placed on leave. Finally, the brand announced that new ad campaigns would feature patriotic themes and emphasize the brand's support for veterans.

   Given the brand's apology, many Democrats and liberal organizations have called for a boycott of Bud Light. News organizations reported that gay bars were removing the brand in protest at the company's lack of support for the transgender community. The Human Rights Campaign, a major LGBT advocacy organization, also removed Bud Light's parent company from it's “Corporate Equality Index” after the company declined to release a statement of support for its transgender customers, shareholders, and employees.

   While it remains to be seen how much these events will affect the company, it is clear that many Democrats will not be drinking Bud Light any time soon.

4. **Democratic Support treatment**

   Democrats Show Support for Bud Light After Controversy

   After a controversial social media campaign featuring a transgender influencer, liberals are are standing up to defend Bud Light.

   The controversy began when Dylan Mulvaney, a transgender actor and activist, posted a sponsored video celebrating the one-year anniversary of her first identifying as a woman. In the video, Mulvaney poses with a specially-wrapped can of Bud Light featuring her face on the label.

   After the brand faced backlash on social media, a group of Democratic lawmakers from California posted to social media to show their support for the brand. Ted Lieu, a liberal Democrat who represents Los Angeles, posted a picture on Facebook of himself drinking a Bud Light alongside three Democratic colleagues. In a follow-up to the initial Twitter post, Lieu wrote that used the post to "mock stupid stuff by MAGA Republicans," making it clear that the photo op was no mere coincidence.

   Many rank-and-file Democrats also appear to be showing their support for Bud Light. A recent survey by polling group Morning Consult showed that two thirds of Democrats would feel favorably towards a beer brand with a transgender spokesperson. Others appear to see supporting Bud Light as a way of pushing back against conservatives. In a post on the Facebook group "Occupy Democrats," one commenter wrote that "I don't drink Bud… however I'm tempted to drink it now just on principle."

   While it remains to be seen how much these events will affect the company, it is clear that many Democrats may soon be cracking open a Bud Light to share with their fellow liberals.

5. **Control treatment**

   New Beer Industry Report Shows $409 Billion in Economic Impact

   Ahead of Memorial Day weekend and the start of the summer selling season, the Beer Institute and the National Beer Wholesalers Association (NBWA) released their biennial Beer Serves America report today on the economic impact and importance of the American brewing industry. The study found that the U.S. beer industry now supports nearly 2.4 million local jobs and contributes more than $409 billion to our economy – equivalent to 1.6% of GDP. The beer industry also pays more than $132 billion in wages and $63.8 billion in taxes.

   These economic figures represent substantial growth in the two years since the beer industry conducted the 2020 Beer Serves America study, which provided a snapshot of the significant pandemic challenges facing the hospitality industry and manufacturers. Since the last Beer Serves America report, the industry has expanded significantly, with an uptick of $78 billion in economic impact while adding nearly 400,000 new beer industry jobs across the country. The beer industry leads the rest of the alcohol industry in economic impact and job footprint, and beer remains America's favorite alcohol beverage.

"Beer continues to be America's favorite alcohol beverage because of its cultural heritage, its important place in our nation's history and its unique ability to bring people together," said Brian Crawford, president and CEO of the Beer Institute. "Americans have more options and ways to enjoy their favorite brews than ever because our $409 billion industry is competitive, vibrant and a crucial part of the American economy. The tremendous growth we've seen since our last report is a true success story that underscores beer's striking recovery coming out of the pandemic and showcases the resilience of our industry and the 2.4 million Americans it employs. The beer industry has always had a storied place in American culture and commerce, and as these new figures confirm, we have an incredibly bright future ahead of us."

C.2 Additional Figures and Tables

Table C.1: Descriptive Statistics of Sample

Variable	N	Percent
Generation	2,992	
... gen_z	509	17%
... millennial	1,211	40%
... gen_x	760	25%
... boomer	475	16%
... silent	37	1%
Gender	3,000	
... female	1,530	51%
... male	1,463	49%
... other	7	0%
LGBT	3,000	
... 0	2,662	89%
... 1	338	11%
Race	2,983	
... white	2,144	72%
... black	501	17%
... other	338	11%
Hispanic	3,000	
... 0	2,491	83%
... 1	509	17%
Education level	3,000	
... hs or less	798	27%
... some college	1,141	38%

(continued)

Variable	N	Percent
... bachelors	678	23%
... postgrad	383	13%
Household income	2,588	
... $< 25k$	466	18%
... 25-49k	775	30%
... 50-74k	653	25%
... 75-100k	438	17%
... 150k+	256	10%
Party ID (including leaners)	3,000	
... ind	411	14%
... dem	1,475	49%
... rep	1,114	37%
Party ID (excluding leaners)	3,000	
... ind	843	28%
... dem	1,235	41%
... rep	922	31%
Ideology	3,000	
... neither	861	29%
... liberal	1,074	36%
... conservative	1,065	36%
Registered voter	3,000	
... 0	332	11%
... 1	2,668	89%
Voted in 2020	3,000	
... 0	604	20%
... 1	2,396	80%
2020 Presidential vote	3,000	
... trump	976	33%
... biden	1,314	44%
... other	710	24%
Consumed beer at home in last 30 days	3,000	
... 0	287	10%
... 1	2,713	90%
Consumed beer somewhere else in last 30 days	3,000	
... 0	1,245	42%
... 1	1,755	58%
Purchases beer off-premises (i.e., at a store)	2,998	
... 0	132	4%
... 1	2,866	96%

(continued)

Variable	N	Percent
Is HH decision maker for beer purchases	2,866	
… 0	29	1%
… 1	2,837	99%

Differences in brand consumption by party

Figure C.1: Differences in beer consumption, by party

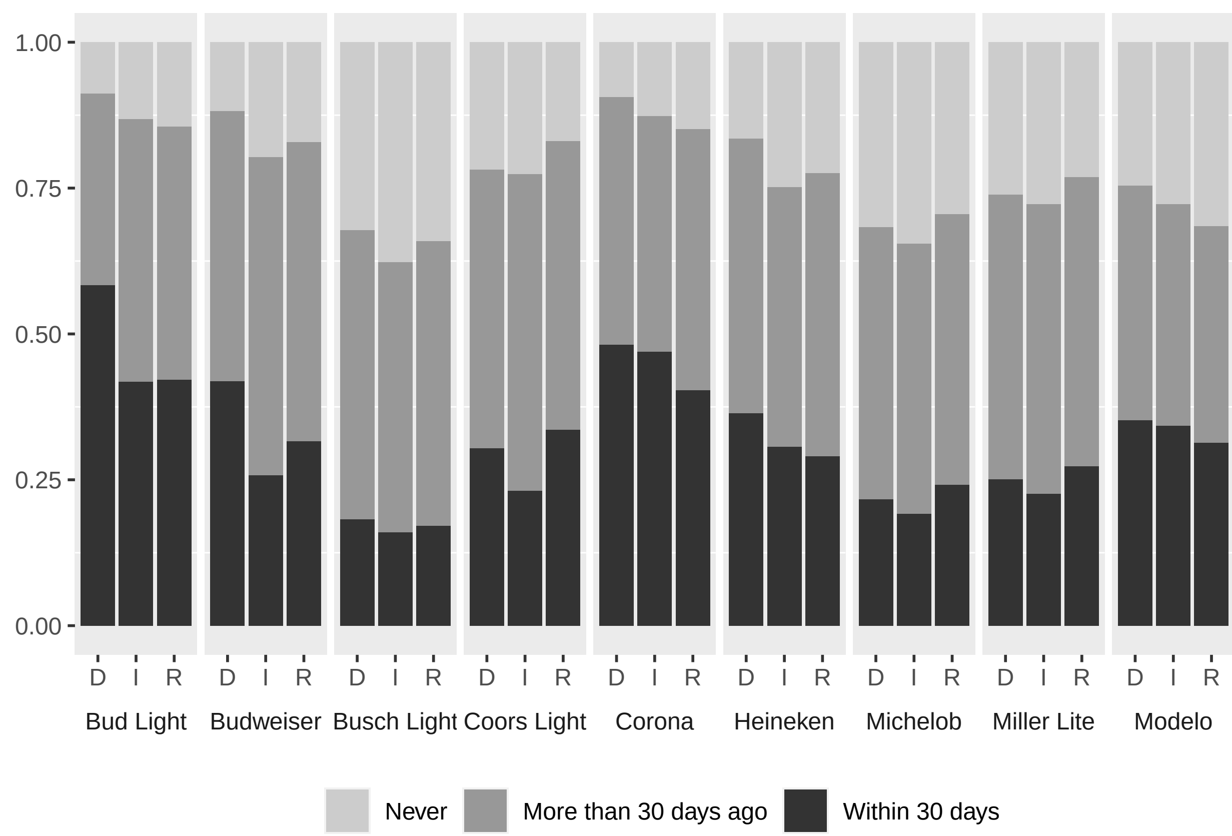

At the beginning of the survey (pre-treatment), respondents were asked about their consumption of nine of the best-selling U.S. beers. Overall, Bud Light was the most commonly-consumed beer. Nearly 90% of those in the sample said that they had consumed Bud Light at some point in their life, with 50% drinking it in the past 30 days. There were significant partisan differences in recent consumption, however: Republicans were over 16 p.p. less likely to have consumed Bud Light in the last 30 days, and nearly 6 p.p. more likely to claim they had *never* consumed the brand. Similar effects were seen in Budweiser.

There were other, smaller differences among other beers. Coors appeared more popular among Republicans, possibly due to the Coors family's association with Republican politics (the company's former chair, Peter Coors, ran unsuccessfully for Congress as a Republican). Heineken and Corona, two foreign beers, are more popular among Democrats in the sample. This may be due to differences in demographics between the parties. For instance, Hispanics in the sample were significantly more likely to drink Corona and Modelo (two Mexican beers), and nearly twice as likely to identify as Democrats.

Table C.2: Differences in beer consumption, by party

Beer	Party	< 30 days	> 30 days	Never
Bud Light	**Overall**	**0.501**	**0.384**	**0.115**
	Democrats	0.584	0.328	0.088
	Independents	0.418	0.450	0.131
	Republicans	0.422	0.434	0.145
Budweiser	**Overall**	**0.359**	**0.493**	**0.149**
	Democrats	0.419	0.463	0.118
	Independents	0.258	0.545	0.197
	Republicans	0.316	0.513	0.171
Busch Light	**Overall**	**0.175**	**0.488**	**0.337**
	Democrats	0.182	0.496	0.322
	Independents	0.161	0.462	0.377
	Republicans	0.171	0.487	0.341
Coors	**Overall**	**0.306**	**0.493**	**0.201**
	Democrats	0.304	0.477	0.218
	Independents	0.231	0.543	0.226
	Republicans	0.336	0.495	0.170
Corona	**Overall**	**0.451**	**0.430**	**0.119**
	Democrats	0.481	0.425	0.094
	Independents	0.470	0.404	0.127
	Republicans	0.404	0.447	0.149
Heineken	**Overall**	**0.329**	**0.472**	**0.199**
	Democrats	0.364	0.471	0.165
	Independents	0.307	0.445	0.248
	Republicans	0.291	0.485	0.224
Michelob	**Overall**	**0.222**	**0.465**	**0.312**
	Democrats	0.216	0.467	0.317
	Independents	0.192	0.462	0.345
	Republicans	0.241	0.464	0.294
Miller	**Overall**	**0.256**	**0.492**	**0.252**
	Democrats	0.252	0.487	0.261
	Independents	0.226	0.496	0.277
	Republicans	0.273	0.496	0.231
Modelo	**Overall**	**0.337**	**0.388**	**0.276**
	Democrats	0.352	0.403	0.245
	Independents	0.343	0.380	0.277
	Republicans	0.314	0.371	0.315

Figure C.2: Baseline preferences in conjoint

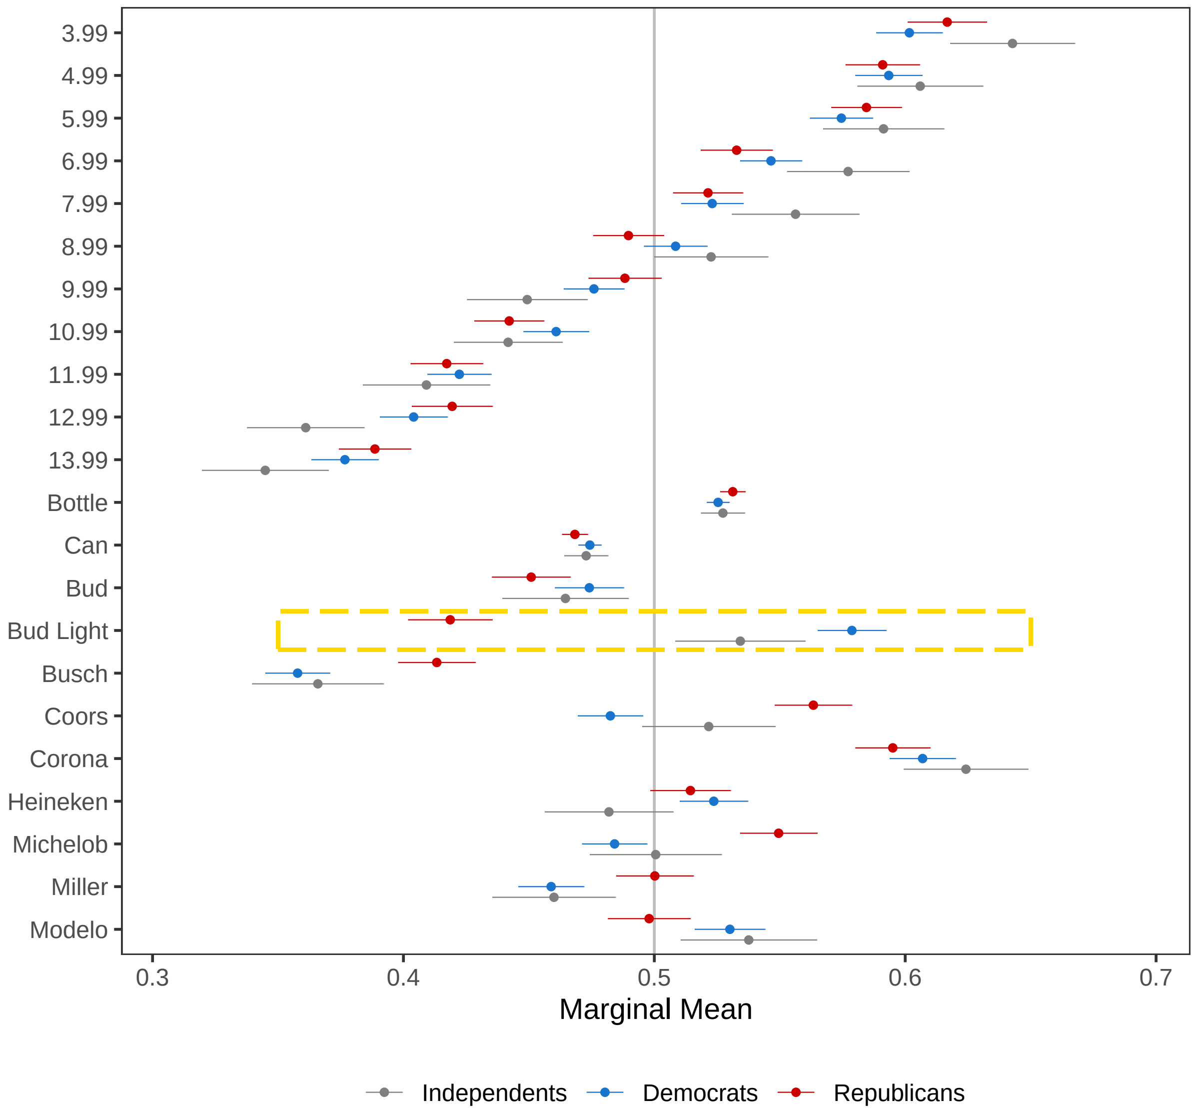

Marginal means split by party. Bars depict 83.4% confidence intervals. Preferences for beer brands largely do not differ across partisans, exception of Bud Light.T he results are not meaningfully different if we consider the control group only; see Table C.4.

Figure C.3: Baseline preferences (control only)

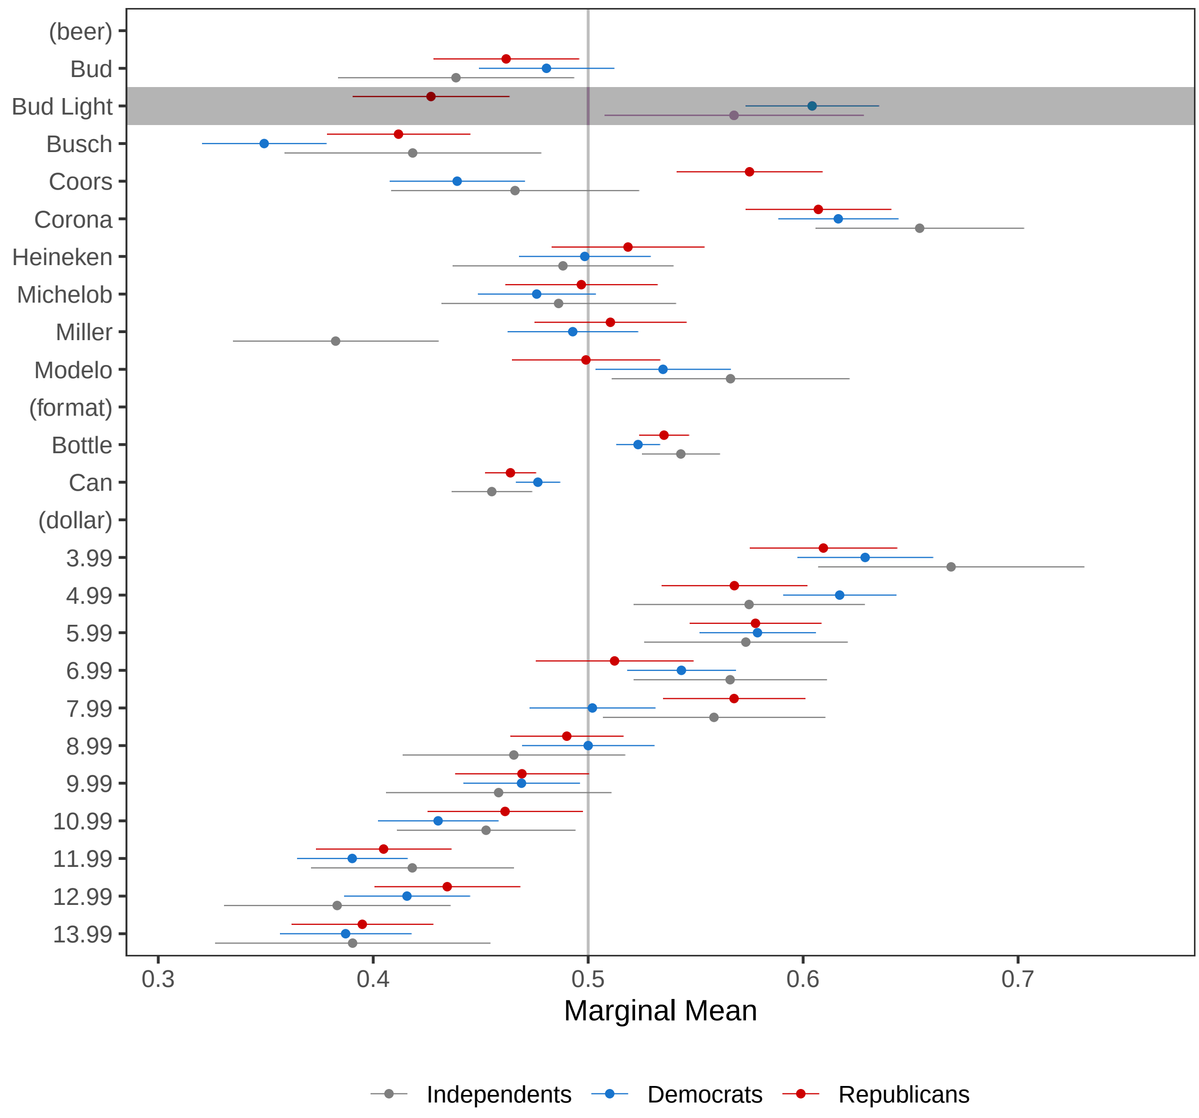

Figure C.4: Baseline preferences (control only). Bars depict 83.4% confidence intervals. Even without treatment, Republicans are significantly less likely to select Bud Light.

Figure C.5: Respondent ratings of brands

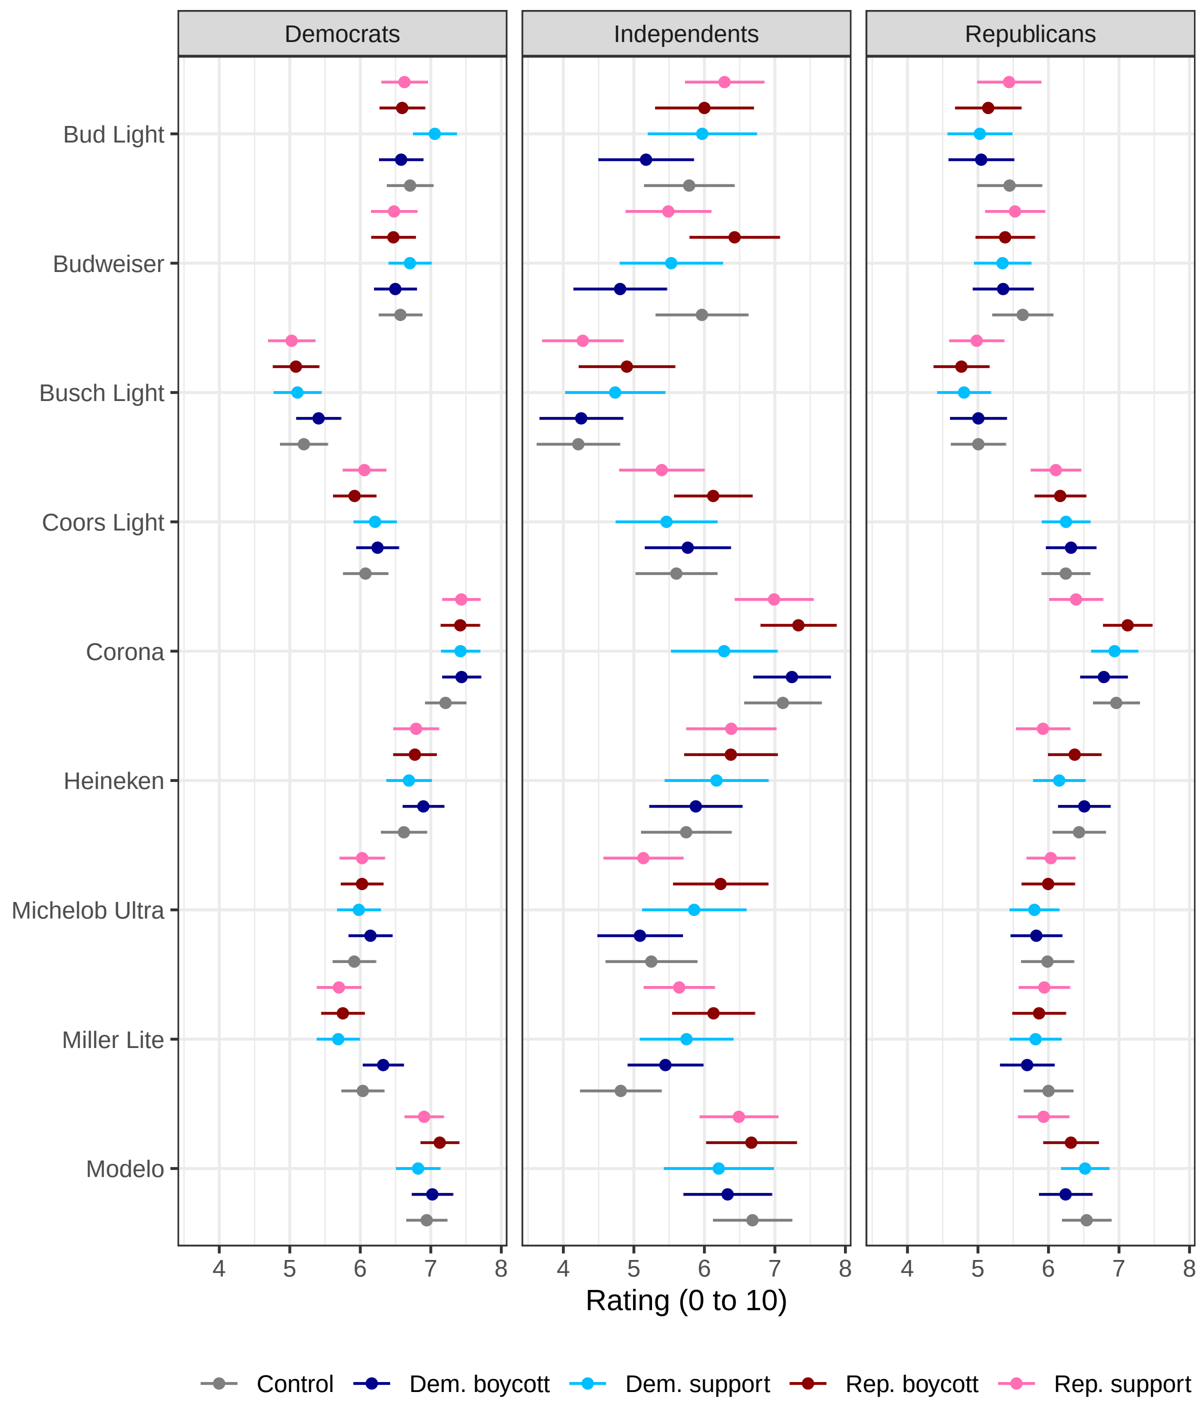

As a supplemental measure, respondents were also asked to rate each brand from 0 (low) to 10 (high). Although Republicans rated Bud Light lower than Democrats, there were no significant treatment effects. Democrats who received the "Democratic Support" treatment effect rated Bud Light slightly higher, but the difference was not significant.

Figure C.6: Treatment effects on politicians mentioned in vignette treatments

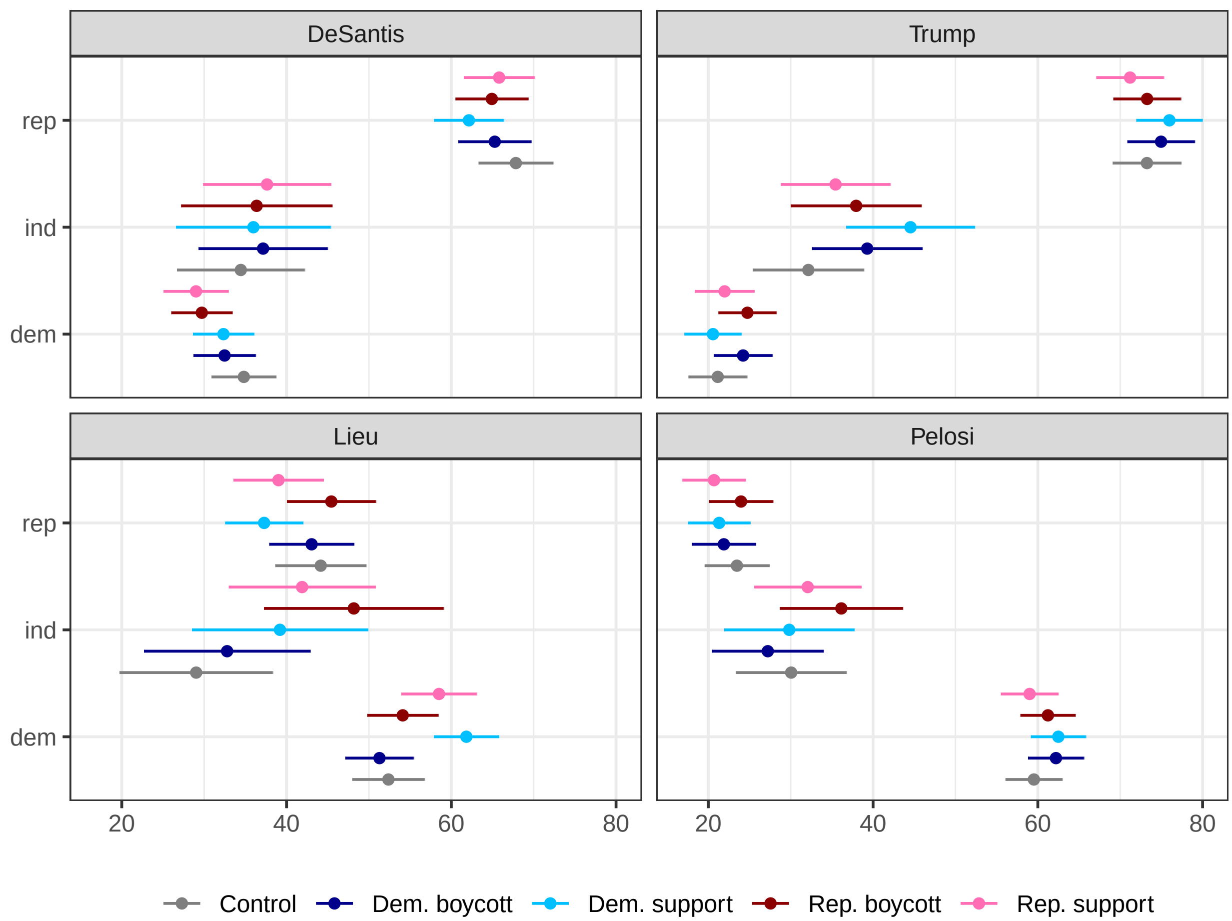

The Democratic Support treatment that mentioned Ted Lieu appeared to modestly increase Democrat's feeling thermometer rating towards Ted Lieu. The Republican Support treatment that mentioned Donald Trump, Jr., and the Republican boycott treatment that mentioned Ron DeSantis did not appear to affect attitudes towards these politicians; the were also no effects on the placebo (Pelosi) who has not mentioned in any treatment.