Political Preferences and Migration of College-educated Workers

Mitch Downey Jinci Liu*

September 4, 2023

Abstract

We study the consequence of political polarization along educational lines in the United States. First, we show an increase in the gap between college and non-college voters' policy views. Today, the average college graduate is far to the left of the average non-college voter on both economic and social issues, and to a degree much larger than 10 years ago. Next, we estimate the causal effects of a Republican governor on college graduates' choice of where to live. Republican governance reduces the in-migration flow of college-educated workers by about 13% per year over the four post-election years. This result changes over time in ways that closely mirror changes in political preferences, is robust to various identification strategies, and cannot be explained by labor demand. Finally, we extend a model of spatial sorting to allow workers to hold preferences over *political* amenities, and we calibrate the model to match our reduced form migration responses. We use the model to simulate various counterfactual changes in political control, with a particular focus on the inter-linkages between different states and the distributional consequences of the effects.

*Downey: IIES, Stockholm University. Liu: IIES, Stockholm University. We are grateful to Jay Lee for invaluable guidance and Simon Görlach, Seth Hill, Kieran Larkin, Kurt Mitman, Alessandra Peter, Giacomo Ponzetto, David Strömberg, Horng Chern Wong, Noam Yuchtman, and seminar participants at SU, IESEG, the UEA, the Barcelona Summer Forum, and the MYPEERS Workshop for helpful feedback.

1 Introduction

How do people choose where to live, and what are the consequences for the pattern of economic activity? Over the last 40 years, US college graduates have become increasingly concentrated in specific places (Moretti, 2012; Diamond and Gaubert, 2022). This is important because many high-growth sectors are skill-intensive, and college graduates generate economic spillovers that raise the wages of local non-graduates, too (Moretti, 2004). Therefore, increased spatial sorting has contributed to the slowdown of regional convergence (Kleinman, Liu, and Redding, 2023) and substantially exacerbated inequality in welfare between college and non-college workers (Diamond, 2016).

A separate literature shows that Democrats and Republicans increasingly live in different places (Brown et al., 2022; Kaplan et al., 2022). This is important because neighbors and neighborhoods do affect individual attitudes (Cantoni and Pons, 2022; Martin and Webster, 2020; Perez-Truglia, 2018; Perez-Truglia and Cruces, 2017), and so spatial sorting can exacerbate political polarization (Bishop, 2009; Brown and Enos, 2021). At the same time, it reduces democratic competition in these places, which further exacerbates the election of political extremists (Hopkins, 2017) and undermines economic growth (Besley, Persson, and Sturm, 2010).

In summary, geographic sorting by education and geographic sorting by political attitudes are both important social, political, and economic phenomena. In this paper, we show how they are related. Divisions in political attitudes increasingly fall along educational lines. In 2020, the college/non-college gap in Biden/Trump voting was as large as the gap between New York and Mississippi. This is part of a long-running worldwide trend in the realignment of political coalitions (Gethin, Martínez-Toledano, and Piketty, 2022), but we show below that the recent growth has been dramatic. We also show that in this period, college graduates are increasingly reluctant to move into Republican governed states, reducing the human capital stock of those states. Our main goal in this paper is to document this effect and quantify its consequences for the pattern of economic activity.

We begin by using state-of-the-art tools borrowed from political science to create comparable indices of policy views on economic and social/cultural issues over the last 15 years. We find that the gap in policy views between college and non-college voters has grown dramatically since 2010. On economic issues, as recently as 2010, there was *no* college/non-college gap in policy views, while today, the average college voter is .4 standard deviations to the left of the average non-college voter. On social and cultural issues, while the gap shrunk during the later years of the Bush Administration and was roughly stable from 2010-2015, it more than doubled from 2015-2020 (to .5 standard deviations). We show that this implies that higher earners are, on average, far to the left of lower earning voters on both social and economic issues.

We next estimate the migration effects of transitioning from a Democratic governor to a Republican one, paying special attention to how these effects have changed over time. As we discuss below, we focus on governors because they are very salient and have a large influence on policy outcomes. During the recent period, amid heightened polarization across education lines, we find that a Democrat-to-Republican transition leads to a large (13% per year) decline in the inflow of college graduates into the state.[1] Put differently, in the average state, around 3% of the college-educated workforce lived in a different state during the previous year. Our estimates imply a 0.4 percentage point decline in the migration flow, which over four years implies a 1.6% decline in the stock of college-educated workers. This magnitude is roughly equal to one year of growth in the average state's college-educated workforce.

We find the same effects from an IV strategy based on the differential timing of gubernatorial elections across states, as well as from a regression discontinuity design, although both are less precise and not statistically significant. We find no effects on non-college workers (who don't show the same aversion to conservative policies), and the pattern of effects that we find on job openings and hiring rule out a labor demand explanation. Remarkably, the over-time pattern we find for migration effects almost perfectly mirrors the non-monotonic over-time pattern that we find for college/non-college gaps in views on social/cultural issues.

Overall, then, we conclude that our migration effects reflect labor supply responses as college graduates are increasingly reticent to live under Republican governance. This idea is not new.[2] Conservative states spend millions advertising low taxes in liberal states (see Moretti and Wilson (2017) for some discussion). As an example, Republican-controlled Indiana ran a billboard in neighboring Democratic-controlled Illinois asking whether residents were "Illinoised by higher taxes" and encouraging them to "Come to Indiana: A state that works." At the same time, however, Indiana Governor Mike Pence was signing into law one of the most controversial anti-LBGTQ laws in the nation, and one of Indianapolis' Democratic City Councilmembers noted "The real harm, to all [of us], is that it makes us all look like backwater hicks. Indiana is losing jobs and young professionals like crazy. How much more can our state government make Indiana uninviting" (Eason, 2015). It is worth noting that his rhetoric is remarkably similar to the rhetoric conservative policymakers use when advocating for tax cuts. Our paper shows

[1] These estimates are identified by states that switch their governors. Thus, they reflect local average treatment effects. They say little about how politics shapes the choice to move to the large set of states never observed switching gubernatorial partisanship during our period, but they are the ideal estimates to inform our structural counterfactuals, where we consider flips among relatively evenly contested swing states like these.

[2] In a 2017 survey, Smith, Hibbing, and Hibbing (2019) find that one-in-four respondents report that politics had led them to consider moving. The challenge of attracting college graduates to conservative states has garnered considerable attention since the Supreme Court's *Dobbs* decision gave states more authority over abortion laws (Cain et al., 2022; Hagelgans and Basi, 2022; Keshner, 2021; Leonhardt, 2021). Our results show that this challenge is real, but broader than abortion laws alone.

that this councilmember was correct and, in a sense, that conservative policies are an implicit preference tax on a large share of college graduates.

To understand the consequences of this, we analyze a structural model of migration that builds on Bryan and Morten (2019). We extend the model in two ways. First, we add two types of workers – college and non-college – who differ in wages, migration costs, and valued amenities, and are combined in state-level production. Second, we add a "political amenity" which differs for college and non-college workers and depends on the partisanship of the governor. We estimate most parameters following the approach developed by Bryan and Morten (2019), and calibrate the political amenity to match our reduced form results. We use the calibrated model to simulate various counterfactuals reflecting plausible shifts in partisan control, and focus on two types of general equilibrium forces that would be difficult or impossible to study using reduced form methods.

First, we study how these labor supply responses affect inequality. Reducing in-migration of college graduates has three separate effects on college/non-college earnings inequality. On the one hand, reducing the relative supply of college graduates raises their relative wages (a price effect). At the same time, the model implies that the marginal migrants deterred from moving should be less positively selected than the inframarginal migrants who are not deterred. That is, the model includes selection effects, and increasing these selection forces should also raise the average earnings of college graduates. On the other hand, all migrants are positively selected relative to non-migrants, so reducing the share of migrants among college graduates reduces the average productivity of college graduates and thereby reduces inequality.

We simulate various plausible changes in political control, and across the states we consider, none of these forces unambigously dominates. Sometimes we estimate that decreasing the number of college graduates will raise their earnings relative to non-college workers (the intuitive effect of a labor supply shock), but sometimes we estimate that it will reduce their relative earnings. This is an interesting result because it implies that in some cases, Democratic governance can simultaneously raise the college share of the workforce *and* increase college/non-college earnings inequality, entirely through labor supply behavior.

The second set of general equilibrium forces we focus are the cross-state spillovers that are inherent to studying migration. Deterred would-be migrants choose to live somewhere else, so changes in the incentives to move to one state ripple throughout other states. Our model lets us quantify these spillovers.

At the broadest level, our results speak to the real-world consequences of political polarization. Despite the widely held view that polarization is an important social phenomenon, there is little evidence of its effects on tangible economic outcomes.[3] As US politics grows more

[3] A noteworthy exception is McConnell et al. (2018), who experimentally show that workers require significant

hostile and divisive (Boxell, Gentzkow, and Shapiro, 2020), it is important to understand its consequences beyond purely political outcomes.

Our results also contribute to work on the role of education in political divides. Recently, Gethin et al. (2022) combine data from 70 years and 21 countries to show that increased college graduates' support for left-of-center parties is a part of a universal and long-run trend. This suggests that the forces we focus on are likely to continue in relevance, regardless of short-term changes in the influence of specific candidates or issues. It also illustrates that these forces are relevant beyond the United States.[4]

The rest of the paper is organized as follows. Section 2 documents our descriptive results. Section 3 estimates the causal effects of governors on migration. Section 4 lays out our model, Section 5 our strategy to identify and estimate key parameters, and Section 6 our results. Section 7 concludes.

2 Education polarization

We begin by documenting trends in "education polarization," by which we mean the difference in policy preferences of college graduates and non-college graduates.

2.1 Data and methods

Our main data is the Cooperative Election Study (CES) operated by Harvard University and conducted by YouGov. The CES began in 2006 with roughly 35,000 respondents, and today includes roughly 60,000 respondents per even-numbered year (much larger than any similar survey). For our purposes, the key feature is that the CES asks a large number of policy questions about diverse topics, with an eye towards comparability over time. From its inception, the CES has asked 15-50 different policy questions during each wave, far more than other political surveys.

Using these policy questions, we create indices of policy preferences on social and economic policies. We classify questions using the topic lists developed by Caughey and Warshaw (2018); Appendix Tables A1 and A2 contain details.[5] While the specific set of covered issues changes from year to year, questions are always chosen to capture the central issues in political debates. Thus, trends in our indices don't necessarily reflect changes in views on the exact same questions,

compensating differentials to work for employers whose political views oppose their own.

[4]For instance, Brox and Krieger (2021) show that far-right protests in Dresden, Germany, led to a decline in in-migration of college graduates there.

[5]Social issues include gun control, abortion, immigration, and policing; economic issues include financial regulation, environmental policy, taxation, and government spending. The average survey wave includes 13.5 social policy questions and 13.9 economic policy questions.

but changes in views on the questions that are important to political divides at the time. We see this as an advantage,[6] although it means we cannot distinguish between changes in specific views and changes in the specific issues debated within the broad classification of "social" and "economic" issues.

To estimate indices, we follow the best practices in political science and estimate ideal points using Item Response Theory (IRT). IRT was developed for use in standardized testing, where not all test takers receive the identical test and not all questions are equally informative about underlying ability. It is widely used in political opinion research because (1) it does not require the analyst to specify, ex ante, which questions are more or less informative about ideology, and (2) unlike other methods (such as principal component analysis), it generates comparable scores across respondents even when the set of questions changes over time, so long as there are "bridge" questions that are continuous from one survey wave to the next (see Caughey and Warshaw (2015) for some discussion). For this reason, the CES intentionally includes bridge questions in each survey wave. IRT estimation does not require assumptions about the sign of individual questions, but researchers must normalize the final index. We follow the common practice and normalize the index so that more negative values indicate more liberal (left-leaning) views and more positive numbers denote more conservative (right-leaning) views. The index is normalized to be mean zero and unit standard deviation across all respondents across all years. We estimate our indices using an Expectations Maximization algorithm (Imai, Lo, and Olmsted, 2016).

Because our interest is in education and migration, we focus on respondents aged 26-45. We focus on respondents over age 26 to ensure that college education has been completed.[7] We end our range at age 45 because of a steep decline in rates of inter-state migration (see Appendix Figure B1). In general, all of the trends that we document are similar when we include the full population; older citizens have different *levels* of views, but the trends are similar.

2.2 Differences in policy preferences by education

Figure 1 presents our main results about "education polarization" (i.e., the gap in policy views between college and non-college respondents). Panels (a) and (b) show the evolution over time of college and non-college respondents' on social and economic policy views.

[6] For example, in 2006 there were major political divides over whether school science classes should teach intelligent design and whether social security should be privatized. By 2020, these issues were irrelevant; voters instead debated whether schools should require face-masks and whether Medicare should be universal. Changing the questions over time is necessary for the index to reflect actual political disagreement. Of course, some of the more permanent divides in US politics (such as questions about tax rates and government spending) do appear continuously in each CES wave.

[7] We only focus on Bachelor's degree completion and never aim to separate those with graduate degrees from the Bachelor-only population.

On social issues, both college graduates and non-college have drifted left considerably over time, consistent with widely recognized changes in views on issues like divorce, gay marriage, and marijuana legalization. But while college graduates have always had more liberal views, the gap hasn't been constant. Panel (c) shows that the gap fell during the end of the Bush Administration (2006-2008), and was roughly stable from 2008-2015, though it has been growing steadily since the candidacy/election of Donald Trump.[8] From 2015 to 2020, it more than doubled from $.2\sigma$ to roughly $.5\sigma$.[9]

[Figure 1 about here.]

Panel (b) shows a somewhat different pattern on economic policy issues. Here, there was historically no gap in average views, but since 2010, college-educated citizens have moved considerably left (by $.2\sigma$-$.3\sigma$), while non-college voters tend to have moved modestly to the right. Panel (c) shows that the gap in views has grown to roughly $.4\sigma$.

One important insight from Figure 1 is that while educational gaps in views on social and economic issues started at somewhat different *levels*, the trends in both are similar and today gaps in both are equally large. Moreover, Panel (d) shows that, across individuals, the correlation between these views has increased from roughly .5 in 2010 to roughly .75 in 2020. For these reasons, when we estimate migration responses to Republican governors, we never attempt to separate between economically and socially conservative governors, which is useful because there is no existing dataset that would allow us to do so.

While differences in *average* views are helpful, they are necessarily a simplifying summary. In Appendix Figure B3, we plot the full distributions of social policy preferences for college and non-college respondents in 2010 and 2020. The change in distributions is striking. Over this period, although the median Democrat moved 0.2σ further to the left, there was a 5 percentage point (12%) increase in the share of college graduates left of the median Democrat, and a 5pp (18%) decrease in the share of non-college voters who were. Thus, college graduates went from being 50% more likely than non-college to fall left of the median Democrat to being 120% more likely.

[8] We see similar trends when we look at the gap in *median* (rather than mean) views in Appendix Figure B2, which is important because the ability to distinguish extreme views has changed over time as more policy questions have been added to the CES.

[9] Our aim in this paper is not to determine the *source* of education polarization, but the coincidence with the timing of Donald Trump's presidential campaign is noteworthy. Various commentators have argued that changes in college-educated political views and campus political climate pre-dated Trump's campaign (Haidt and Lukianoff, 2021; Yglesias, 2019).

2.3 Differences in policy preferences by earnings

In general, states care about attracting and building a skilled, productive workforce, and education is correlated with measures of productivity like earnings. However, it is possible that the college/non-college gaps we see are driven by a subset of left-leaning graduates who are have low earnings. In this case, a conservative policymaker might be unconcerned about these workers away, and reasonable minds can disagree about the costs to economic growth. If, on the other hand, the gap in policy views holds broadly when comparing high- and low-earners, then virtually all policymakers would be concerned about the consequences of pushing away these workers.

One challenge in assessing this is that the CES does not collect information on individual-level earnings or income.[10] For most recent years, however, it does collect respondents' sector of employment (2-digit NAICS), in addition to education and age.[11] This information is sufficient to impute earnings with high accuracy. To do so, we rely on an approach that adjusts earnings for geography. Coastal areas and large cities are more liberal and have higher nominal wages, and we would not want to conflate political differences by earnings with a spurious correlation between regional earnings differences and local political dynamics. Using the American Community Survey (ACS) we regress log annual earnings on geography-by-time fixed effects and age-by-education-by-industry fixed effects:

$$\ln(y_i) = \alpha_{m,s,t} + \theta_{a,n,e} + \varepsilon_i \tag{1}$$

where i denotes individual, m denotes MSA, s denotes state, t denotes year, a denotes age range (25-32, 33-40, 41-55), n denotes industry (2-digit NAICS), and e denotes education (college graduate vs. non-college). The MSA-by-state-by-year fixed effects ($\alpha_{m,s,t}$) account for heterogeneous levels and trends in nominal earnings in a flexible way, and the age-by-industry-by-education fixed effects ($\theta_{a,n,e}$) therefore represent persistent earnings differences between different types of workers, regardless of where they live. We see these estimated $\hat{\theta}_{a,n,e}$ fixed effects as CES respondents' *earnings capacity*, which is relevant for a state attempting to attract high productivity workers.[12] In this way, we determine the earnings capacity for all employed CES respondents (ignoring non-employed respondents) and normalize our measure to have unit standard deviation across all employed respondents across all years.

[10]The CES *does* collect information on household income, but *i*) it is collected only for relatively broad income ranges, and *ii*) household income includes income from both partners in married or cohabiting families, and patterns of assortative matching and marriage rates differ by education and have changed over time.

[11]Industry is collected in 2011-2014 and 2016-2020.

[12]One advantage this approach relative to simple average earnings is that changes over time in the relationship between policy views and earnings capacity reflect changes in different groups' views, holding fixed the earnings position of any given "group" (i.e., age-industry-education cell).

Figure 2 shows the basic patterns of earnings polarization. In Panel (a), we plot the coefficients on our earnings measure from year-by-year regressions of social and economic policy views on earnings capacity (again for workers age 26-45). Earnings had little correlation with policy views through 2014, though in 2011 higher earners were slightly more conservative on economic policy issues. But since then, higher earners have become substantially more left-leaning on both social and economic issues. By 2020, a 1σ increase in earnings capacity is associated with a $.12\sigma$ shift left in policy views on both sets of issues.

This relationship is entirely driven by education. Panel (b) presents coefficients that come from regressing views on both earnings and education. Once we control for education, high earners are more conservative on both economic and social policy issues, and this relationship has not changed at all over time. The trends in "earnings polarization" arise entirely from the college/non-college gap which, once we control for earnings, has increased from roughly $.25\sigma$ in 2011 to $.55\sigma$-$.65\sigma$ in 2020.

[Figure 2 about here.]

Finally, panels (c) and (d) show the distribution of earnings capacity and 2020 policy views across all age-education-industry cells. There is a striking separation between college and non-college workers, with almost zero overlap. At virtually any point in the earnings distribution, all types of college graduates are to the left of all groups of non-college workers, on both economic and social policy issues. The figure shows that differences between education groups are dramatically larger than any differences within the college and non-college population, at least among groups that we can measure in the data.[13] For this reason, when we study migration responses to gubernatorial partisanship, we focus only on differences by education and do not seek to explore differences by age or industry.

2.4 Explaining education polarization

Our goal in this paper is not to explain the rise of education polarization. Nonetheless, some features of our data do speak to some potential explanations. In Appendix Section B.2.1

[13] Earnings and political views do, of course, differ within age-industry-education cell, and we can say nothing about the within-cell correlation in views. Existing causal evidence from the University of California's admissions policies, however, shows that attending a more prestigious university increases earnings (Bleemer, 2021) *and* pushes students further to the left (Firoozi, 2022). These effects are consistent with the correlation across all universities, which shows that students at more prestigious universities are generally further to the left (Firoozi, 2022). Thus, it is plausible that within these cells, higher earning college graduates are even further to the left of their peers, although we cannot test this in our data. In other data, we can test differences across college majors, and there is very little correlation across majors between earnings and political views (see Appendix Figure B13).

we provide suggestive evidence on three prominent explanations for the growth in the gaps between college and non-college voters.

First, we show that some of the change is driven by a growth in the gaps between successive cohorts, which is consistent either with universities having stronger "treatment effects" over time, or with changes in selection into university attendance over time. However, these between-cohort effects are much smaller than the within-cohort over time changes, so these forces are not important for explaining the aggregate patterns. Second, one prominent explanation is that political debate has shifted from economic issues to social issues (Gethin et al., 2022). We show that views on social issues have become more important (and those on economic issues less important) for explaining party preferences. However, these effects are also fairly small, and above we showed that gaps in views themselves (and not simply the priority given to different types of views) have changed considerably. Third, we find some evidence that the states which saw the greatest growth in college share of the workforce also saw the largest change in education polarization, which is consistent with the idea that increased sorting of like-minded individuals reinforces polarization (Bishop, 2009), but this effect is also fairly weak.

Overall, then, we find some modest support for three common claims, although none offers a complete explanation. Indeed, we do not believe there *is* a singular explanation for the rise in education polarization. For this reason, we do not seek to model endogenous polarization.

2.5 Strength of policy preferences

We have shown large and growing differences in the views of college and non-college voters. We have, thus far, said nothing about the *strength* of those views. It is possible that different groups hold different views, but that these views are only weakly held and unlikely to shape actual behavior. While it is more difficult to quantitatively assess evidence on the strength of preferences, here we summarize a range of evidence illustrating that many US citizens, and particularly college-educated liberals in recent years, care deeply about politics.

First, Americans devote substantial time to politics. Standard time use datasets do not collect information on political activity or media consumption, but Hersh (2020) fielded a nationally representative survey in 2018 which found that one-third of Americans averaged more than two hours *per day* following politics. College-educated respondents and Democrats were both more likely to fall into this group. In a 2017 nationally representative survey, Smith, Hibbing, and Hibbing (2019) find that 26% of Americans spend more time thinking about politics than they would like, and 17% report their lives would be better if they focused on politics less.

Second, Americans spend substantial money on politics. In the 2020 CES, among college graduates aged 26-45, 29% report contributing to a political campaign, compared to 9% of non-college respondents.[14] One-quarter of these contributors spent more than $300. A classic perspective in scholarship on campaign finance is that individual contributions are "consumption" rather than any sort of investment (Ansolabehere, de Figueiredo, and Snyder, 2003). This appears even more true during the recent period. Candidates routinely raise millions of dollars – largely from out-of-state donors who cannot vote in the election – by challenging influential conservative senators, only to lose by large margins.[15]

Third, Americans report that politics is having meaningful effects on their mental health. Smith et al. (2019) find that politics has led 38% of survey respondents to experience stress, 18% to lose sleep, 20% to damage valuable friendships, 12% to experience adverse effects to physical health, and 4% to consider suicide. The battery of such health problems is more severe among younger, better educated, more left-leaning, and more politically engaged respondents. These are self-reported associations between politics and mental health, but experimental evidence shows that exposing participants to mainstream political news (the type heavily consumed by college graduates) causally increases stress and reduces psychological well-being (Ford et al., 2023). Accordingly, recent editions of the American Psychological Association's annual "Stress in America" report have focused extensively on politics.

Finally, it is worth noting that our evidence on differences in political views may well understate polarization. Research in political psychology emphasizes that even modest differences in policy views can lead to strongly held partisan identities that reinforce hostility towards the opposite party (Mason, 2015). This can lead to "social polarization" (which individuals feel and act upon) vastly greater than observable "issue polarization" (which we can measure in our data), and it is particularly likely to occur when gaps fall along salient, easily observable lines (Mason, 2015), such as educational attainment. In summary then, politics has become a dominant force in individuals' lives, personal identity, and mental health. We next turn to whether people respond to politics in deciding where to move.

3 Migration responses

Our main interest is in how college-educated workers respond to changes in the partisan control of state government, and how this response has changed over time. For this reason, we

[14]Bouton et al. (2022) estimate there were roughly 20 million independent individuals making federal campaign contributions in 2020 (roughly 10% of the adult population), with 8 million contributing more than $200.

[15]For instance, in 2020, Jamie Harrison raised $132 million challenging Senator Lindsey Graham in South Carolina, and Amy McGrath raised $96 million challenging Senator Mitch McConnell in Kentucky. They lost by 10.2 and 19.6 points, respectively.

focus on the partisanship of the governor. It has long been recognized that citizens tend to know more about their governor than other elected officials, and that governors garner more media attention (Hinckley, Hofstetter, and Kessel, 1974; Squire and Fastnow, 1994). Even today, although politics has nationalized, 78% of 2020 CES respondents could correctly identify the partisanship of their governor, compared to only between 45-60% for their Senators or Representatives in the House.

Moreover, governors have significant effects on state policy outcomes, and the size of these effects has grown over time (Caughey, Xu, and Warshaw, 2017).[16] In part, this is because governors face very few checks and balances or other constraints on authority and the prevalence of these checks has declined over time (Seifter, 2017), in part it is because they use their veto power much more actively than the President does, and in part it is because they are typically the *de facto* leaders of their political party within their state, giving them responsibility for setting agendas and driving legislation. We quantify their relevance for policy using the difference-in-difference strategy discussed below and updated data from Caughey and Warshaw (2016) – who create policy indices based on 148 separate state-level policies – and we find that a Republican governor leads states' economic and social policies to become .2-.3σ more conservative (see Appendix Figure B6). At the same time, given the salience of governors and their high-profile stances on controversial issues, we do not expect effects to operate *solely* through policy. We make no attempt to separate between governors' policy effects and other high-salience political theatrics.

3.1 Data and methods

Our primary migration data is from the American Community Survey (ACS), a nationally representative 1% sample of the US population conducted by the Census Bureau, accessed via IPUMS (Ruggles et al., 2022). Microdata includes respondents' state-of-residence, age, education, employment details, and last year's state-of-residence. This allows us to calculate annual inter-state migration rates by education, and we mostly focus on those holding at least a Bachelor's degree. Relative to studies using state of birth and residence to measure migration (e.g., Bryan and Morten (2019)), a key advantage is that we observe the specific timing of the move.[17] We use the ACS back to 2000 (when it was introduced), although from 2000-2004 the

[16]In an influential paper on *mayors*, Ferreira and Gyourko (2009) use a regression discontinuity to show that partisanship does not matter for policy outcomes. Using updated and expanded data, de Benedictis-Kessner and Warshaw (2016) show that it does. Moreover, the same regression discontinuity design shows that *governor* partisanship is consequential for policy (Caughey et al., 2017).

[17]Some migration studies use the IRS migration files. This data does not contain education, and the college/non-college earnings distributions are not sufficiently distinct to impute it. Other recent work studies migration using voter registration data or change of address data (Brown et al., 2023). This data also does

sample size was considerably smaller than the 1% sample design used since 2005.

Our primary focus is on migration flows, calculated separately by education. We define migration inflows into state s in year t as the number of respondents living in state s in year t who report living in a different state in year $t-1$. We always use ACS sample weights when calculating migration flows. Throughout the paper, we never analyze migration from 2020 onwards because we are concerned that during the main pandemic years, respondents' state-of-residence is either not well defined, not well measured, or is only temporary. Thus, we always end our migration analysis window in 2019. Finally, our primary analyses focus on respondents who are employed in the private sector (and thus most relevant for economic activity), citizens (and thus more engaged in politics), and aged 26-45 (where inter-state migration is concentrated). We also present estimates using other samples, and the figure we use for calibrating our structural model is based on all private sector employees age 26 or older.

With these annual migration flows, we estimate the effects of gubernatorial transitions on migration. Our primary analyses use a difference-in-difference specification, although we show similar point estimates and time trends using a regression discontinuity specification and an instrumental variables specification where we instrument for outcomes using variation across states in the timing of their gubernatorial election cycle. Since gubernatorial terms in 48 of the 50 states are four years long, we focus on four-year treatment effect windows.

Switches in governor partisanship (what we call "flips") occur in different places at different times and are thus a form of "staggered rollout" problem that has received attention in the recent econometrics literature. We use a Callaway and Sant'Anna (2021) estimator, in which each state experiencing a flip is paired with other states sharing the same pre-flip partisanship, but which did not experience a flip during the five years before and after the treated state's flip. For example, in 2012 North Carolina elected Republican Governor Pat McCrory to replace the outgoing Democratic governor. McCrory took office in 2013. At the time, there were six other states which had Democratic governors in office during the full ten years surrounding this election (i.e., 2008-2017). We estimate the effects of McCrory taking office on migration into North Carolina by comparing the post-McCrory change in migration inflows into North Carolina with the simultaneous changes experienced in these six other states.

As Callaway and Sant'Anna (2021) show, this yields a consistent estimate of effects of North Carolina's 2013 flip, avoiding the documented concerns about negative weights in staggered rollout designs (e.g., De Chaisemartin and d'Haultfoeuille (2020), Goodman-Bacon (2021)). Adding other states experiencing flips during the same year yields a consistent estimate of the average treatment effect for all of those states, which Callaway and Sant'Anna call a "treatment cohort." Weighting the control states for each cohort so that their collective weight is propor-

not contain education, and voter registration data only includes partisanship for around half the country.

tional to the number of treated states in the cohort and then combining multiple treatment cohorts into the same regression yields the average treatment effect (ATE) of a given type of flip. Our main approach is to use this pooled estimator, combining different sets of treatment cohorts to study dynamic effects. Unless otherwise noted, we always cluster standard errors at the state level to account for serial correlation in migration and the fact that the same state-year can appear as controls for different treatment cohorts. Finally, we always weight state-year observations by the number of ACS respondents.

When we aim to summarize a single number, we regress log migration flows on a standard post-treatment dummy. That is, we use a standard two-way fixed effects estimate,[18] but after using the Callaway-Sant'Anna approach to ensure that we are only comparing treated states to untreated states. In general, however, we prefer to compare the levels of our outcome variables between treated states and control states. We normalize each state's level by dividing by its average level during the four pre-election years, and then plot normalized outcomes across event time. This allows a more transparent inspection of the dynamics of migration around the election.

Figure 3 shows the number of flips over time, separately for Republican-to-Democrat and Democrat-to-Republican transitions. One notable fact that shapes our analysis is that during recent years, D-to-R transitions have been much more common than R-to-D transitions. For instance, since 2015 (when we observe the dramatic increase in education polarization) there have been eight transitions from Democratic to Republican governors, and only three transitions the other way around. Thus, for recent years, we are naturally better powered to study D-to-R transitions than R-to-D transitions.

[Figure 3 about here.]

3.2 Main results

We begin by focusing on the effects of a transition from a Democratic to a Republican governor during the recent period where the increase in education polarization has been concentrated. Specifically, we consider flips where the new Republican governor took office in 2015-2017 (since we see migration data as unreliable from 2020 on).[19] Figure 4 shows that college graduates' migration into the state falls considerably upon the election of a Republican governor. Prior to the election, migration patterns were extremely similar between the eight treated "flip states"

[18] $\ln(Mig)_{s,t} = \alpha_s + \delta_t + \beta\left(Flip_s^{DR} \times Post_t\right) + \varepsilon_{s,t}$

[19] Technically, this is an unbalanced panel. As Figure 3 shows, this sample includes three flips in 2017 where, because we do not use mid-pandemic 2020 migration rates, we observe only three (2017, 2018, 2019) and not four post-election years.

and the 11 states that remained Democrat-run. After the election, however, in-migration falls by nearly 20% in those states relative to the control states.

[Figure 4 about here.]

Interestingly, when we examine the source of migration, we find that the effects are completely driven by in-migration from Republican-led states. These results are shown in Figure 5. Leading up to the election, both treatment states and control states saw in-migration flows from these states increase by about 25%. When treated states elect a Republican governor, this growth halts almost completely, while control states see it continue to rise by another 25% of the next four years. In the appendix, we present various additional results where we find no responses among non-college workers, no response of out-migration to gubernatorial flips, and no effects of Republican-to-Democrat flips.[20]

[Figure 5 about here.]

It is not unreasonable for a reader to be concerned about adverse economic shocks biasing our estimates. After all, economic shocks have effects on both election outcomes and migration. Below, we present a series of robustness checks to rule out various labor demand explanations for our results. Here, we simply note that *i*) 2015-2019 (our post-election window) was a period of broad economic growth around the country, shared by all regions and all demographic groups and *ii*) given the immediate timing of college graduates' response, these economic shocks would have had to precisely coincide with the election. Below, we present more systematic evidence.

First, however, we focus on changes in migration responses over time. Above, we showed that the size of the gap in political views between college and non-college has changed in important ways over the last 20 years. Here, we investigate whether migration behavior shows the same dynamics. To do so, we split our 15-year sample into five three-year periods: flips occurring 2003-2005, 2006-2008, 2009-2011, 2012-2014, and 2015-2017. For each period, we estimate the effects of a D-to-R flip on in-migration of college-educated workers. These period-specific estimates are plotted in Figure 6 along with the average level of education polarization along social issues during the post-election years. The relationship is striking. In the mid-aughts, when the college/non-college gap in views was large, college graduates' migration responded strongly to a Republican's election as governor. During the middle three periods, when the

[20] See Figure B8. It is difficult to say whether our null effects for non-college workers are because they care less about politics or because base rates of migration are much lower. It is also difficult to interpret our null findings on R-to-D transitions since we have far fewer R-to-D transitions than D-to-R transitions during this period (see Figure 3). We interpret our null effects on out-migration as evidence that elections do not affect individuals' decision about *whether* to move, only *where* to move, and note that Monras (2020) shows that economic shocks, in general, only affect in-migration and not out-migration.

college/non-college gap was relatively small and stable, migration did not respond to these elections. During the recent period, when education polarization has risen again, migration has again become quite responsive.[21]

[Figure 6 about here.]

The above estimates are based on private sector employed citizens age 26-45. However, these workers are not the full set of a state's college-educated workforce, and for our structural model it is the total change in college-educated labor that is important. Thus, Table 1 estimates our main difference-in-difference specification for less restrictive samples of college-educated workers. In our calibration, we match our model-generated migration behavior to that of all private sector employees aged 26 and above (column 2), where we estimate a Republican governor reduces annual migration flows by 14 log points (13%). In the appendix, we show corresponding plots for our estimated effects for these alternative samples of workers (Figure B9) as well as for the full sample of years (Figure B10).

[Table 1 about here.]

How large are these effects? For the average state, roughly 3% of college-educated private sector workers have moved across state lines in the last year. Thus, the 13% decrease in annual flows that we estimate corresponds to a roughly .4 percentage point decline per year, which over four years accumulates to a 1.6% decline in the stock of college-educated workers. This is roughly equal to the annual growth that the average state sees in its college-educated workforce (see Appendix Figure B4), implying that four years under a Republican governor sets back the state's human capital accumulation by about one year.

3.3 Alternative identification strategies

It is not unreasonable for a reader to worry about reverse causality. We have shown that college-educated workers would generally prefer liberal candidates. Thus, any decrease in the stock of college-educated workers should, all else equal, decrease Democratic candidates' vote shares and increase the probability of a Democrat-to-Republican gubernatorial transition. In principle, one would expect this to be reflected in the pre-election migration trends – which we do not see – but it is nonetheless a legitimate concern.

[21] Nonetheless, it is important to acknowledge that the labor market was healthier in the mid-aughts and post-2015 period than at any point from 2007-2014. Thus, a potential alternative explanation for our *heterogeneous* effects is that migration is only responsive to politics when the labor market is strong enough to afford workers decent employment options. We stress that this is a possible alternative explanation for the treatment effect *heterogeneity* that we document, but cannot explain the main treatment effect that we estimate.

In Appendix Section B.2.3, we propose two alternative identification strategies. The first is based on an instrumental variables strategy based on pre-determined variation in the timing at which different states hold their gubernatorial elections. Nearly all states give governors four-year terms, but some states hold these elections during presidential years (e.g., 2008, 2012) while others hold them during off-presidential years (e.g., 2010, 2014). This means that some governors were elected during 2008 (a record year for Democrats, according to Congressional election returns, in part because of the Obama campaign's success), while others were elected during 2010 (a record year for Republicans who came to power on the backs of the right-wing Tea Party wave). Our instrument shows that these national swings in partisan sentiment translate into large effects on gubernatorial election outcomes for governors who are lucky/unlucky enough to be running during these years. Our other identification strategy is a standard regression discontinuity design (RDD) based on close gubernatorial elections.

In both cases, we estimate substantially similar effects of a Democrat-to-Republican transition on in-migration of college-educated workers. These effects are similar not only in their levels, but also in their trends over time. However, none of these estimates are statistically significant because they only use a relatively small share of the data. Thus, the standard errors are much larger than our preferred difference-in-difference estimates. We interpret this as evidence that our difference-in-difference approach does not suffer from biases that could potentially arise from the fact that election outcomes are endogenous with respect to demographic changes.

To us, this is conclusion is unsurprising. While it is clearly true that a change in the composition of the electoral will affect vote shares and election outcomes, electoral politics is an inherently noisy and chaotic process, and changes in the composition of the electorate are unlikely to be a primary determinant of outcomes.[22] While a classic perspective in political science held that perceptions of economic conditions were key to incumbents' reelection (which would be concerning, given they they are also related to in-migration), this relationship is quite weak in recent years, as Democrats and Republicans now hold very different perceptions of economic conditions (Jones, 2020).

Considering some recent transitions within our sample illustrates the importance of "random" shocks in modern electoral politics. In 2016, Missouri elected a Republican to replace term-limited Democrat Jay Nixon. Many argued at the time that Democrats performed poorly in the election because of the 2014 murder of Michael Brown at the hands of police in Ferguson, and mistakes made by party leaders (including Nixon and the Democrat's 2016 gubernatorial candidate, who had been the Attorney General at the time) in responding to the protests

[22] A large literature in political science shows that voters' choices are significantly influenced by truly random events like college football matches (Healy, Malhotra, and Mo, 2010) and shark attacks (Achen and Bartels, 2017), though this literature is controversial (Fowler and Hall, 2018; Fowler and Montagnes, 2015).

(Ortiz, 2016). In 2012, North Carolina elected a Republican governor for the first time since 1988. The Democratic candidate was the incumbent Lieutenant Governor, and commentary at the time focused on the challenges he faced campaigning at the same time that several of the Governor's staff members were facing obstruction of justice charges (Catanese, 2012) and the Democratic Party's senior staff was embroiled in sexual harassment proceedings (Bass, 2012). In 2015, Louisiana elected Democrat John Bel Edwards to replace term-limited Republican Bobby Jindahl. Many argued that Edward's defeat of Republican David Vitter was because of Vitter's earlier prostitution scandal (McAfee, 2015), which led key Republicans to either endorse Edwards or no one (Barnes, 2015).

Importantly, many of these races were not close, though they appear to be driven by events exogenous to migration incentives. Put differently, the set of election outcomes useful for identification is larger than only the close elections used by an RDD strategy. For this reason, in our model, we do not attempt to endogenize political outcomes, although we acknowledge that our model is fundamentally about endogenous shifts in the composition of the electorate.

Finally, it is worth noting that we are not the first to conclude that difference-in-difference estimates of gubernatorial transitions do well at replicating RDD estimates. In their study of the effects of governor partisanship on policy outcomes, Caughey, Xu, and Warshaw (2017) also find that difference-in-difference estimates are similar to the RDD estimates, both in levels and in over-time heterogeneity.

3.4 Alternative explanations

We have now established that the decline of college in-migration following a D-to-R gubernatorial transition is the causal effect of the election outcome. However, this might not necessarily imply that it reflects these workers' political preference (despite the fact that we have shown a strong correlation over time between preferences and responses to elections).[23] In Appendix B.2.4, we consider two alternative reasons why college in-migration might respond to election outcomes: *i*) effects on economic opportunities and *ii*) effects on perceptions of citizens' political views.

First, rather than reducing labor supply of college graduates who disprefer conservative governance, Republican governors could reduce demand for college graduates via spending cuts that induce an economic contraction and reduce job opportunities.[24] In Appendix Figure B17,

[23] As an additional test of our political preferences explanation, in Appendix Section B.2.2 we compare different types of majors depending on the average political lean of college seniors. We find stronger migration responses among graduates holder further left-leaning majors.

[24] We see this explanation as, *ex ante*, unlikely. First, economic contractions tend to be disproportionately experienced by non-college workers rather than college graduates (Patterson, 2020). Second, Republican governors tend to lead to more pro-growth business-friendly policies like deregulation and tax cuts.

we test for effects on job openings and hiring. We find no short-run effects on job openings, while hiring falls immediately. This suggests a labor supply effect rather than a labor demand one.

Second, Republican governors might reduce in-migration of college graduates not because those graduates are averse to conservative *policies*, but because the election signals the prevalence of conservative *voters*, and college graduates do not want to live near these conservative voters. This matters for policy because it suggests that the actual decisions of Republican governors about what to do once in office have no bearing.[25] To test for this, we test whether the six states that switched from voting for Barack Obama in 2012 to Donald Trump in 2016 saw decreased college in-migration relative to the 20 Obama-Clinton states. Because all of these states had the same actual president, this switch does not affect policy, but likely does affect voters' perceptions of the states' residents. We find no migration effects of these placebo switches. We conclude that our estimated effects reflect labor supply responses to governors, and we model them accordingly.

4 Model setup

In this section, we build a static general equilibrium model of migration based on Bryan and Morten (2019).[26] College and non-college workers have different preferences over political amenities, in addition to the other amenities more standard in the migration literature.

Workers initially live in origin o and sort across destinations (denoted d) based on wages, amenities, migration costs, and partisanship of governors, as well as an idiosyncratic skill draw for each potential destination.

4.1 Utility

Let the utility of worker i belonging to education group g and moving from origin o to destination d be given by:

$$U_{ido}^{g} = c_{id}^{g}\alpha_{d}^{g}(1-\tau_{do}^{g})(1-\gamma_{p(d)}^{g})$$

where c_{id}^{g} is her consumption and α_{d}^{g} is the amenity value of living in destination d for education group g. We consider amenities to be fixed across individuals with the same education, but we

[25] Although stated choice experiments over neighborhoods do show that voters prefer living near ideologically similar neighbors (Mummolo and Nall, 2017), these effects are fairly small, and it is important to note that all cities (even in conservative states) have liberal neighborhoods that residents can choose.

[26] There are other models of spatial sorting (e.g., Diamond (2016)). An advantage of the Bryan-Morton model is that identification and estimation are based on observed migration flows, whereas some other models infer migration behavior from relative changes in city size over long periods. Given that our reduced form estimates focus specifically on migration flows, this model provides a more natural framework to incorporate those results.

allow all parameters to vary across education groups, denoted by $g = C$ for college graduates and $g = N$ for non-college. Let τ_{do}^{g} be the migration cost of moving from o to d. We assume there are no costs of not migrating ($\tau_{oo} = 0$) and that migration costs are symmetric ($\tau_{do} = \tau_{od}$).

The term $\gamma_{p(d)}^{g}$ captures a group-specific preference wedge emerging from group g's disutility of living under a governor of partisanship $p(d) \in \{\text{Republican}, \text{Democrat}\}$. When this wedge is positive, workers must be compensated with $1/(1-\gamma_{p(d)}^{g})$ times greater consumption in order to be indifferent. Since college-educated workers are to the left of non-college ones, we normalize $\gamma_{Dem}^{C} = 0$ and $\gamma_{Rep}^{N} = 0$ so that γ_{Rep}^{C} captures college graduates' disutility of Republican governance, relative to Democratic governance, and conversely for non-college workers' disutility of Democratic governance.[27]

Following Bryan and Morten (2019) and Hsieh et al. (2019), each worker gets one idiosyncratic skill draw for each potential destination state. We think of this draw as reflecting the best employment opportunity available in that state, and we sometimes refer to it as a match-specific human capital draw. We assume these idiosyncratic opportunities, denoted s_{id}^{g}, are drawn from a multivariate Fréchet distribution:

$$F(s_1^g, s_2^g, ..., s_N^g) = exp\left(-\left[\sum_{d=1}^{N} s_d^{-\theta^g}\right]\right)$$

The distribution of s_{id}^{g} is governed by the scale parameter θ^g. A higher value of θ^g implies less skill dispersion across locations, such as would be the case if all states afforded the worker an equally good employment match. As θ^g decreases, there is a greater difference between the match-specific human capital realizations a worker sees in different locations. This parameter is key because it determines how close to indifferent workers are between the employment opportunities available in different states. Note that this formulation assumes each worker's draws are independent across states, although state-specific factor supplies and wage rates will lead to a correlation in *earnings* across different workers within the same state.

We derive the indirect utility function for worker i in group g who moved from o to d as:

$$V_{ido}^{g} = \alpha_d^g (1 - \tau_{do}^{g})(1 - \gamma_{p(d)}^{g}) w_d^g s_{id}^{g} \tag{2}$$

[27] This parameter varies across education groups but not across individuals within the group. It is obviously stark and unrealistic to assume that all individuals of the same education have the same political preferences. More realistically, we could assume that only some share $\chi < 1$ of college graduates are liberal. But with this formulation, the decline in college in-migration that we estimated above would have been generated entirely by this χ share of liberal college graduates. This would require a larger value of the γ disutility parameter, since this subset's behavioral responses would have to be larger, but would not change the aggregate size of the migration response (since this is a targeted moment). This would change the welfare implications of our model (since it would increase the intensity and heterogeneity in preferences), but likely would not affect other implications, which is where our analysis focuses.

where w_d^g is the human capital price in d for group g. That is, if $w_d^g > w_{d'}^g$ then d has higher "wages" (human capital prices) than d': a worker with the same quality of employment opportunity in both d and d' (i.e., the same human capital draw: $s_{id}^g = s_{id'}^g$) will have higher earnings and consumption in d since it is a higher wage state.

It is convenient to define the overall utility returns of destination d for a worker of group g and origin o as $\tilde{w_{do}^g} = \alpha_d^g(1-\tau_{do}^g)(1-\gamma_{p(d)}^g)w_d^g$ so that $V_{ido}^g = \tilde{w_{do}^g} s_{id}^g$.

Our key assumption is that workers choose the location that yields the highest indirect utility. Note that the only idiosyncratic component that varies across individuals is their vector of skill draws. Since these are Fréchet distributed, a property of the Fréchet distribution is that the share of individuals from origin o who choose to work in destination d can be written as:

$$\pi_{do}^g \equiv \frac{L_{do}^g}{L_o^g} = \frac{\tilde{w_{do}^g}^{\theta^g}}{\sum_{j=1}^N \tilde{w_{jo}^g}^{\theta^g}} \tag{3}$$

where L_o^g denotes the number of group-g workers from origin o, and L_{do}^g denotes the number who move to d. These migration flows are driven by relative returns, including non-wage returns in terms of higher amenities or lower migration costs.

A property of the Fréchet distribution is that, conditional on the workers' optimal choice of where to live, the average skill of workers choosing to move from o to d can be written as:

$$\mathbb{E}\left(s_d^g \mid \text{choose } d \text{ from } o\right) = \bar{\Gamma^g} \pi_{do}^{-\frac{1}{\theta^g}} \tag{4}$$

where $\bar{\Gamma^g} = \Gamma(\frac{\theta^g-1}{\theta^g})$ and $\Gamma(.)$ is the gamma function. This property is important, as it allows us to infer unobserved human capital from observed migration rates.

4.2 Production

Firms produce a single final good, which is costlessly traded and is chosen as the numeraire ($p = 1$). Output is produced by perfectly competitive firms. They combine the effective labor (i.e., total human capital) of the two education groups using a Constant Elasticity of Substitution (CES) production function.

Total output in state d is given by:

$$Y_d = A_d \left[\left(H_d^C\right)^{\frac{\sigma-1}{\sigma}} + \left(H_d^N\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} \tag{5}$$

where A_d is the exogenous location-specific total factor productivity (TFP) and $\sigma \geq 1$ is the elasticity of substitution between college and non-college workers. H_{jd}^g is the total efficiency units of labor employed by firm j, which is simply the sum of skill draws for workers working at firm j:

$$H^g_{jd} = \sum_{i \in P_j} s^g_{id}$$

where P_j denotes the set of workers employed at firm j.

The profits for a representative firm in location d are given by:

$$\Pi_{jd} = Y_{jd} - w^C_{jd} H^C_{jd} - w^N_{jd} H^N_{jd}$$

where w_{jd} is the wage paid per effective unit of labor by firm j. In this economy, firms compete for each type of worker by setting wages. Since the labor market is perfectly competitive, in equilibrium, $\Pi_{jd} = 0$ and $w^g_{jd} = w^g_d \ \forall j$. This prevailing wage (or the price of human capital) is the same w^g_d discussed in the worker's problem above.

4.3 General equilibrium

A competitive equilibrium in this economy consists of destination choices, total efficiency of labor in each destination H^g_d, and a wage w^g_d such that:

- Workers choose the workplace that maximizes their utility
- Firms choose efficient labor H^g_d to maximize profit
- w^g_d clears labor market for each destination

5 Identification and estimation

In this section, we describe the procedure for estimating the parameters $\left\{\theta^g, w^g_d, \tau^g_{do}, \alpha^g_d, \gamma^C_{\text{Rep}}, \gamma^N_{\text{Dem}}, A_d\right\}$. Unless otherwise specified, all calculations are based on the sample of ACS respondents from 2011-2019, age 26 or older, and employed in the private sector. In all cases, o refers to state of residence last year, and d refers to current state of residence. In all calculations based on worker earnings, we restrict to full-time, full-year employees since ACS earnings are reported annually.

5.1 Identification

5.1.1 Elasticity of skills substitution: σ

A large literature estimates the elasticity of substitution between skills. We use $\sigma = 2.6$, which is the estimate from Jerzmanowski and Tamura (2020) for college and non-college workers.

It is worth noting that we use a value such that $\sigma > 1$. This implies that a decrease in the supply of college-educated workers will increase inequality through two forces. First, it will raise the relative wages of college graduates (i.e., college and non-college workers are gross substitutes). Second, it will reduce the real wage of non-college workers (i.e., they are q-complements).[28]

5.1.2 Fréchet parameter: θ^g

Workers' earnings are determined by state-specific human capital prices (the wage rate paid per effective unit of labor), and the individual human capital draw with which workers move to the state. Looking at equation (4), once one accounts for the cross-state variation in the returns to human capital and variation across *od* in the degree of worker self-selection, the remaining variation in earnings is informative about the underlying dispersion of human capital draws. This insight was pointed out by Hsieh et al. (2019), and our approach to estimation follows theirs. Specifically, we regress log annual earnings in the ACS on destination-by-origin fixed effects, while also controlling for age, gender, and fixed effects for year, occupation, and industry. We do this separately for college and non-college workers, and calculate residual log earnings.

Since we control for destination-by-origin fixed effects, these residualized earnings have been purged of differential cross-state selection and heterogeneous returns to human capital. Since we control for individual-level covariates, the remaining dispersion reflects the variation in place-adjusted earnings that one worker could plausibly see across different potential states of residence. In this way, we can write a non-linear function of θ as a non-linear function of the distribution of residual wages:

$$\frac{\text{Variance}(\tilde{s}_i^g)}{(\text{ Mean}(\tilde{s}_i^g))^2} = \frac{\Gamma\left(1 - \frac{2}{\theta^g}\right)}{\left(\Gamma\left(1 - \frac{1}{\theta^g}\right)\right)^2} - 1 \tag{6}$$

where $\tilde{s}_i^g$ is individual i's residual log earnings, exponentiated.[29]

5.1.3 Moving cost: τ_{do}^g

As Bryan and Morten (2019) point out, moving costs are identified because they depress migration flows between two states *symmetrically*, while differences in wages or amenities would

[28] These results mirror historical discussions in labor economics about the wage inequality implications of changes in the aggregate supply of college graduates. See Acemoglu and Autor (2011) for some discussion, who calibrate a value of $\sigma = 2.9$ for the 1964-2008 period.

[29] Residual log earnings are mechanically mean zero across *od*, but exponentiated residual log earnings are not mean one because the exponential function is non-linear.

only depress flows in one direction. We take the log of equation (3):

$$\ln(\pi^g_{do}) = \theta^g \ln(\alpha^g_d) + \theta^g \ln(w^g_d) + \theta^g \ln(1-\gamma^g_{p(d)}) + \theta^g \ln(1-\tau^g_{do}) - \ln\left(\sum_{j=1}^{N} \tilde{w^g_{jo}}^{\theta^g}\right) \tag{7}$$

In this expression, most parameters are destination-specific, and so they cancel out when we compare migration from two different origins into the same destination. This includes when considering "migration" from origin d to destination d (i.e., the probability of staying in d), which is useful because d-to-d migration costs do not appear because of our assumption that there are no costs to *not* migrating (i.e., $\tau_{dd} = 0$). Thus, the difference between rates of o-to-d migration and rates of staying in d can be written as the migration costs and the difference in the logged sum of o-specific and d-specific returns to migration:

$$\ln \pi^g_{do} - \ln \pi^g_{dd} = \theta^g \ln(1-\tau^g_{do}) - \ln\left(\sum_{j=1}^{N} \tilde{w^g_{jo}}^{\theta^g}\right) + \ln\left(\sum_{j=1}^{N} \tilde{w^g_{jd}}^{\theta^g}\right)$$

The same derivation can be done to calculate $\ln \pi^g_{od} - \ln \pi^g_{oo}$ (migration from d to o relative to the probability of staying in o), and then summing these terms, the differences in the sum of returns cancel out, leaving only:

$$[\ln \pi^g_{do} - \ln \pi^g_{dd}] + [\ln \pi^g_{od} - \ln \pi^g_{oo}] = \theta^g \ln(1-\tau^g_{do}) + \theta^g \ln(1-\tau^g_{od}) \tag{8}$$

Because we assume that migration costs are symmetric ($\tau_{od} = \tau_{do}$), this reduces to one equation and one unknown (conditional on the estimate of θ^g from above) per origin-destination pair. In this way, the matrix of observed migration flows and non-migration decisions identifies all migration costs. Because it is identified only by the *symmetry* of flows, differences in attractiveness of the two states is irrelevant, and so identification does not depend on who is governor.

5.1.4 Productivity and wages: A_d, w^g_d

We jointly identify productivity and wages using two conditions. First, the assumption of perfect competition in the product market and the law of one price implies that price equals marginal cost, which yields:

$$\left(\frac{w^C_d}{A_d}\right)^{1-\sigma} + \left(\frac{w^N_d}{A_d}\right)^{1-\sigma} = 1 \tag{9}$$

Second, labor market clearing implies that total human capital demanded equals total human capital supplied. Human capital demanded can be written from the CES production function as a function of wages and total output. Human capital supplied is simply the product of the equilibrium skill conditional on migrating from o to d (derived from the Frechet above in equation (4)) and the number of workers migrating from o to d ($L_o^g \pi_{do}^g$), summed over all origins.

$$H_d^g(demand) = \left(\frac{A_d^{\frac{\sigma-1}{\sigma}}}{w_d^g} \right)^{\sigma} Y_d = \sum_{o=1}^{N} L_o^g \pi_{do}^g \bar{\Gamma}^g \pi_{do}^{-\frac{1}{\theta^g}} = H_d^g(supply) \tag{10}$$

Given the above estimate of θ, and observable output and migration, this delivers three equations (since equation (10) must hold for $g \in \{C, N\}$) and three unknowns (A_d, w_d^C, w_d^N) for each state.

Fundamentally, TFP and wages are pinned down by variation in GDP. Given an estimate of skill dispersion, our model tells us how human capital stocks can be determined by population size and migration (since migration entails selection: higher rates of in-migration crowd-in marginally lower human capital draws). Thus, surprisingly high levels of GDP (given the human capital stock) imply a high total factor productivity, and because all firms are assumed to compete in the same product market and face the same production function, there is a stable relationship between productivity and wages across states. Thus, variation in productivity combined with the available supply of human capital determines the wages.

5.1.5 General amenities: α_d^g

When governors in d and d' belong to the same party, politics will not affect workers' decisions about whether to migrate from o to d as opposed to o to d'. Considering equation (3) and taking the difference in migration between two same-governor destination states, we have

$$[\ln(\pi_{do}^g) - \ln(\pi_{d'o}^g)] = \theta^g \ln(\alpha_d^g/\alpha_{d'}^g) + \theta^g \ln(w_d^g/w_{d'}^g) + \theta^g \ln\left((1-\tau_{do}^g)/(1-\tau_{d'o}^g)\right) \tag{11}$$

With θ^g, w_d^g, τ_{do}^g at hand, we can identify relative amenities, relative to some reference state d' (discussed below). Essentially, state d's amenities are identified by higher migration into d and lower out-migration out of d than one would expect given wages and estimated migration costs.

5.1.6 Political amenities: $\gamma_{\textbf{Rep}}^C$, $\gamma_{\textbf{Dem}}^N$

To calculate our values of political amenities, we perform a grid search over different values for the pair γ_{Rep}^C and γ_{Dem}^N such that our model-implied effects of switching governors match our reduced form estimates. More specifically, our primary reduced form estimates are identified from eight states switching from a Democratic to a Republican governor between 2015 and

2017. For each value pair of values for γ^C_{Rep} and γ^N_{Dem}, we simulate two counterfactuals from our model. First, we simulate outcomes where all governors are held fixed as they were in 2014. Second, we simulate outcomes where all governors are held fixed as they were in 2017.

Let S denote the set of states that switched from a Democratic to a Republican governor between 2015 and 2017. The purpose of difference-in-difference is to identify the expected difference in potential outcomes between treated states under treatment and treated states in the absence of treatment: $E\big[Y_d(Rep) - Y_d(Dem)|d \in S\big]$, which can be rewritten as $E\big[Y_d(Rep)|d \in S\big] - E\big[Y_d(Dem)|d \in S\big]$. Given some choice of the parameters γ^C_{Rep} and γ^N_{Dem}, our two simulations allow us to recover these two quantities by averaging (over the eight treated states that identify our reduced form estimates) the change in log in-migration rates between the counterfactual where governors are as they were in 2014 – an estimate of $E\big[Y_d(Dem)|d \in S\big]$ – and where governors were as they were in 2017 – an estimate of $E\big[Y_d(Rep)|d \in S\big]$.

For each education group g and each choice of γ^C_{Rep} and γ^N_{Dem}, these two counterfactuals are used to generate one moment:

$$\varphi^g = \frac{1}{|S|} \sum_{d \in S} \left[\underbrace{\ln \left(\sum_{o=1, o \neq d}^{N} L^g_o \pi^g_{do}(p(d)') \right)}_{Y_d(Rep)} - \underbrace{\ln \left(\sum_{o=1, o \neq d}^{N} L^g_o \pi^g_{do}(p(d)) \right)}_{Y_d(Dem)} \right] \tag{12}$$

According to our difference-in-difference analysis, when a Democratic governor is replaced by a Republican governor, the migration flow of college-educated workers decreases by 12.8%, while the migration flow of non-college-educated workers does not show a significant change. Therefore, we choose γ^C_{Rep} and γ^N_{Dem} to match the target of $\hat{\varphi^C}$ = -12.8% and $\hat{\varphi^N} = 0$.

5.2 Estimation

5.2.1 Migration flow

The matrix of migration flows is central to our estimation procedure. Using the American Community Survey (ACS), we calculate the matrix of cross-state flows, year-by-year, separately for college and non-college workers. However, this matrix, calculated from raw data, has two issues for which we make adjustments.

First, since the ACS is only a 1% sample, flows are measured with error. This causes a problem when flows are low, in which case there is a chance that the 1% sample misses the people who moved from o to d, and we observe zero flows. This is an issue because many key moments derive from the *log* of migration flows.

As is well-known, ignoring these missing flows can bias our estimates, and so we build on the approach suggested by Silva and Tenreyro (2006) and assume that the realized number of college graduates observed in the ACS migrating from o to d in year t is drawn from a Poisson process. We can then estimate the expected number of true flows by imposing a functional form assumption on how these flows relate to state characteristics.

Key to this approach is that we have administrative data made publicly available by the Internal Revenue Service (IRS) on the total number of tax filers who moved from origin o to destination d in year t. This data is not suitable for our main reduced form regressions because it includes no demographic information (such as education), but it can be used to infer flows for our structural estimates. To do so for each origin-destination pair in each year, we assume that the *ratio* of college graduate movers to IRS movers is a function of an origin fixed effect, a destination fixed effect, a time fixed effect, the log distance between o and d, and the ratios of per capita GDP, group-specific wages, and college populations between the two states. This is a very flexible approach that assumes that we can describe very well the *relationship* between college migration and total tax-filer migration, and then imposes no assumptions or structure on the realizations of tax-filer migration across pairs or over time. Under these assumptions, we can write the expected number of college migrants per tax-filing migrant as:

$$\frac{E\left[L_{dot}^{g}\right]}{IRS_{dot}} = exp\left(\theta_o + \eta_d + t + \beta_1 \log(distance_{do}) + \beta_2 \frac{GDPpc_d}{GDPpc_o} + \beta_3 \frac{Wage_{dt}^{g}}{Wage_{ot}^{g}} + \beta_4 \frac{L_{dt}^{g}}{L_{ot}^{g}} + \epsilon_{dot}\right) \quad (13)$$

where IRS_{dot} refers to observed migration in the IRS administrative data. Equation (13) can be estimated using a Poisson pseudo-maximum likelihood procedure (Correia, Guimarães, and Zylkin, 2019, 2020), and we can calculate the expected number of college graduate migrants $\tilde{L}_{dot}^{g}$ from the observed IRS migrants (where we do not know education). We do this for all $o \neq d$, and then calculate π_{oot} as $1 - \sum_{d \neq o} \pi_{dot}$.

Second, these flows are one-year flows. They represent the number of people moving from o to d from one year to the next. However, our goal is to use our model to simulate the steady state that prevails four years after a gubernatorial switch. Thus, we need to adjust our one-year migration rates to be four year migration rates.

If we assume that cross-state migration is a Markov process, then we can calculate the four-year migration rate from o to d as the probability that an individual moves from o to d in year $t + k$ and then remains in d for the remaining $4 - k$ years. This can be written as $\pi_{do}\pi_{dd}^{3} + \pi_{do}\pi_{dd}^{2} + \pi_{do}\pi_{dd} + \pi_{do}$, where π_{dd} is the probability of staying in destination d. We average our year-by-year estimates over the nine years we can calculate them (π_{dot} estimated using equation (13) above), and then calculate these four-year flows from the resulting averages of one-year flows.

5.2.2 Estimation

We provide a detailed explanation of estimating all the parameters in five steps. In all of these steps, we use the estimated migration flows based on the two adjustments to the raw ACS flows which we described in the previous section.

The first step involves estimating θ^g using equation (6) based on the dispersion of residual wages. We use workers' real earnings, which refers to the nominal wage obtained from ACS divided by the price deflators estimated by Bureau Economic Analysis at year and state levels. We calculate the wage residuals from a cross-sectional regression of log annual wage on age, age squared, gender, marital status, education, race and fixed effects for year, occupation, and industry.

Next, we estimate the matrix of migration costs, τ^g_{do}, based on equation (8). This relies on the estimate of θ^g from above and the matrix of migration flows.

After estimating θ^g and τ^g_{do}, the following step is to recover A_d and w^g_d which are determined using the system of equations (9) and (10). For output Y_d, we first adjust nominal state-level GDP using the year-level and state-level price deflators that we used to adjust workers' nominal earnings. We then measure Y_d as the average of real GDP from 2011 to 2019.

To determine α_d, it is necessary to select a reference state that experienced a change in the governor's party between 2011 and 2019, as α_d can only be estimated if the governors of d and the reference state belong to the same party. We compile a list of states that have experienced such a switch and select as a reference state the state with the highest migration flows for both college-educated and non-college workers: North Carolina. Given this, we can recover α_d using equation (11). For any state d, it is often the case that there exist multiple years in which the same governor serves with the reference state. In such situations, we compute the average of $\hat{\alpha_{dt}}$ values to recover α_d.

The final step is to recover the last two parameters, γ^C_{Rep} and γ^N_{Dem}, which are chosen to match the reduced form results. We use a grid search method that allows each γ to take on values ranging from -0.5 to 0.5, with a grid size of 0.005. We estimate γ^C_{Rep} and γ^N_{Dem} that produce the minimum sum of absolute errors when we solve for equation (12).

5.2.3 Resulting estimates

We present the estimation results in Table 2.[30] Our estimate of γ^C_{Rep} (college workers' preference against living under a Republican governor) is 0.065, which can be compared with

[30] For brevity, we do not present extensive validations of our estimated parameters. Our estimates of γ^C and γ^N are global optima, our estimated migration costs are positively correlated with the distance between states, and our estimated amenities are positively correlated with the estimates of others' (Albouy, 2016; Diamond, 2016; Kleinman, Liu, and Redding, 2023), as well as with raw data on net migration residualized of wages.

both non-political amenities and migration costs, since all enter the indirect utility function in equation (2) together. Our estimate is about 1/3 of the cross-state standard deviation in non-political amenities. For the average state, comparing the alternative destinations with the lowest and highest migration costs, the difference is approximately .11 (almost twice $\hat{\gamma}^{C}_{\text{Rep}}$). Of course, the average costs incurred from moving at all, relative to staying, are about 14 times larger than the political wedge. To compare to wages, note that a college graduate would need about 7% higher consumption to be indifferent about a Republican governor. This is only one-third of the increase in consumption that would come from a one standard deviation increase in a state's TFP or human capital price.[31] In other words, our estimate of $\hat{\gamma}^{C}_{\text{Rep}}$ is not large. The partisanship of the governor is a substantially less meaningful than other non-wage amenities, the costs of migrating to one state vs. another, or the wages available in different states. However, as we will show, even this modest disutility of Republican governance can generate meaningful effects on the college population and state-level output.

[Table 2 about here.]

It is worth noting that we estimate a positive but small disutility for non-college workers living under a Democratic governor. Our estimate is only one-twelfth as large as the college-educated workers' political wedge, but it is not zero. Above, we estimate no statistically or economically significant effects of governor's party on migration of these workers. As we show below, however, college-educated workers sorting away from Republican-governed states reduces output and wages there, which increases non-college workers' incentives to avoid those states, too. Thus, our model can only reconcile zero migration responses for non-college workers by concluding that non-college workers do have some preference for living under Republican governors, consistent with our descriptive facts about political opinions above.

6 Model-implied effects of governors

In this section, we use counterfactuals generated by our model to understand the effects of gubernatorial transitions on equilibrium output, wages, and inequality, as well as how these effects spill over onto other states. In all cases, we compare steady state equilibria under one set of governors with that under another set. Since our estimation matches our reduced form estimates that take place over the four post-election years, we view these as relatively short-run (four-year) effects. Thus, we do not attempt to model potential changes in TFP that might result from changes in industry structure or firm location, changes in amenities that might

[31] This calculate holds fixed the human capital draw, of course, and compares to the average. Thus, $99.6/482 = .21$ and $58.3/269 = .22$

be endogenously influenced by the educational composition of the population (as in Diamond (2016)), or changes in future politics that might result from changes in the composition of the electorate. We view all of these as plausible long-run outcomes which could be influenced by the governor's party and its effects on migration, but where further empirical work would be needed to discipline a quantitative model.

6.1 Heterogeneity

An interesting result from our analyses is that our model implies significant heterogeneity in the predicted effects of a gubernatorial transition. This is despite the fact that we impose substantial homogeneity by using only two disutility parameters: one that applies to *all* college graduates and *all* Republican governors, and one that applies to all non-college workers and all Democratic governors. Figure 7 presents the results from 50 separate counterfactuals, each of which "flips" the governor of a single state, leaving all other governors unchanged. We map the percent change in the number of college graduates that our model predicts the treated state would experience.

Blue states have Democratic governors in 2023, and we are therefore simulating the effects of a switch to a Republican governor (which reduces the number of college graduates). Red states currently have a Republican governor, and so we're simulating the effects of a Democratic governor. In both cases, darker colors reflect larger predicted proportional effects.

[Figure 7 about here.]

Implied decreases in the number of college graduates range from 2% to 5.3%, and implied increases range from 2.4% to 8.5%.[32] In Appendix Table B1, we show that this heterogeneity is mostly driven by differences in the baseline fraction of college graduates that come from other states. States that rely on in-migration to generate their college-educated workforce are obviously more susceptible to the impacts of changing migration incentives. This generates an interesting pattern in the map in Figure 7, in which nearly all of the states predicted to have the smallest effects are those with very strong public university systems,[33] though our estimation

[32]It is worth noting that these effects are somewhat larger than the back-of-the-envelope calculation we presented in the reduced form section. This is mainly because those calculations were based only on the effects on *inflows*. Our model also predicts effects on *outflows* in response to a change in governor, and the total effect in our model comes from both the combination of the in-migration and out-migration responses. We see little empirical evidence for out-migration responses (see Figure B8), though the estimate is noisy. A generally shortcoming of static models is that they cannot match the asymmetric in-/out-migration effects seen in response to many shocks (Monras, 2020), but we see it as beyond the scope of this paper to develop a more complicated dynamic model.

[33]Democrat-governed states: California, Minnesota, Wisconsin, Michigan, and Pennsylvania. Republican-governed states: Texas, Indiana, and Ohio.

approach did not use this information. This highlights an important policy implication of our results. Republican governors interested in blunting these effects might invest in universities to produce "home-grown" college-educated workers, although in practice, Republican governors and legislatures tend to reduce university funding.

In Appendix Figure B11, we show that our model-implied heterogeneity is realistic. Using our same reduced form specification above, we estimates separate effects for each of the eight treated states driving identification in our primary specification (Democrat-to-Republican transition in 2015-2017). We then calculate these same predicted outcomes from our model. The state-specific empirical estimates are positively correlated with the model-implied predictions. In general, that states predicted by our model to show smaller effects also realized smaller effects empirically. While the empirical and model-implied effects do not line up perfectly, we see this as valuable validation that our model-implied heterogeneity is realistic. This is particularly noteworthy since our model abstracts from many types of realistic heterogeneity across different voters or different governors.[34]

6.2 Counterfactuals

Because different states show different implied effects, it is important that our counterfactuals reflect realistic scenarios. For example, Wyoming shows among the largest effects of a Democratic governor, but the last Democrat running for governor there received only 17% of the vote. Thus, these large effects are irrelevant.

In choosing our counterfactuals, we recognize that voters show increasingly strong partisan identities, are increasingly reluctant to vote for different parties for different offices, and increasingly focus on national (not state) political debates and then evaluate state candidates according to those same terms. Grumbach (2022) calls this the "nationalization" of state politics. Recognizing this, we consider counterfactuals in which the recent 2020 Presidential election results imply that a different state-level gubernatorial outcome is plausible.

Figure 8 plots the Democratic Party vote share from the most recent Presidential election, along with the most recent gubernatorial election. The two are strongly correlated (.73), and dropping Vermont makes the correlation even stronger (.86). However, the arbitrariness inherent in winner-take-all elections means that outcomes differ, and since many elections were close, modest shifts in public sentiment could easily sway the outcomes in many of these states (similar to the IV strategy we developed above).

[Figure 8 about here.]

[34] It is notable that our largest empirical *underestimate* of the model-implied effects is in Massachusetts, where Republican Governor Charlie Baker was rated as the single most popular governor in the country (Leins, 2019).

In total, we consider four counterfactuals, all of which are shown in Figure 8. First, we consider a "weak red wave," in which a modest tide of conservative national sentiment could flip the four states won by Donald Trump but which currently have Democratic governors. Second, we consider a "strong red wave" which also flips an additional four states that Biden won narrowly and which have Democratic governors. Similarly, we consider a "weak blue wave" that flips the five states won by Biden but represented by a Republican governor, and a "strong blue wave" that also flips four battleground states narrowly lost by Biden and represented by a Republican governor.

6.3 General equilibrium effects

Our first question is how these plausible "waves" of Republican and Democratic gubernatorial elections might affect equilibrium economic activity and inequality. Panels A and B of Table 3 summarize our model's prediction for how the college-educated workforce and GDP per worker would change in the states experiencing the gubernatorial flip (we return to spillovers on other states below).

Across these counterfactuals, we estimate that the average state switching its governor would see a 3% change in the stock of college graduates. But in each counterfactual, we estimate significant heterogeneity, with the most affected states seeing roughly twice the response of the least affected states. These 3% shifts in college workforce translate into roughly 1% shifts in GDP per worker.[35] These are not large effects. A 1% decline in a state's equilibrium GDP per capita would not cripple its economy. Nonetheless, it would be meaningful. The median state saw average annual GDP growth of 1.25% during the 2005-2019 period, so our estimated effects roughly correspond to one year of lost growth.

[Table 3 about here.]

We next turn to a consideration of inequality. First, within our model, increasing the number of college graduates should unambiguously reduce the relative "wage" (i.e., earnings per unit of human capital) of college graduates. This occurs both because the college graduates are becoming less scarce (reducing college wages), but also because the CES production function implies that an increase in college graduates raise non-college workers' productivity and wages. In Panel C of Table 3, we present the joint effects on inequality in human capital price (i.e., how equilibrium $\ln w_d^C - \ln w_d^N$ changes when the governor changes). We find that the log wage gap falls by an average of 11-18% in response to a blue wave, while is rises by 11-13% in response to a red wave. In Appendix Table B2 we decompose this into changes in non-college wages and

[35] Our model does not include non-workers, so GDP per worker is equivalent to GDP per capita.

changes in college graduate wages. Across all four counterfactuals, we find that about 45% is explained by changes in the real wages of non-college workers, with the remainder explained by changes in college graduate wages.

These estimates represent a model-consistent notion of wage inequality (the returns to a given skill draw from the Fréchet distribution), but are unobservable to a researcher who could not measure workers' idiosyncratic productivity draws. Changes in *earnings* inequality are observable, but will differ from the theoretically unambiguous effects on wage inequality because of two separate effects of changing the composition of workers in the state. Both of these effects emerge because migrants in our model are positively selected; in order to move to d, their idiosyncratic human capital draw in d must justify the migration costs and must be better than their opportunities in any other state.

This gives rise to two separate channels by which human capital in state d changes when in-migration rates change. First, marginal migrants will always be less positively selected than the average migrant. This means that increasing the migration rate (due to a Democratic governor) will crowd in migrants with lower human capital draws than the average migrant, reducing the average human capital of migrants. This is a selection effect and its implications for wage inequality are the same as the price effects documented in Panel C. On the other hand, because all migrants are positively selected relative to non-migrants (since non-migrants' are willing to choose d even with a lower skill draw because they do not need to compensate for migration costs), increasing the share of college graduates who are migrants increases the average skills of college graduates. We refer to this as a composition effect.[36]

Panel D shows that neither of these forces is unambiguously dominant. In some states, a Democratic governor will reduce earnings inequality by up to 2.3%. On the other hand, in some states, despite the fact that "wage" inequality has fallen (Panel C), earnings inequality is expected to *rise*, sometimes by as much as 1.6%. Likewise, despite that fact that all Republican governors are expected to reduce the stock of college graduates (Panel A) and increase the relative price of college graduates human capital (Panel C), some are expected to also reduce earnings inequality by as much as 1.2%. These results are surprising because they contrast with standard intuition about the effects of a supply shock (which is the only response in our model), but they can arise because reducing the share of migrants among college graduates reduces the overall human capital level of those graduates. It is worth noting that many of

[36]It is worth noting that the best empirical evidence for the causal effect of the college share on inequality comes from Fortin (2006), who finds that an increase in the college share reduces college/non-college wage gaps. But Fortin's instruments are based on state funding policies for higher education. This means that her effects are driven by those who were educated within the state (i.e., non-migrants). Thus, they do not include the selection or composition effects discussed here, since those are specific to the selection of in-migrants. Our estimated price effects (Panel C) are consistent with her findings.

the most progressive areas in the country (e.g., California, Seattle, New York) have seen rising shares of college graduates at the same time as rising inequality. Our model can reconcile this even without agglomeration effects or housing market distortions.

6.4 Cross-state spillover effects

Finally, our model lets us study how changes in the governor in one state affect other states. All states are linked through the migration decision, since workers are deciding between different states in choosing where to live. Thus, changes in the incentives to live in one state might push those workers towards another state or pull workers away from that state.

In Figure 9, we map these spillovers. In Panel (a), we present the direct effects on college-educated workforce for the nine states we consider electing Democratic governors in a strong blue wave.[37] These effects vary, but it is worth noting that we estimate reasonably strong effects for the three Southeastern states that we consider plausible flips (Florida, Georgia, and Virginia). In Panel (b), we then plot the indirect spillover effects on the college-educated workforce in states that *do not* change their governor. Most states show small effects, but in the Southeast, some estimates of spillovers can be nearly as large as the direct effects we estimate for treated states. Thus, while spillovers are typically small, when several large states in the same area all flip at once, it can cause a substantial ripple effect throughout that region.

[Figure 9 about here.]

In Figure 10 we perform the same exercise to calculate the spillovers induced by a strong red wave. Because these states aren't concentrated in a major area the way they are in the case of the strong blue wave, they don't tend to be amplified into a large regional effect. However, some states are highly exposed to one or two particularly large states that we simulate flipping. Thus, we predict that Arizona flipping will have large effects on its smaller neighbor New Mexico, and North Carolina flipping will have large effects on its smaller neighbor South Carolina. One of the largest increases predicted by our model would be West Virginia, where we predict that the college graduate population could increase by more than 1% if its larger state neighbors of Pennsylvania and Kentucky were to flip to Republican governors.

[Figure 10 about here.]

In conclusion, our estimates suggest only a modest disutility of politics. It is not the dominant force driving migration decisions, paling in comparison to the distribution of non-wage

[37] Throughout this section, we focus on the implied spillover effects on the size of the college-educated workforce. Since this is the proximate driver of all downstream equilibrium outcomes, we find this the most straightforward to interpret.

amenities, migration costs, wages, or TFP. Nonetheless, it is substantial enough for political swings to have meaningful effects on economic outcomes, not only on the states experiencing a change but on neighboring states as well. In our simulations, we have explored the effects of changing college-educated workers' disutility parameter, and have found that effects are close to linear in $\gamma^C_{d(Rep)}$. This implies that a continuation of recent patterns in education polarization can exacerbate the effects we estimate here. A doubling of $\gamma^C_{d(Rep)}$ (which is the same amount of change that we see within our sample period, from the times when migration was non-responsive to the strong responses we find post-2015) would roughly double the effects we calculate here. As such, political divides across educational lines are important to monitor in the future.

7 Conclusion

College education is increasingly the fulcrum of political disagreement in the United States. On both social and economic issues, college graduates are well to the left of non-college voters, and to a degree much greater than just 15 years ago. As a result, Republican governance deters college graduates from moving to a state, driving down total human capital in the state, and reducing economic growth. Of course, to the extent that Republican policies themselves are pro-growth, these migration responses offset some of those gains.

What might conservative politicians do to reduce this growth penalty? One option is to increase support for universities in order to increase the stock of "home-grown" college graduates (as in Fortin (2006) or Kennan (2015)) who would presumably show greater sympathies for the policies preferred by local voters. Instead, however, Republican governors and legislatures tend to cut funding and support for universities, exacerbating the local supply shortages induced by lower in-migration.

As an alternative, conservative politicians might aim to win back the college-educated electorate. Thus far, this does not appear to be a priority. Influential party leaders tend to use "college-educated," "liberal," and "elites" as almost interchangeable terms to refer to their enemies.[38] If anything, the sort of regional inequality exacerbated by education polarization and migration appears to be a rhetorical victory for conservative politicians, who frequently emphasize these gaps to win votes. Moreover, much of the conservative appeal to college-educated voters in the past derived from pro-market low-tax policy positions. With social and economic attitudes increasingly correlated, it is unclear whether libertarian economic policies can still

[38]For example, in announcing a controversial bill to reduce public universities' teaching of liberal views on race and gender issues, Florida Governor Ron DeSantis said "Nobody wants this crap. This is an elite-driven phenomenon being driven by bureaucratic elites, elites in universities, and elites in corporate America. And they're trying to shove it down the throats of the American people... They really want to tear at the fabric of our society." (Farrington, 2021)

appeal to college graduates.

Finally, an open question is how the COVID-19 pandemic and its aftermath will interact with education polarization. We intentionally end our analysis in 2019 because we see migration data from 2020-2022 as being unreliable, but a salient feature of the post-pandemic United States is that working remotely will be common, particularly for better educated workers (Barrero, Bloom, and Davis, 2021). This is likely to exacerbate our findings, since it allows college-educated workers to live in the states they prefer while holding jobs in conservative states where the firm is located, without the general equilibrium wage pressures to offset the political migration incentives.

The implications for Republican-led states could be dramatic. Conservative states tend to raise more revenue through sales taxes, and typically have low income and corporate tax rates, while liberal states rely more on income taxation for revenue. Thus, even if the firm locates in a conservative state, the state will receive little revenue from taxing the company and little revenue from sales taxes since many of the state's workers do not actually live and consume within the state. Liberal states reliant on income taxes, on the other hand, will essentially benefit from a revenue windfall as more high-earning workers relocate there than otherwise would, given the relatively high corporate tax rates disincentivizing firms from locating there. With this in mind, future research should account for education polarization (and the plausibility that it continues to grow) when assessing the implications of post-pandemic work arrangements for heterogeneous growth in different regions of the country.

Figure 1: Differences in policy views by education

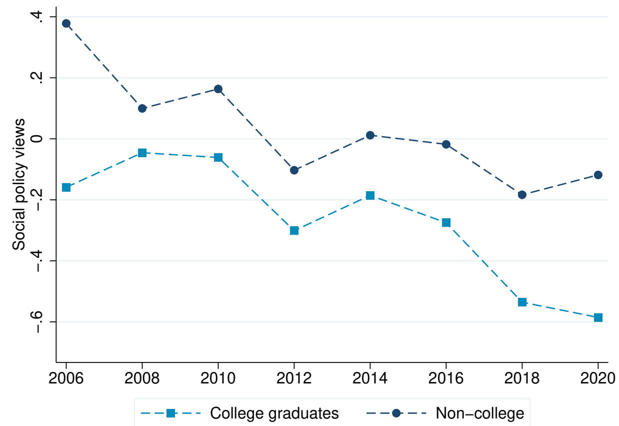

(a) Social policy views

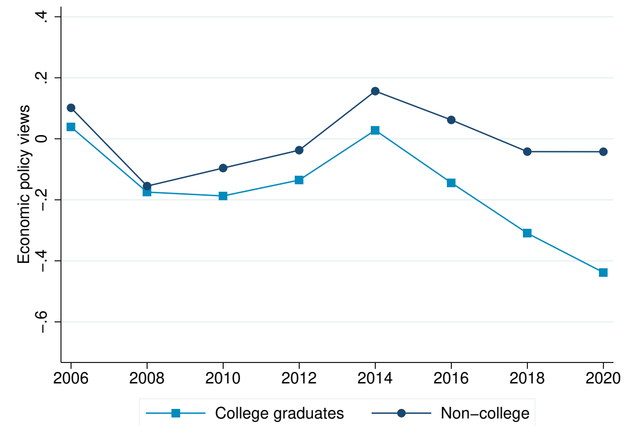

(b) Economic policy views

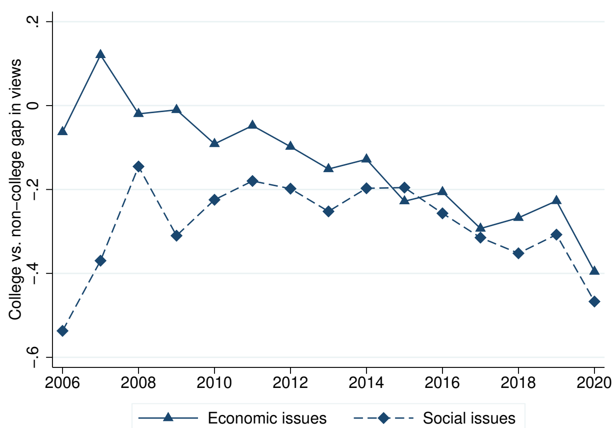

(c) College vs. Non-college gaps

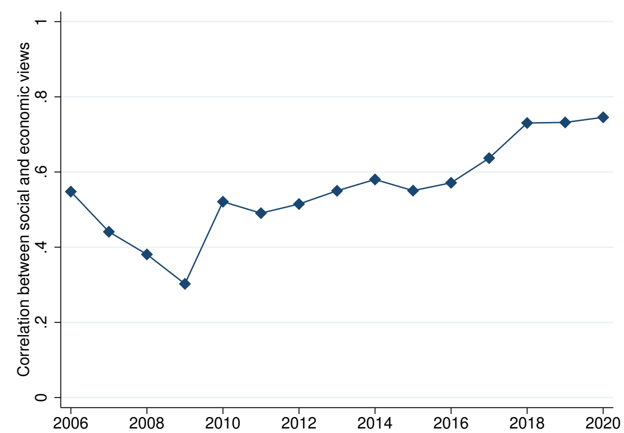

(d) Correlation between social and economic views

Notes: All calculations based on Cooperative Election Study (CES) respondents aged 26-45. Policy view indices are based on Item Response Theory estimates using the approach developed by Imai et al. (2016). Policy questions are divided into social and economic policy issues using the scheme developed by Caughey and Warshaw (2018). Indices are normalized to have mean zero and unit standard deviation across all respondents across all years, with more positive values indicating more conservative views. All calculations use sample weights.

Figure 2: Differences in policy views by earnings

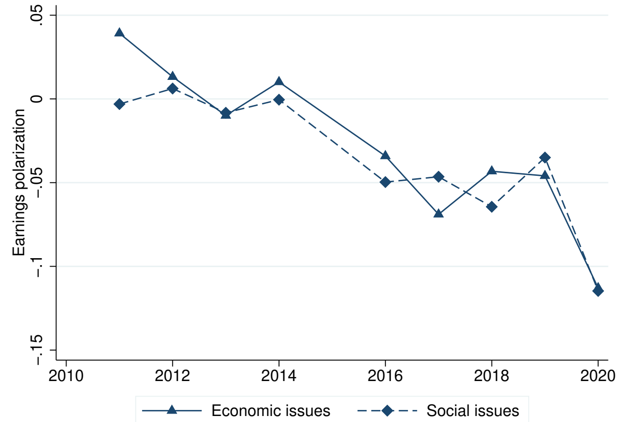

(a) Earnings polarization over time

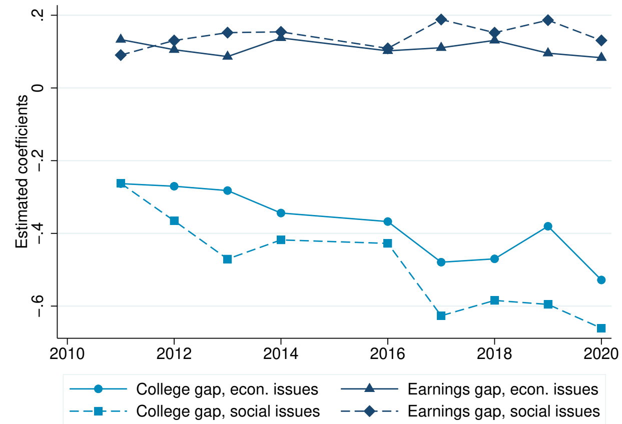

(b) Coefficients on earnings and education

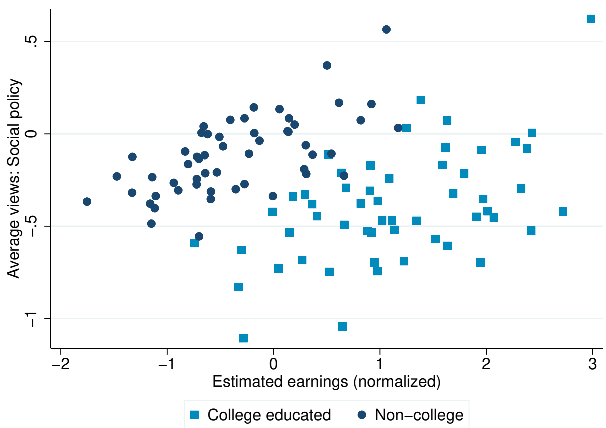

(c) Social policy views by age-industry-college

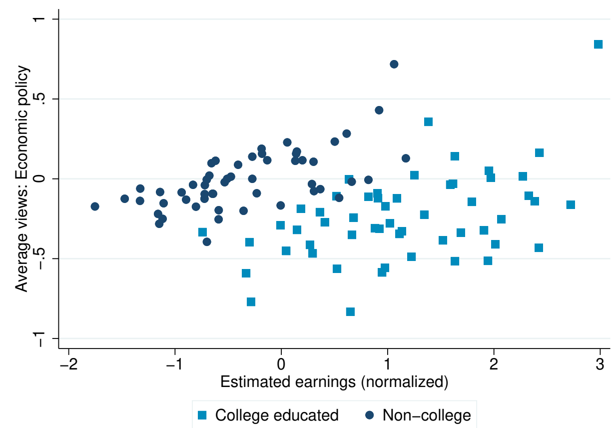

(d) Economic policy views by age-industry-college

Notes: All calculations based on Cooperative Election Study (CES) respondents aged 26-45 and employed in the private sector. Earnings are imputed from the American Community Survey (ACS) based on age, education, and industry, using the approach proposed in the text that accounts for geographic pay differences (see equation (1)). Policy view indices are based on Item Response Theory estimates using the approach developed by Imai et al. (2016). Policy questions are divided into social and economic policy issues using the scheme developed by Caughey and Warshaw (2018). Indices are normalized to have mean zero and unit standard deviation across all respondents across all years, with more positive values indicating more conservative views. All calculations use sample weights.

Figure 3: Number of gubernatorial party transitions over time

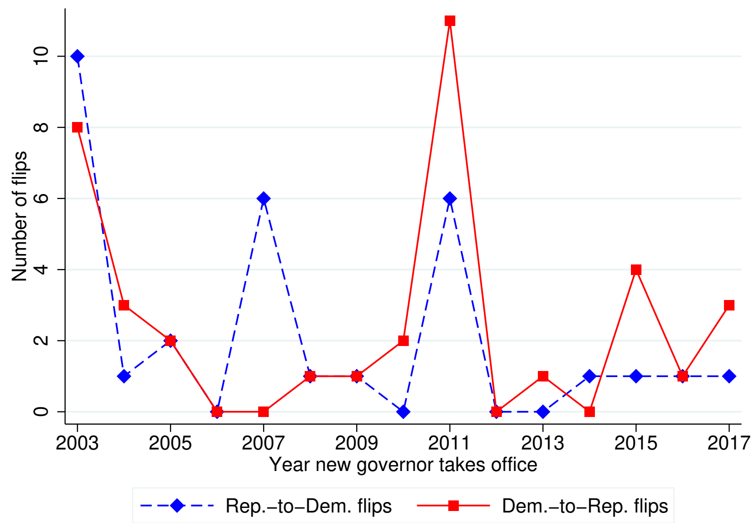

Notes: Figure plots the total number of gubernatorial transitions by type and year.

Figure 4: Democrat-to-Republican gubernatorial transitions and college in-migration

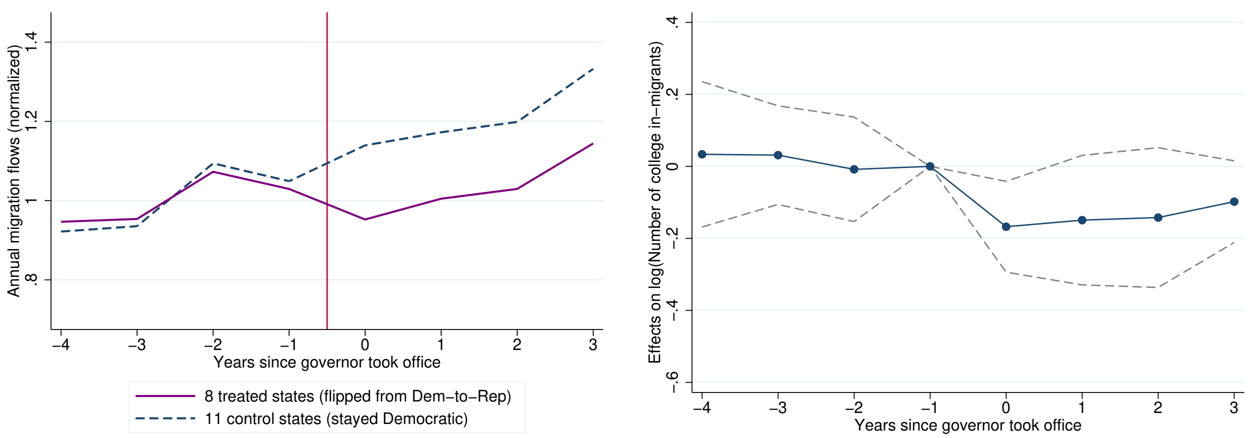

(a) Levels: Treated and control states

(b) Difference-in-difference estimates

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican between 2015 and 2017) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is never used. Panel (a): Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. Panel (b): Dependent variable is log number of collage graduate in-migrants (measured in the ACS). See column 3 of Table 1 for estimates corresponding to Panel (b).

Figure 5: Migration responses by state-of-origin

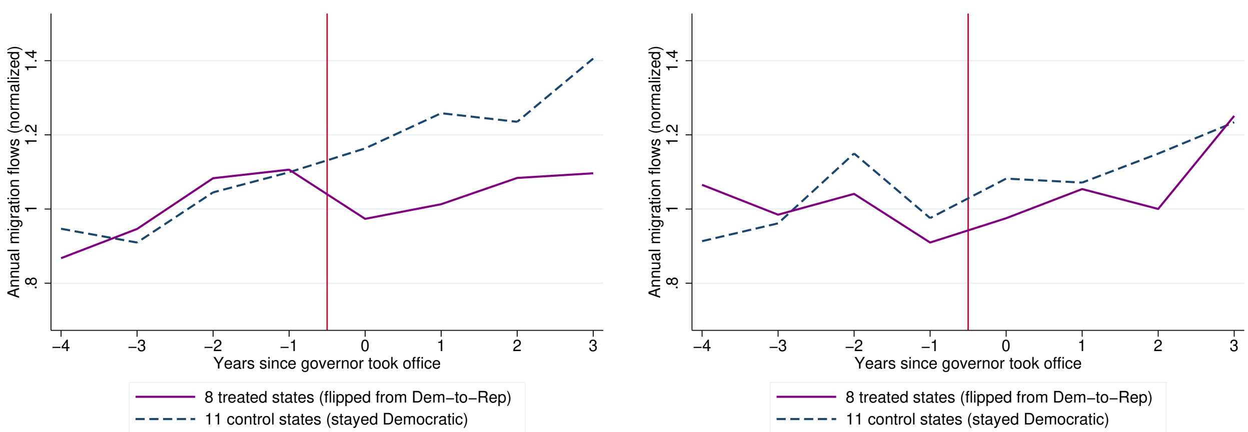

(a) In-migration from Red states

(b) In-migration from Blue states

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. In-migration is calculated separately for migration from a state-year with a Republican governor (Panel (a)) and a state-year with a Democratic governor (Panel (b)). Migration from 2020 onwards is never used.

Figure 6: College graduates' migration responses and education polarization over time

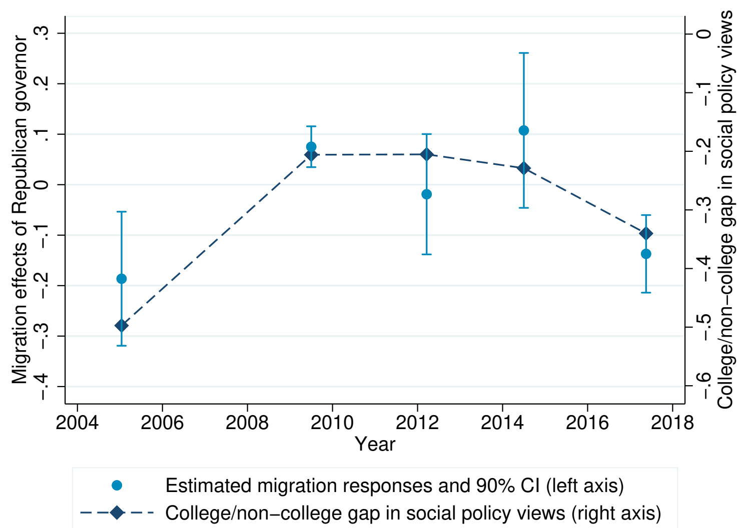

Notes: Figure plots difference-in-difference estimates (based on the Callaway-Sant'Anna-type estimator) for the effects of a Democrat-to-Republican gubernatorial transition on college graduate in-migration, separately for each three-year treatment cohort of our sample (i.e., 2003-2005, ..., 2015-2017), and the average gap in social policy views between college and non-college voters during those same years. Our calculation of average education polarization accounts for the fact that years have different numbers of treated states. For instance, in the 2009-2011 period, there were 11 times as many DR transitions in 2011 (post-treatment years: 2011-2014) as in 2009 (post-treatment years: 2009-2012). Thus, in calculating average polarization during the period, we give 11 times as much weight to the level in 2011-2014 as the level in 2009-2012. Migration data from 2020 onwards is never used.

Figure 7: Heterogeneous effects of single states switching governors

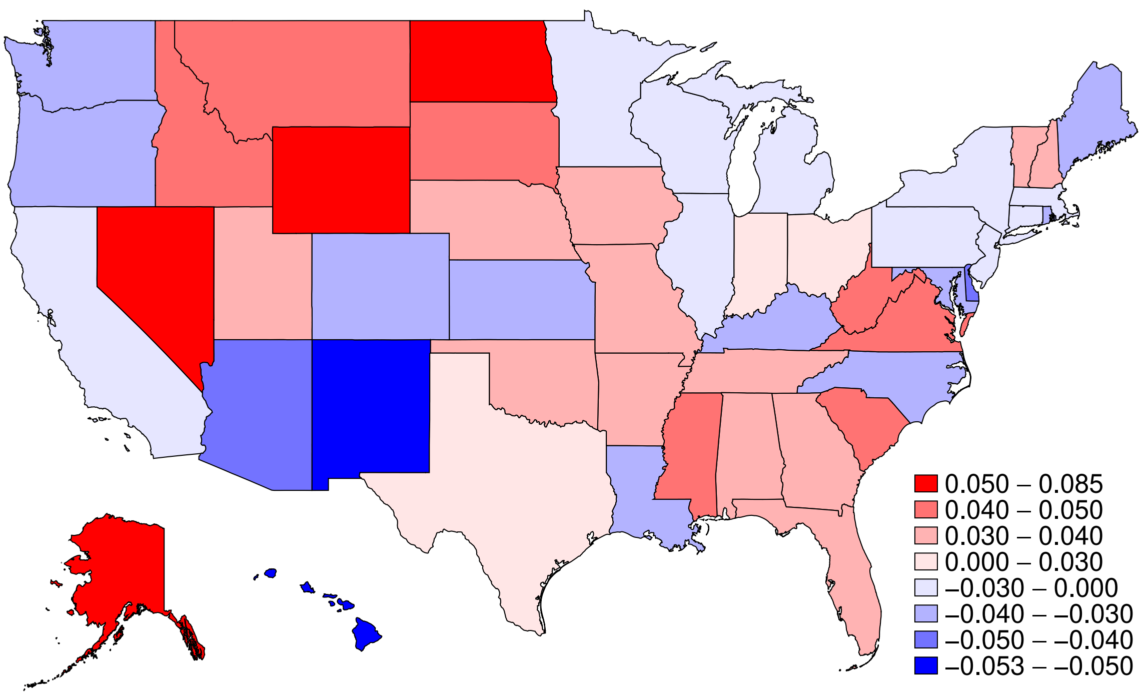

Notes: Map based on 50 separate counterfactuals, each of which flips only one state independently, relative to the set of governors in office in 2023. Red states have Republican governors in 2023, and so we simulate the effects of a Democrat taking office. Blue states have Democratic governors in 2023, and so we simulate the effects of a Republican taking office. Darker colors indicate a larger percentage change in the college graduate workforce predicted by our model.

Figure 8: Counterfactual switching states relative to 2023 gubernatorial partisanship

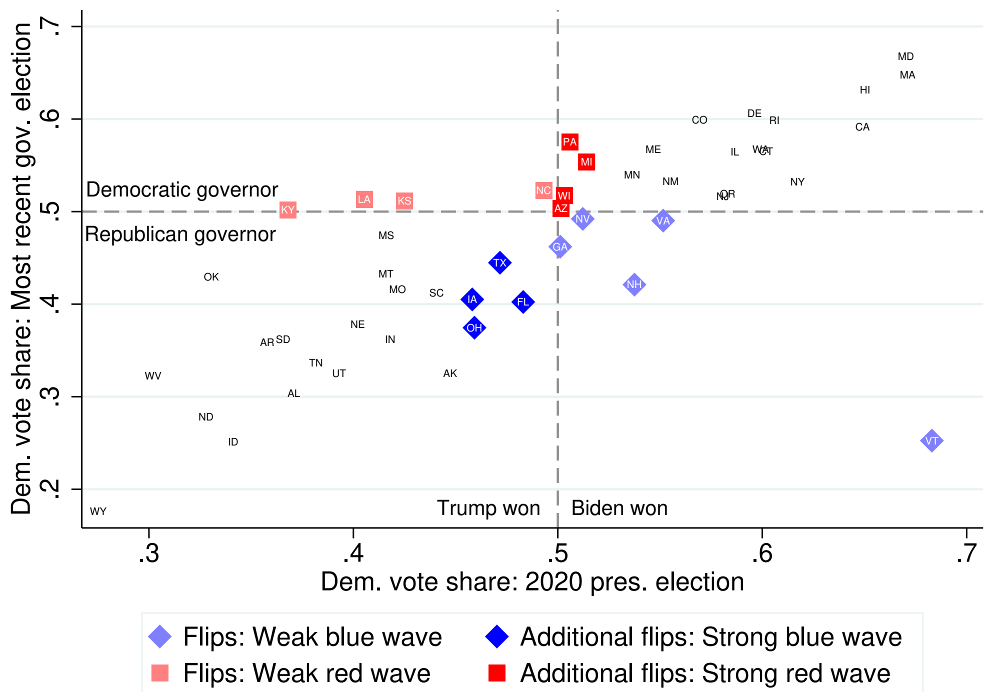

Notes: Figure plots the relationship between 2020 Presidential vote share and the most recent (as of 2023) gubernatorial election vote share. All vote shares based on two-party vote share. Figure identifies the 17 states we consider plausible flips in our four counterfactuals.

Figure 9: Direct and spillover effects of big blue wave

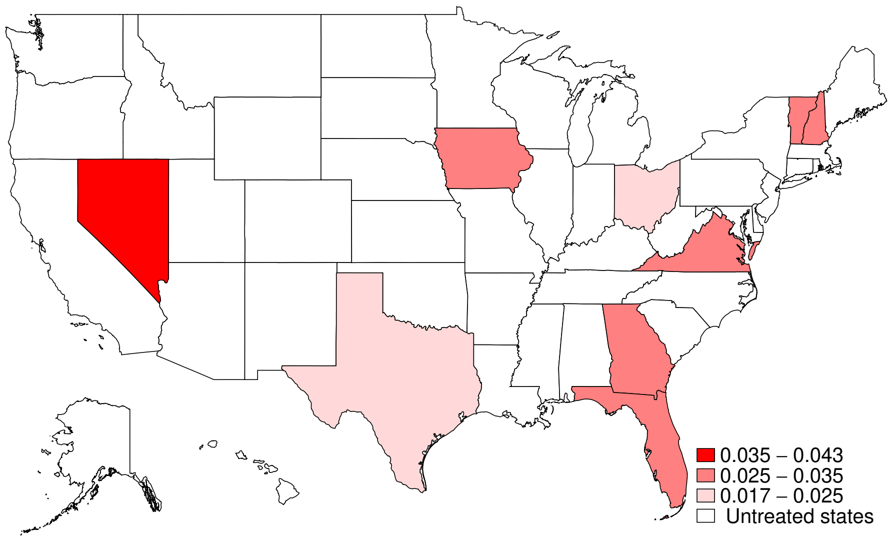

(a) Direct effects: Treated states

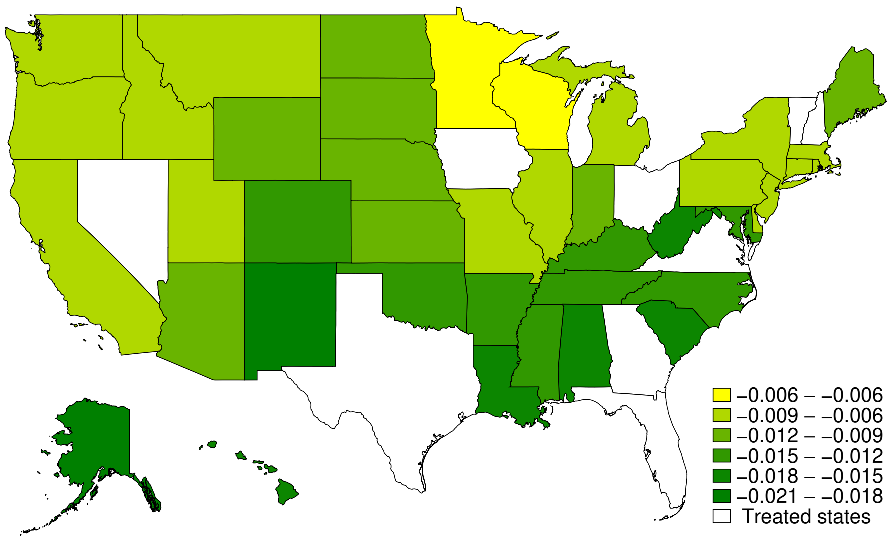

(b) Indirect spillover effects

Notes: Map based on the strong blue wave (see Figure 8) in which nine states governed by Republicans in 2023 are simulated as flipping to Democratic governors. Percentage changes in the college-educated workforce for the treated states are mapped in Panel (a), and percentage changes for un-treated states are mapped in Panel (b).

Figure 10: Direct and spillover effects of big red wave

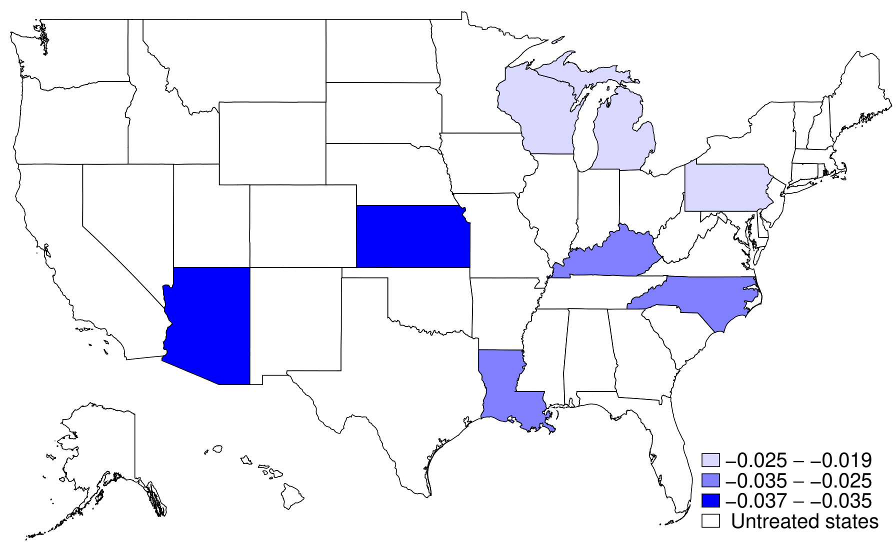

(a) Direct effects: Treated states

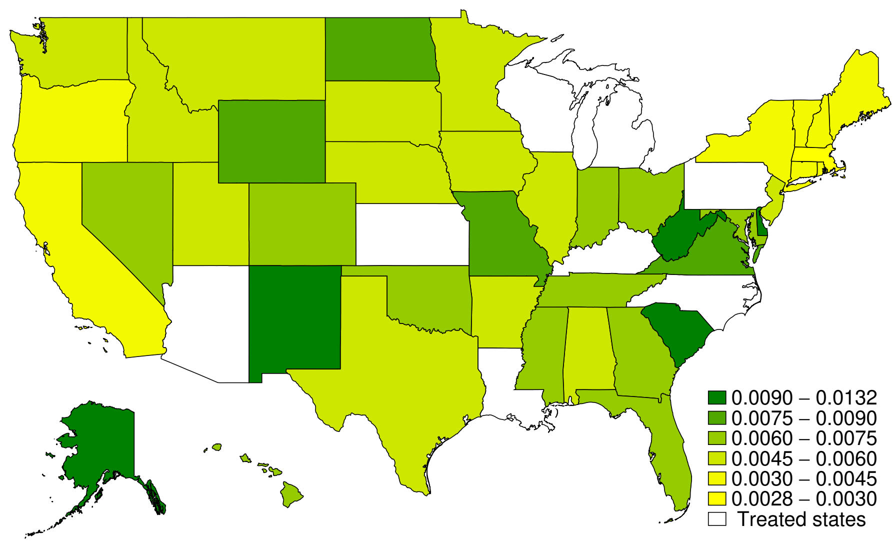

(b) Indirect spillover effects

Notes: Map based on the strong red wave (see Figure 8) in which eight states governed by Democrats in 2023 are simulated as flipping to Republican governors. Percentage changes in the college-educated workforce for the treated states are mapped in Panel (a), and percentage changes for un-treated states are mapped in Panel (b).

Table 1: Effects of Democrat-to-Republican transitions on college in-migration

	(1)	(2)	(3)	(4)	(5)	(6)
Flip years:		2015-2017			2003-2017	
Treat × Post	-0.105***	-0.137***	-0.157***	-0.077**	-0.083**	-0.117**
	(0.032)	(0.047)	(0.051)	(0.032)	(0.039)	(0.045)
R^2	0.987	0.983	0.980	0.976	0.965	0.955
N	274	274	274	853	853	853
Avg. migration rate	.027	.030	.040	.028	.033	.042
Number of states						
Treated		8			33	
Control		11			16	
Sample						
Employed	Yes	Yes	Yes	Yes	Yes	Yes
Age	18+	26+	26-45	18+	26+	26-45
Private sector		Yes	Yes		Yes	Yes
US Citizen			Yes			Yes

* $p < .10$, ** $p < .05$, *** $p < .01$. Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is log number of college graduate in-migrants (measured in the ACS). All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is never used. Standard errors are clustered at the state level.

Table 2: Parameter values

Description	Determination	Parameter	Value	
EoS across groups	Jerzmanowski-Tamura '20	σ	2.6	
Skill dispersion	Wage dispersion	θ^C	2.85	
		θ^N	2.96	
Political wedge	Reduced-form moments	$\gamma^c_{p(d)}$	0.065	
		$\gamma^N_{p(d)}$	0.010	
			Mean	Standard Deviation
Migration cost	Migration flow difference	τ^C_d	0.88	0.13
		τ^N_d	0.90	0.13
Amenities	Same party governors	α^C_d	0.81	0.17
		α^N_d	0.85	0.18
Human capital price	Market clear + common price	w^C_d	482	99.6
		w^N_d	369	83.8
TFP	Market clear + common price	A_d	269	58.3

This table reports the parameter values for the model. σ is from the literature and the rest parameters are estimated from target moments. Human capital price and TFP are reported in thousands of dollars.

Table 3: General equilibrium effects of counterfactual gubernatorial switches

	(1)	(2)	(3)	(4)
Counterfactual	Small blue	Big blue	Small red	Big red
Number of states	5	9	4	8
Panel A: Total college-educated labor force				
Average effects	0.039	0.028	-0.034	-0.028
(min, max)	(0.031, 0.051)	(0.017, 0.043)	(-0.038, -0.031)	(-0.037, -0.019)
Panel B: GDP per worker				
Average effects	0.014	0.009	-0.013	-0.011
(min, max)	(0.010, 0.019)	(0.006, 0.015)	(-0.014, -0.010)	(-0.013, -0.008)
Panel C: Unobservable wage inequality (human capital price)				
Average effects	-0.180	-0.109	0.129	0.107
(min, max)	(-0.238, -0.100)	(-0.191, -0.055)	(0.092, 0.171)	(0.076, 0.156)
Panel D: Observable wage inequality (college vs. non-college earnings)				
Average effects	-0.006	-0.004	0.001	-0.001
(min, max)	(-0.023, 0.016)	(-0.020, 0.012)	(-0.006, 0.011)	(-0.012, 0.011)

Table presents changes in equilibrium outcomes for "treated" states (i.e., those changing their governors) in our four counterfactuals (Figure 8 presents the states we simulate as flipping in each counterfactual). Panels A and B present percent changes in the college-educated workforce and GDP per worker, respectively. Panels C and D present changes in the college/non-college difference in log human capital price and log average earnings per worker, respectively.

References

Acemoglu, D. and D. Autor (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, Volume 4, pp. 1043–1171. Elsevier.

Achen, C. H. and L. M. Bartels (2017). Democracy for realists. In *Democracy for Realists*. Princeton University Press.

Albouy, D. (2016). What are cities worth? land rents, local productivity, and the total value of amenities. *Review of Economics and Statistics 98*(3), 477–487.

Ansolabehere, S., J. M. de Figueiredo, and J. M. Snyder (2003). Why is there so little money in us politics? *Journal of Economic perspectives 17*(1), 105–130.

Barnes, J. (2015). How edwards won in louisiana. *Ballotpedia 2015 Election Analysis: November 23, 2015*.

Barrero, J. M., N. Bloom, and S. J. Davis (2021). Why working from home will stick. *NBER Working Paper*.

Bass, D. N. (2012). N.c. democratic party chair david parker refuses to resign. *The Carolina Journal: April 19, 2012*.

Besley, T., T. Persson, and D. M. Sturm (2010). Political competition, policy and growth: theory and evidence from the us. *The Review of Economic Studies 77*(4), 1329–1352.

Bishop, B. (2009). *The big sort: Why the clustering of like-minded America is tearing us apart*. Houghton Mifflin Harcourt.

Bleemer, Z. (2021). Top percent policies and the return to postsecondary selectivity. *Working Paper*.

Bouton, L., J. Cagé, E. Dewitte, and V. Pons (2022). Small campaign donors. *NBER Working Paper 30050*.

Boxell, L., M. Gentzkow, and J. M. Shapiro (2020). Cross-country trends in affective polarization. *The Review of Economics and Statistics*, 1–60.

Brown, J. R., E. Cantoni, S. Chinoy, M. Koenen, and V. Pons (2023). Where you grow up shapes your political behavior: Evidence from childhood moves. *Working Paper*.

Brown, J. R., E. Cantoni, R. D. Enos, V. Pons, and E. Sartre (2022). The increase in partisan segregation in the united states. *Working paper*.

Brown, J. R. and R. D. Enos (2021). The measurement of partisan sorting for 180 million voters. *Nature Human Behaviour 5*(8), 998–1008.

Brox, E. and T. Krieger (2021). Far-right protests and migration. *Working Paper*.

Bryan, G. and M. Morten (2019). The aggregate productivity effects of internal migration: Evidence from indonesia. *Journal of Political Economy 127*(5), 2229–2268.

Cain, , M. Ward, M. Hoff, and T. Dua (2022). The country's biggest employers, including walmart and amazon, should 'say goodbye to attracting top female talent' in abortion 'trigger law' states. *Business Insider: June 24, 2022*.

Callaway, B. and P. H. Sant'Anna (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics 225*(2), 200–230.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica 82*(6), 2295–2326.

Cantoni, E. and V. Pons (2022). Does context outweigh individual characteristics in driving voting behavior? evidence from relocations within the united states. *American Economic Review 112*(4), 1226–72.

Catanese, D. (2012). Gov. perdue won't seek reelection. *Politico: January 26, 2012*.

Caughey, D. and C. Warshaw (2015). Dynamic estimation of latent opinion using a hierarchical group-level irt model. *Political Analysis 23*(2), 197–211.

Caughey, D. and C. Warshaw (2016). The dynamics of state policy liberalism, 1936–2014. *American Journal of Political Science 60*(4), 899–913.

Caughey, D. and C. Warshaw (2018). Policy preferences and policy change: Dynamic responsiveness in the american states, 1936–2014. *American Political Science Review 112*(2), 249–266.

Caughey, D., Y. Xu, and C. Warshaw (2017). Incremental democracy: The policy effects of partisan control of state government. *The Journal of Politics 79*(4), 1342–1358.

Correia, S., P. Guimarães, and T. Zylkin (2019). Verifying the existence of maximum likelihood estimates for generalized linear models. *arXiv preprint arXiv:1903.01633*.

Correia, S., P. Guimarães, and T. Zylkin (2020). Fast poisson estimation with high-dimensional fixed effects. *The Stata Journal 20*(1), 95–115.

Danieli, O., N. Gidron, S. Kikuchi, and R. Levy (2022). Decomposing the rise of the populist radical right. *Available at SSRN 4255937*.

de Benedictis-Kessner, J. and C. Warshaw (2016). Mayoral partisanship and municipal fiscal policy. *The Journal of Politics 78*(4), 1124–1138.

De Chaisemartin, C. and X. d'Haultfoeuille (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review 110*(9), 2964–96.

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980–2000. *American Economic Review 106*(3), 479–524.

Diamond, R. and C. Gaubert (2022). Spatial sorting and inequality. *Annual Review of Economics 14*, 795–819.

Eason, B. (2015). Ballard, council to legislature: Repeal law, protect lgbt from discrimination. *The Indianapolis Star: March 30, 2015.*

Farrington, B. (2021). Desantis: Critical race theory is "crap," vows to fight it. *AP News: December 15, 2021.*

Ferreira, F. and J. Gyourko (2009). Do political parties matter? evidence from us cities. *The Quarterly journal of economics 124*(1), 399–422.

Firoozi, D. (2022). The effect of research universities on student partisanship and turnout. *Job Market Paper*.

Ford, B. Q., M. Feinberg, B. Lassetter, S. Thai, and A. Gatchpazian (2023). The political is personal: The costs of daily politics. *Journal of Personality and Social Psychology*.

Fortin, N. M. (2006). Higher-education policies and the college wage premium: Cross-state evidence from the 1990s. *American Economic Review 96*(4), 959–987.

Fowler, A. and A. B. Hall (2018). Do shark attacks influence presidential elections? reassessing a prominent finding on voter competence. *The Journal of Politics 80*(4), 1423–1437.

Fowler, A. and B. P. Montagnes (2015). College football, elections, and false-positive results in observational research. *Proceedings of the National Academy of Sciences 112*(45), 13800–13804.

Gethin, A., C. Martínez-Toledano, and T. Piketty (2022). Brahmin left versus merchant right: Changing political cleavages in 21 western democracies, 1948–2020. *The Quarterly Journal of Economics 137*(1), 1–48.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics 225*(2), 254–277.

Grumbach, J. (2022). *Laboratories against democracy: How national parties transformed state politics*, Volume 182. Princeton University Press.

Hagelgans, A. and S. Basi (2022). Roe v. wade's demise is a turning point for corporate america. *Harvard Business Review: June 30, 2022*.

Haidt, J. and G. Lukianoff (2021). The polarization spiral. *Persuasion: October 29, 2021*.

Healy, A. J., N. Malhotra, and C. H. Mo (2010). Irrelevant events affect voters' evaluations of government performance. *Proceedings of the National Academy of Sciences 107*(29), 12804–12809.

Hersh, E. (2020). *Politics is for power: How to move beyond political hobbyism, take action, and make real change*. Simon and Schuster.

Hinckley, B., R. Hofstetter, and J. Kessel (1974). Information and the vote: A comparative election study. *American Politics Quarterly 2*(2), 131–158.

Hopkins, D. A. (2017). *Red fighting blue: How geography and electoral rules polarize American politics*. Cambridge University Press.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The allocation of talent and us economic growth. *Econometrica 87*(5), 1439–1474.

Imai, K., J. Lo, and J. Olmsted (2016). Fast estimation of ideal points with massive data. *American Political Science Review 110*(4), 631–656.

Jerzmanowski, M. and R. Tamura (2020). Aggregate elasticity of substitution between skills: estimates from a macroeconomic approach. *Macroeconomic Dynamics*, 1–31.

Jones, P. E. (2020). Partisanship, political awareness, and retrospective evaluations, 1956–2016. *Political Behavior 42*(4), 1295–1317.

Kaplan, E., J. L. Spenkuch, and R. Sullivan (2022). Partisan spatial sorting in the united states: A theoretical and empirical overview. *Journal of Public Economics 211*, 104668.

Kennan, J. (2015). Spatial variation in higher education financing and the supply of college graduates. *NBER Working Paper*.

Keshner, A. (2021). 'they've been silent': Many companies are on the sidelines about the new texas abortion law — will they stay there? *Market Watch: September 3, 2021.*

Kleinman, B., E. Liu, and S. J. Redding (2023). Dynamic spatial general equilibrium. *Econometrica 91*(2), 385–424.

Leins, C. (2019). The most popular and least popular governors. *US News: July 19, 2019.*

Leonhardt, M. (2021). Hard right turn: Companies are toeing the political line in business-friendly red states. *Fortune: October 2, 2021.*

Martin, G. J. and S. W. Webster (2020). Does residential sorting explain geographic polarization? *Political Science Research and Methods 8*(2), 215–231.

Mason, L. (2015). "i disrespectfully agree": The differential effects of partisan sorting on social and issue polarization. *American journal of political science 59*(1), 128–145.

McAfee, T. (2015). Prostitutes, spying and strip joints have dominated louisiana governor's race – now voters will finally pick bobby jindal's replacement. *People Magazine: November 20, 2015.*

McConnell, C., Y. Margalit, N. Malhotra, and M. Levendusky (2018). The economic consequences of partisanship in a polarized era. *American Journal of Political Science 62*(1), 5–18.

Mijs, J. J. and E. L. Roe (2021). Is america coming apart? socioeconomic segregation in neighborhoods, schools, workplaces, and social networks, 1970–2020. *Sociology Compass 15*(6), e12884.

Monras, J. (2020). Economic shocks and internal migration.

Moretti, E. (2004). Workers' education, spillovers, and productivity: evidence from plant-level production functions. *American Economic Review 94*(3), 656–690.

Moretti, E. (2012). *The new geography of jobs.* Houghton Mifflin Harcourt.

Moretti, E. and D. J. Wilson (2017). The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *American Economic Review 107*(7), 1858–1903.

Mummolo, J. and C. Nall (2017). Why partisans do not sort: The constraints on political segregation. *The Journal of Politics 79*(1), 45–59.

Ortiz, F. (2016). Scars of ferguson protests shape missouri governor's race. *US News: July 29, 2016*.

Patterson, C. (2020). The matching multiplier and the amplification of recessions. *Working paper*.

Perez-Truglia, R. (2018). Political conformity: Event-study evidence from the united states. *Review of Economics and Statistics 100*(1), 14–28.

Perez-Truglia, R. and G. Cruces (2017). Partisan interactions: Evidence from a field experiment in the united states. *Journal of Political Economy 125*(4), 1208–1243.

Ruggles, S., S. Flood, R. Goeken, M. Schouweiler, and M. Sobek (2022). Ipums usa: Version 12.0 [dataset]. *https://doi.org/10.18128/D010.V12.0 Minneapolis, MN: IPUMS*.

Seifter, M. (2017). Gubernatorial administration. *Harvard Law Review 131*, 483.

Silva, J. S. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and statistics 88*(4), 641–658.

Smith, K. B., M. V. Hibbing, and J. R. Hibbing (2019). Friends, relatives, sanity, and health: The costs of politics. *PloS one 14*(9), e0221870.

Squire, P. and C. Fastnow (1994). Comparing gubernatorial and senatorial elections. *Political Research Quarterly 47*(3), 705–720.

Yglesias, M. (2019). The great awokening. *Vox: April 1, 2019*.

A Data

[Table A1 about here.]

[Table A2 about here.]

Table A1: Social Policy Questions

Category	Number of questions ever	Number of questions each year	Most common question
Abortion and birth control	5	2.7	Do you support banning abortion after the 20th week of pregnancy?
Affirmative action	1	0.7	Do you support affirmative action programs that give preference to racial minorities in employment and college admissions in order to correct for past discrimination?
Crime and policing	12	1.1	Do you support eliminating mandatory minimum sentences for non-violent drug offenders?
Education	3	0.2	Do you support the Student Success Act, which would end more than 70 federal education programs and decentralize education decision-making?
Gender and sexuality	7	1.5	Do you favor or oppose allowing gays and lesbians to marry legally?
Gun control	6	2.1	Do you support banning assault rifles?
Immigration	23	4.7	Do you support increasing the number of border patrols on the US-Mexico border?
Impeachment	2	0.1	Do you support removing President Trump from office for abuse of power?
Supreme Court	5	0.4	Do you support confirming Brett Kavanaugh to become a Justice of the Supreme Court?

Table displays characteristics of economic policy questions. Topics are classified as economic policy based on scheme developed by Caughey and Warshaw (2018). Many of the questions appear in grid format, where the literal question itself is simply the last part of a longer sentence (the first part being shown above). We have shortened and re-worded the questions we present here into standard language (also sometimes eliminating preambles before the questions). It is sometimes difficult to cleanly separate questions across these categories (e.g., a question related to taxes *and* government spending, or to immigration *and* crime), however our analysis does not use these category labels *within* social/economic policy groups. It is typically simple to determine whether a question belongs to the social group or the economic group. “Number of questions ever” refers to the total number of unique questions appearing about the topic between 2006 and 2020. “Number of questions each year” refers to the average number of questions asked about the topic during any given year, 2006-2020.

Table A2: Economic Policy Questions

Category	Number of questions ever	Number of questions each year	Most common question
Deregulation	2	0.3	Would you support an executive action requiring that with each new regulation enacted, two must be cut?
Environment	10	2.9	Do you support giving the Environmental Protection Agency the power to regulate Carbon Dioxide?
Government spending	23	4.7	If your state were to have a budget deficit this year, what would you prefer more, raising taxes or cutting spending?
Health care	14	2.3	Do you support repealing the Affordable Care Act?
Minimum wage	2	0.4	Do you support raising the minimum wage to $15 an hour?
Taxes	5	2.1	If the state had to raise taxes, what share of the tax increase should come from increased income taxes and what share from increased sales taxes?
Trade	8	1.2	Should the United States withdraw from the Trans-Pacific Partnership

Table displays characteristics of economic policy questions. Topics are classified as economic policy based on scheme developed by Caughey and Warshaw (2018). Many of the questions appear in grid format, where the literal question itself is simply the last part of a longer sentence (the first part being shown above). We have shortened and re-worded the questions we present here into standard language (also sometimes eliminating preambles before the questions). It is sometimes difficult to cleanly separate questions across these categories (e.g., a question related to taxes *and* government spending, or to immigration *and* crime), however our analysis does not use these category labels *within* social/economic policy groups. It is typically simple to determine whether a question belongs to the social group or the economic group. “Number of questions ever” refers to the total number of questions about the topic ever asked from 2006-2020. “Number of questions each year” refers to the average number of questions asked about the topic during any given year, 2006-2020.

B Additional results

B.1 Additional figures and tables

B.1.1 Descriptive facts

[Figure B1 about here.]

[Figure B2 about here.]

[Figure B3 about here.]

[Figure B4 about here.]

[Figure B5 about here.]

B.1.2 Causal effects

[Figure B6 about here.]

[Figure B7 about here.]

[Figure B8 about here.]

[Figure B9 about here.]

[Figure B10 about here.]

B.1.3 Structural analysis results

[Figure B11 about here.]

[Table B1 about here.]

[Table B2 about here.]

B.2 Further analysis

B.2.1 Explaining education polarization

Our goal in this paper is not to explain the rise of education polarization. Nonetheless, some features of our data do speak to some potential explanations.

First, many conservatives argue that universities are increasingly sites of indoctrination of leftist ideas. Thus, education polarization could reflect either a "treatment" effect of attending college, or a "selection" effect as conservatives are deterred from attending.[39] We find very limited support for this. In panel (a), we plot education polarization estimated separately by year and cohort, among cohorts turning 26 from 2008-2020.

[Figure B12 about here.]

For the period of 2012-2016, we find modest evidence supporting the "indoctination" claim: The college/non-college gap among new cohorts grew by roughly $.1\sigma$, while the gap among already-educated cohorts stayed constant. From 2016-2020, however, there is no support for this hypothesis. This period is where the dramatic increase in education polarization was concentrated, and that increase was identical between new cohorts and already-educated cohorts. During this period, each of the five cohorts that had already turned age 26 saw the college/non-college gap widen by roughly $.3\sigma$, much larger than the modest between-cohort effects from 2012-2016. Thus, while there is some support for either changes in universities' treatment effects or changes in selection into universities, these effects are modest relative to the overall growth in college/non-college gaps. Most of the growth (and all of the recent growth) instead reflects changes in the political environment.

One of the leading changes in the political environment that has garnered the most attention is changes in the relative importance of social and economic issues (Danieli et al., 2022; Gethin et al., 2022). Similar to Danieli et al. (2022), who focus on Europe, in panel (b) we predict whether a respondent identifies as a Republican (as opposed to a Democrat and excluding Independents) using our indices for social and economic policy views. We do this separately by education and year, and plot the resulting coefficients.

Consistent with this view, the relative importance of social issues for explaining partisan affiliation has increased, while the relative importance of economic issues has declined (for both college-educated and non-college voters). In 2010, a 1σ increase in conservative views on economic issues increases the probability of identifying as a Republican by 22 percentage points,

[39] Little academic evidence has assessed this explanation because of both measurement and identification challenges. One important exception is Firoozi (2022), who estimates the causal effect of shifting from a less prestigious university to a more prestigious one, and finds sizeable impacts on political activity.

while a 1σ increase in social issues conservatism implies only a 11-13pp increase. By 2020, this had flipped, and a 1σ increase in economic issues implies only a 10-15pp increase in Republican identification, compared to a 20-22pp increase from a 1σ increase in social issue conservatism.[40] Thus, for both college-educated and non-college voters, social issues have become more important and economic issues have become less important. Given that education polarization has always been larger for social issues than economic issues, this does lead to some widening in partisan preferences between the two groups. However, as we showed above, there has been widening polarization on both sets of issues independently and, moreover, the correlation between views has risen, making it less and less important to distinguish between them. Thus, while these shifts in priorities exacerbate education polarization, they do not explain the overall increase.

Finally, in an influential book, Bishop (2009) presents evidence that liberals and conservatives increasingly live in different places, and argues that a decline in interpersonal interactions drives polarization. Given the growth of sorting across different metropolitan areas by education (Diamond and Gaubert, 2022), as well as rising socioeconomic segregation at the neighborhood level (Mijs and Roe, 2021), it seems natural to connect these hypotheses. Indeed, segregation by political affiliation has increased for virtually any geography (Brown et al., 2022), and neighbors and communities do appear to have effects on individuals' political views (Cantoni and Pons, 2022; Martin and Webster, 2020; Perez-Truglia and Cruces, 2017; Perez-Truglia, 2018).

Consistent with this, panel (c) shows that the states with the highest college shares in 2005-2010 also saw the greatest growth in education polarization on social issues from 2010-2020. While this is consistent with an "echo chamber" effect, the magnitude is relatively small and only marginally significant, and the same effects do not show for polarization on economic issues, although this polarization has been equally extreme.

B.2.2 Heterogeneity by college major

As a "sanity check," it is useful to know whether the effects are driven by college graduates with more left-leaning majors. As we have emphasized throughout, the gap in average views between college and non-college voters is large, but there are meaningful variations within different types of college-educated voters.

To characterize these differences, we match the field of degree reported by college graduates in the ACS from 2009 to data on average political views by major from UCLA's Higher Education Research Institute (HERI). Specifically, we use the public use files from HERI's

[40] Interestingly, while college and non-college voters show nearly identical *levels* for the two issues in both 2010 and 2020, the timing of the change has been different. The shift from economic to social issues among non-college voters began following 2010, while for college graduates it began only after 2016 (and was very rapid).

College Senior Survey, which asks students "How would you characterize your political views?" and offers five options ranging from "Far left" to "Far right." For our purposes, there are two substantial limitations of the HERI data. First, publicly available data is from 1994-2008, which only modestly overlaps our main sample. Second, self-reported political identification is always relative to one's political context (in this case, other college majors at that specific point in time), which makes this a much more difficult metric to compare over time or with non-college respondents than our index above, which was based on actual policy views. Figure B13 characterizes some of the variation in political leanings across different college majors, as well as showing the relationship with median earnings (which is generally weak). Overall, many of the majors which lean furthest to the left and the right are intuitive examples, such as the strong left-lean of sociology, ethnic studies, and environmental sciences, or the far right lean of military studies, agricultural studies, and theology.

[Figure B13 about here.]

We split college majors into above and below median based on the share of HERI respondents who identify as either "Far left" or "Liberal." In Figure B14 we then study the migration effects for these two different types of degree holders. We find significant responses for both sets of college graduates (consistent with our arguments above that the *overall* differences in views are quite stark), but they are somewhat larger for graduates holding more left-leaning majors.

[Figure B14 about here.]

B.2.3 Alternative identification strategies

One potential concern is reverse causality; shocks to migration incentives might affect election outcomes, for instance by changing economic conditions or shifting the composition of the electorate. Here, we present two alternative identification strategies which are immune to these concerns.

We first use an instrumental variables (IV) strategy that builds on two insights: First, American politics is very "nationalized," and voters state and local vote choices largely reflect views on salient national controversies (Grumbach, 2022), which change over time. Second, states differ in the timing of their gubernatorial elections, which are largely pre-determined (barring deaths, mid-term resignations, etc.). As a result, states differ in the national political mood that happens to prevail at the time in which they hold their gubernatorial election, and this has effects on election outcomes.

The main intuition for our IV strategy is given in Figure B15, which shows *i*) the share of Democratic-led states holding an election during the year in which the Republican candidate wins, and *ii*) Republicans' share of the *Congressional* vote (i.e., not gubernatorial) in "purple" states during the year.[41] The over-time correlation between these two series is .78, suggesting that Republican gubernatorial candidates do well during the years when Republican Congressional candidates do well.[42] Our IV strategy takes advantage of the fact that while all states and districts elect their *Congressional* Representative every two years, 48 states elect their *Governor* only once every four years, and with staggered timing set well in advance. For example, six states with Democratic Governors held gubernatorial elections in 2008, a year in which Congressional Republicans were dominated by Democrats amidst a worsening recession (Republicans received less than 43% of the vote, the within-sample minimum). None of those six states flipped to a Republican governor. However, two years later, amid the "Tea Party wave," 19 states with Democratic governors held gubernatorial elections. Republicans captured 51% of the vote in purple states, and nearly 60% of Democrat-led states holding gubernatorial elections that year saw their state flip to Republicans. It is plausible that many of the Democratic governors winning in 2008 would have been defeated if they had been unlucky enough for their state's election to be in 2010 instead. Our IV strategy takes advantage of this insight.

[Figure B15 about here.]

Specifically, we instrument for a D-to-R flip using the interaction between the time-specific nationwide Republican vote share for Congressional candidates and an indicator for whether the state's regular gubernatorial election schedule fell in that year.[43] We control separately for year effects (absorbing the time-specific Congressional vote share) and whether the state had a regular gubernatorial election during that year, and so identification is driven solely by the above interaction; in other words, identification comes only from idiosyncratic cross-state variation in whether the regularly scheduled election fell on a "bad year" for Democrats:

$$Flip_{st}^{DR} = \delta_t + \beta_1 ElectYear_{st} + \beta_2 \big(ElectionYear_{st} \times RepVoteShareCongress_t \big) + \varepsilon_{st} \quad (14)$$

where again, $RepVoteShareCongress_t$ is the share of the Congressional vote going to Republicans during the year, among all "purple" states which experience both Democratic and Republican governors during the sample (i.e., the sample of "purple" states is fixed across time). The estimation sample is restricted to states with a Democratic governor at time t, since it is otherwise impossible to transition to a Republican.

[41] We define "purple" states as those we observe with both a Democratic and a Republican governor during our 20 year period. Thus, the definition of "purple" states does not change year-to-year.

[42] Interestingly, Congressional voting patterns among the same states are far less predictive of R-to-D transitions (also shown in Figure B15). Thus, we cannot instrument for those flips.

[43] Gubernatorial transitions and special elections can occur during other years because of deaths, resignations, etc. We focus on the years of regularly scheduled elections.

In the second stage, we estimate the change in migration rates, as a function of this instrumented D-to-R flip. Specifically, the second stage estimating equation is given by:

$$\Delta Mig_{st} = \theta_t + \gamma_1 ElectYear_{st} + \gamma_2 \hat{Flip}_{st}^{DR} + \nu_{st} \tag{15}$$

where ΔMig_{st} is the average annual log inflow of college-educated workers during the next four years (i.e., post-flip) minus the average annual log inflow during the previous four years (i.e., pre-flip). In this way, our IV strategy still identifies only within-state changes in migration.

Our second alternative identification strategy is to use a standard regression discontinuity design (RDD), in which identification comes solely from close election outcomes. In this regression, we restrict only to state-years with a close election, and predict the pre/post change in migration as a function of the Republican candidate's vote share, whether the Republican won, and the interaction. This estimating equation is given by:

$$\begin{aligned} \Delta Mig_{st} = (\delta_t +)\ \beta RepWins_{st} + \gamma_1 RepVoteShare_{st} \\ + \gamma_2 \big(RepVoteShare_{st} \times RepWins_{st}\big) + \varepsilon_{st} \end{aligned} \tag{16}$$

where $RepVoteShare$ is "centered" at zero so that the core coefficient of interest, β, reflects the change in migration for an asymptotically close election won by a Republican. We use both the biased adjusted estimator from Calonico, Cattaneo, and Titiunik (2014) with a triangular kernel and no controls, as well as a standard OLS estimate on the Calonico et al. (2014)-selected bandwidth with a uniform kernel and year fixed effects.

Table B3 presents the results from our main difference-in-difference specification, as well as those from the IV specification and the RDD specifications. Panel A presents estimates from the full sample (all flips occurring between 2003 and 2017), while Panel B restricts to the most recent 10-year period (flips from 2008-2017). Unfortunately, we cannot restrict further to even more recent years because our IV strategy loses power in the first stage (see Figure B15) and the RDD sample would fall to too few close elections.

Overall, IV and RDD point estimates are very similar to our main specification. For the full sample period, the four estimates range from an 8 log point decline to an 11.7 log point decline.[44] The key difference is that these alternative identification strategies are far less precise, and none yield a statistically significant estimate. It is worth noting that each Panel A estimate would be significant if it had the same standard error as our main specification does. Our key emphasis, then, is that Table B15 does not seem to suggest that our primary identification strategy over-estimates the effects on college-educated migration. Instead, it seems that the

[44] Above, we presented the difference-in-difference plot for the 2015-2017 flip period. Appendix Figure B10 shows the corresponding plot for all flips 2003-2017.

main difference between our preferred specification and these more conservative ones is that our preferred specification has much more statistical power.

[Table B3 about here.]

While our full-sample difference-in-difference estimate is slightly larger than that generated by the alternative identification strategies, this is no longer true in the more recent time period, where alternative strategies generate both larger and smaller estimated effects (again, none of which are statistically significant).

Beyond point estimates, it is also important to know whether the over-time patterns we identify hold for these alternative strategies. Because both strategies inherently have lower statistical power than our preferred difference-in-difference specification, we cannot estimate over-time changes as flexibly as we have done above (our instrument becomes very weak and the RDD sample becomes very small). Instead, in Figure B16, we present estimates for six 10-year rolling windows. In general, both the IV strategy and the RDD strategy generate an inverted-U-shaped pattern like the difference-in-difference specification does, with the exception of the RDD estimate in the final time period.

[Figure B16 about here.]

B.2.4 Alternative explanations

We interpret the decline in in-migration of college graduates after a Democrat-to-Republican gubernatorial transition as being driven by those workers' political preferences. Here, we consider two possibilities: *i*) effects on economic activities and the incentives to migrate and *ii*) effects on perceptions of citizens' political views.

First, we have argued that college graduates have a general preference against Republican Governors. This can reduce migration via a labor supply effect, and this is our primary interpretation. It is possible, instead, that Republican Governors reduce migration through a labor *demand* effect, such as through contractionary spending cuts. It is worth keeping in mind that economic contractions tend to be more experienced by non-college workers than college graduates (e.g., Patterson (2020)), and that Republican Governors tend to lead to more pro-growth business friendly policy Caughey et al. (2017). Thus, ex ante, one would expect a Republican Governor to *increase* labor demand, and one would also expect any economic incentives to be more pronounced for non-college workers than college graduates. Nonetheless, this is a concern worth exploring.

We use data from the Bureau of Labor Statistics' Job Openings and Labor Turnover Survey (JOLTS) to estimate effects on job openings and hiring, proxies for labor demand.[45] If Republican Governors implement contractionary spending cuts, we would expect job openings to go down. We do not find this. Panel (a) of Figure B17 shows that job openings are unaffected (if anything, rising in the very short-term). In panel (b), when we consider effects on hiring, we do find a small non-significant immediate decline in hiring. Recall that we found immediate migration responses during that year. Thus, our JOLTS results are consistent with a labor supply response in which fewer in-migrants make it more difficult to fill an unchanging number of vacancies.

[Figure B17 about here.]

Second, it is possible that the election of a Republican governor reduces in-migration of college graduates not because those graduates are averse to conservative *policies*, but because the election signals the prevalence of conservative *voters*, and college graduates do not want to live near these conservative voters. This matters for policy because it suggests that the actual decisions of Republican governors about what to do once in office have no bearing, the penalty emerges simply from being elected. Stated choice experiments over neighborhoods do show that voters prefer living near those who are ideologically similar, and this could explain why we find effects on in-migration but not out-migration. Nonetheless, we are skeptical of this explanation, since estimated preferences for same-ideology neighbors have been fairly small (Mummolo and Nall, 2017) and anyway, most cities (even in very conservative states) have liberal neighborhoods that residents can choose.

Nonetheless, we can test for the effects of shifts in *perceptions* – rather than policy – by estimating the effects of flips in the Presidential vote. Since all states have the same President, a shift from supporting the Democrat to supporting the Republican does not actually differentially affect the policies in that state; all states are affected. However, these electoral outcomes are highly salient, so it should have the same signalling effects about the voters that electing a Republican governor does. In the 2016 Presidential election, Donald Trump flipped six states that Barack Obama had won in 2012. We estimate the effects of these "flips" on college workers' in-migration in Figure B18. We find no effects on college graduates' in-migration, though we note the $t = -3$ coefficient suggests some caution in interpreting the pre-trends. It is important to note that this specification as well as time frame are directly comparable to our main estimates of 2015-2017 flips presented above in Figure 4. Overall, then, we conclude that the evidence is more consistent with the importance of Republican *governors* who, as we have discussed, have meaningful influence on policy (Caughey et al., 2017; Seifter, 2017).

[45] Unfortunately, at the state level, the BLS does not make such data available separated by industry or any other proxy for education.

[Figure B18 about here.]

Figure B1: Inter-state migration rates by age and education (2005-2019)

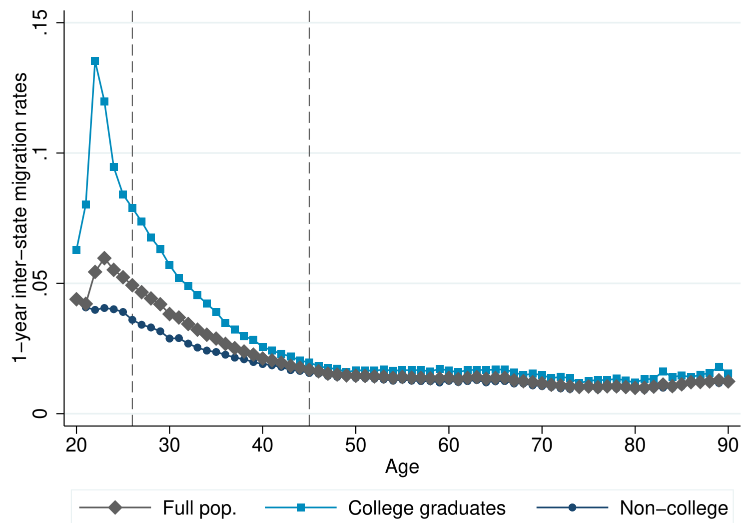

Notes: Calculations based on one-year inter-state migration rates from the American Community Survey (ACS). All calculations use sample weights.

Figure B2: College vs. Non-college gaps in median views

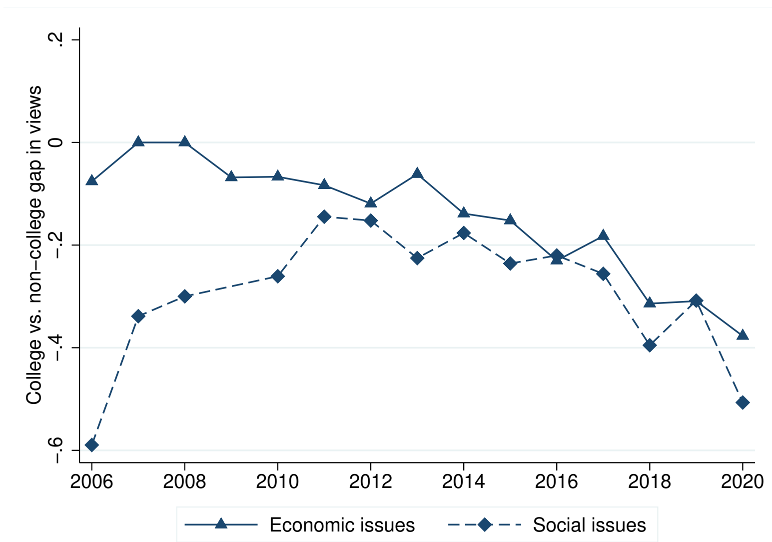

Notes: All calculations based on Cooperative Election Study (CES) respondents aged 26-45. Policy view indices are based on Item Response Theory estimates using the approach developed by Imai et al. (2016). Policy questions are divided into social and economic policy issues using the scheme developed by Caughey and Warshaw (2018). Indices are normalized to have mean zero and unit standard deviation across all respondents across all years, with more positive values indicating more conservative views. All calculations use sample weights.

Figure B3: Distribution of social policy views: 2010-2020

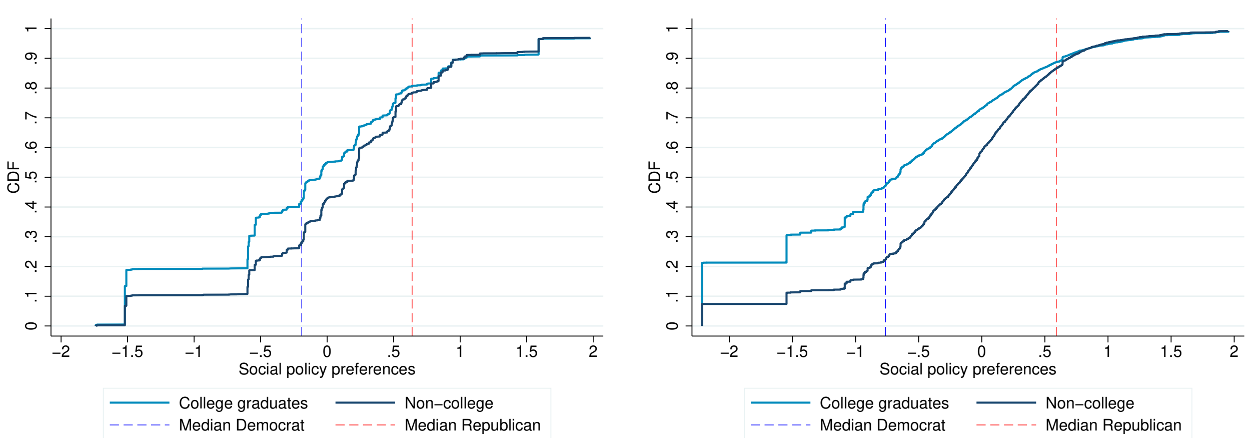

(a) 2010

(b) 2020

Notes: All calculations based on Cooperative Election Study (CES) respondents aged 26-45. Policy view indices are based on Item Response Theory estimates using the approach developed by Imai et al. (2016). Policy questions are divided into social and economic policy issues using the scheme developed by Caughey and Warshaw (2018). Indices are normalized to have mean zero and unit standard deviation across all respondents across all years, with more positive values indicating more conservative views. All calculations use sample weights.

Figure B4: Changes in college-shares by state over time

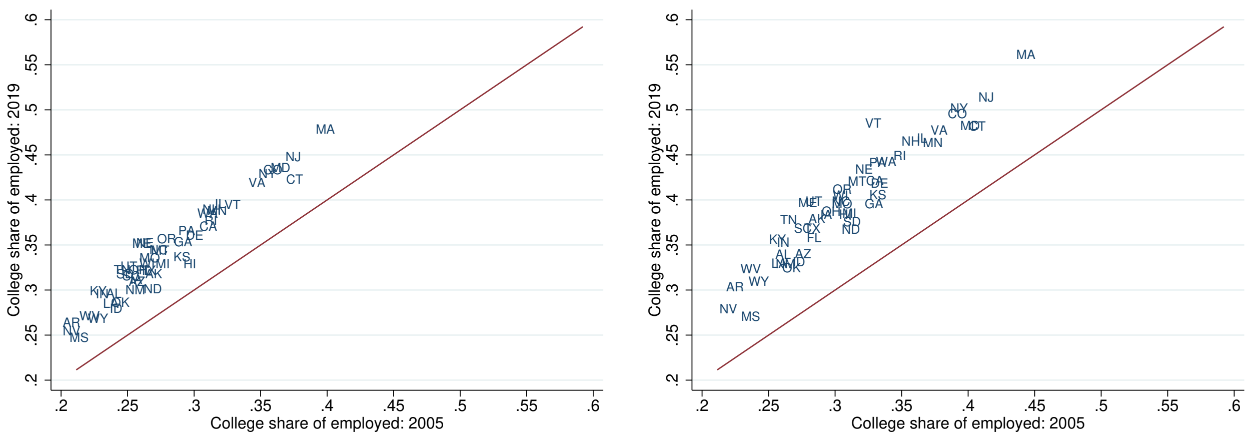

(a) Share of all employed

(b) Share of "young" (26-45) employed

Notes: Calculations based on American Community Survey (ACS). All calculations use sample weights.

Figure B5: College share by state policy over time

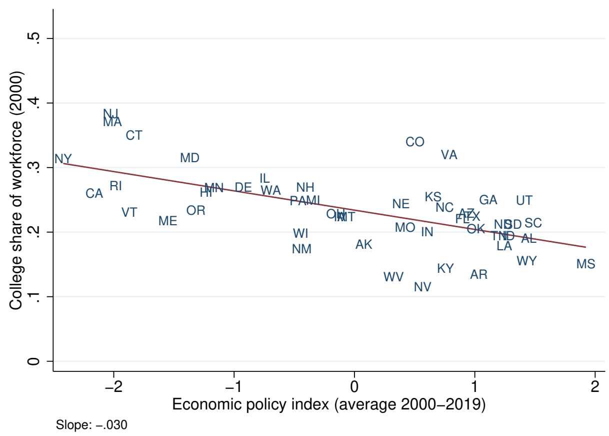

(a) By economic policy, 2000

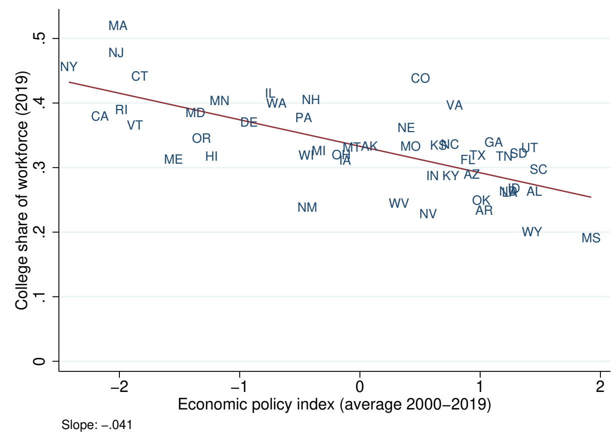

(b) By economic policy, 2019

(c) By social policy, 2000

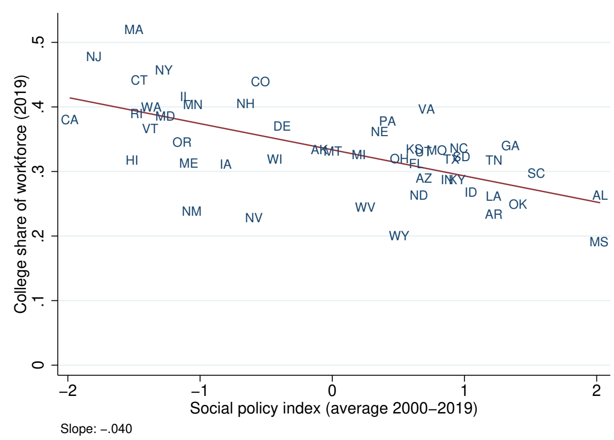

(d) By social policy, 2019

Notes: Calculations based on American Community Survey (ACS). All calculations use sample weights. Policy indices are drawn from Caughey and Warshaw (2016).

Figure B6: Policy effects of Democrat-to-Republican gubernatorial transitions

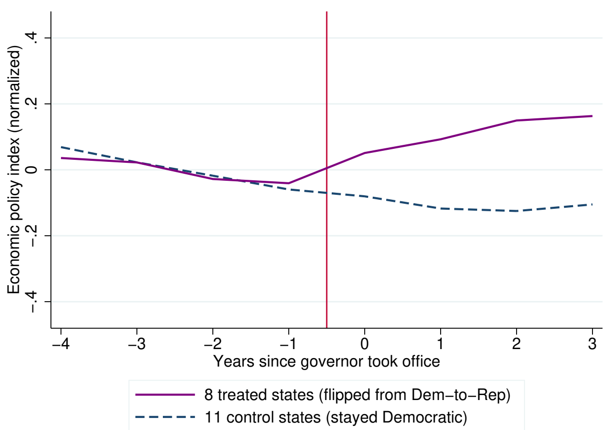

(a) Economic policy

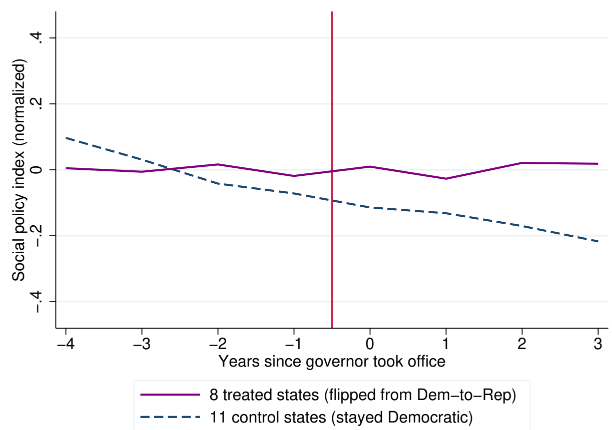

(b) Social policy

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is index of policies in place at the state-year level, taken from Caughey and Warshaw (2016) and normalized to be mean-zero in the pre-period. Higher values indicate more conservative policies. All state-years are re-weighted to ensure that years are proportional to the number of treated states.

Figure B7: Effects on college in-migration by timing of D-to-R flip

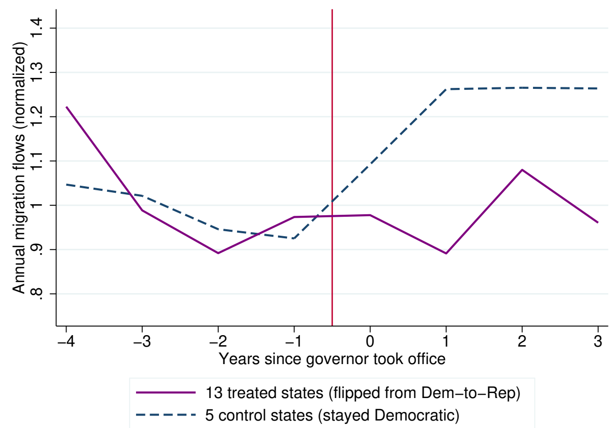

(a) 2003-2005

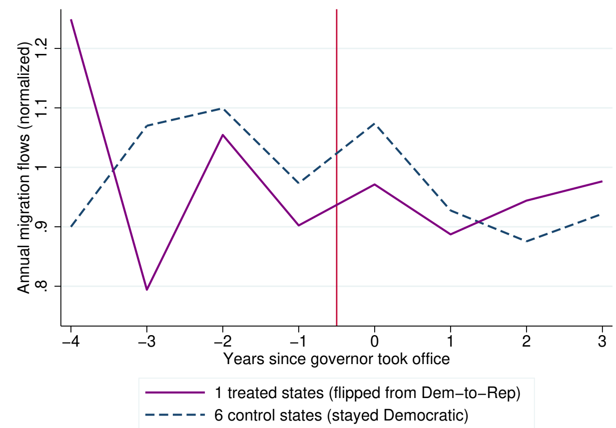

(b) 2006-2008

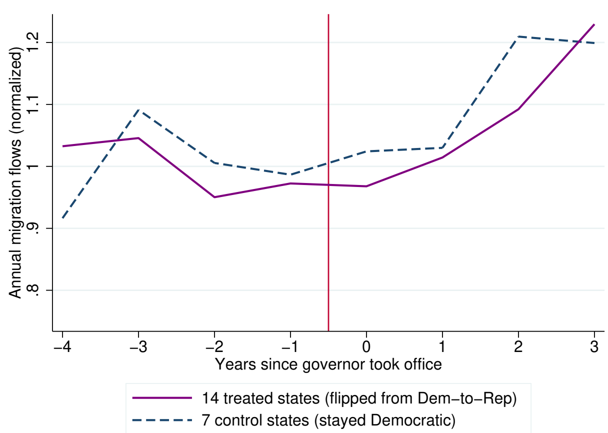

(c) 2009-2011

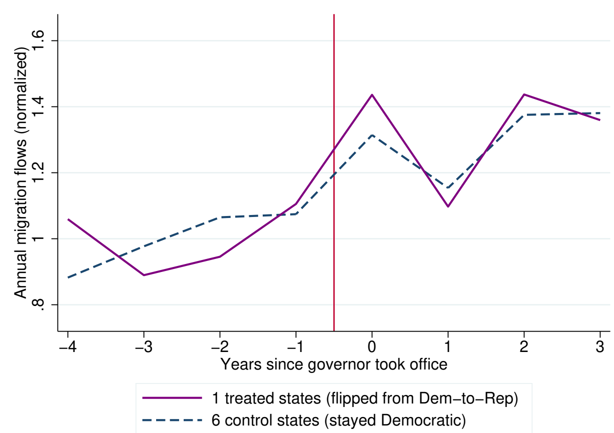

(d) 2012-2014

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is never used.

Figure B8: Other margins of migration adjustment

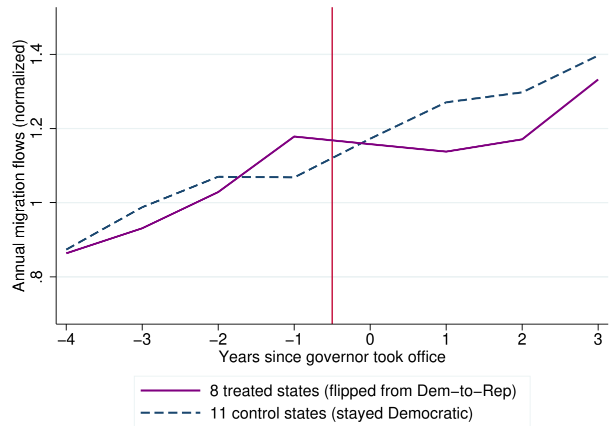

(a) D-to-R, college, out-migration

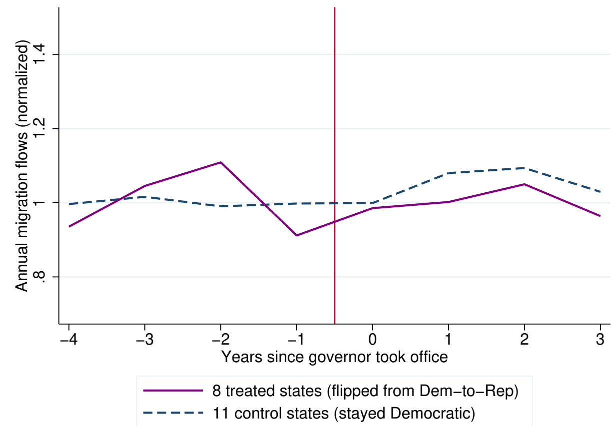

(b) D-to-R, non-college, in-migration

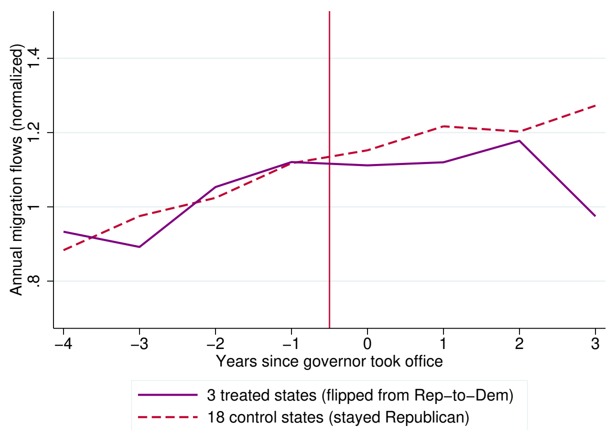

(c) R-to-D, college, in-migration

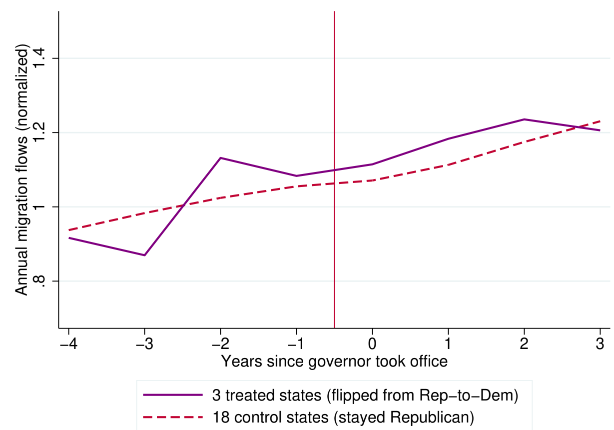

(d) R-to-D, college, out-migration

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states are matched only with "never treated" states (which had the same pre-flip party of governor as the treated states and kept that party for the previous and subsequent 5 years). Dependent variable is normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is never used.

Figure B9: Migration responses for different samples of college-educated workers

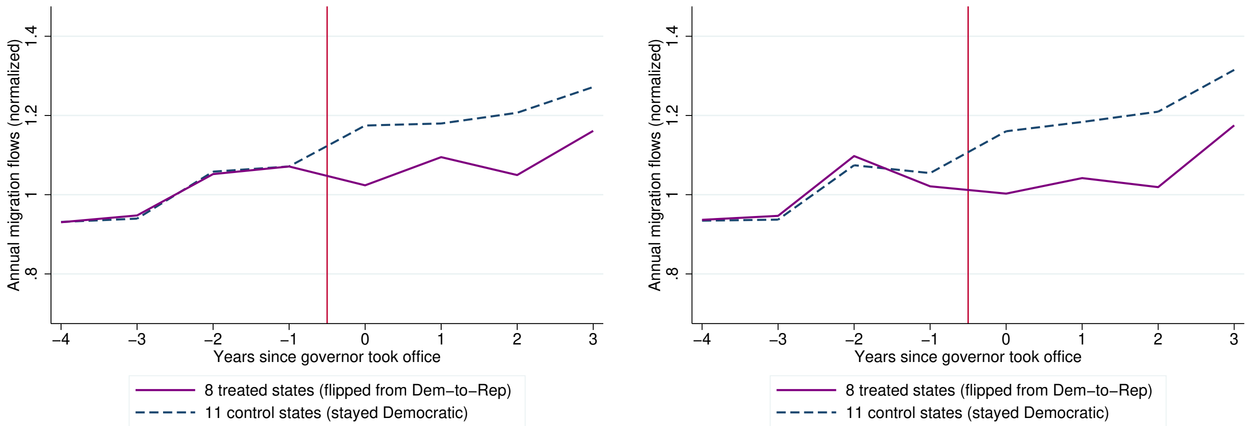

(a) All employed, age 18+

(b) Private sector employed, 26+

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is never used.

Figure B10: Democrat-to-Republican gubernatorial transitions (all years)

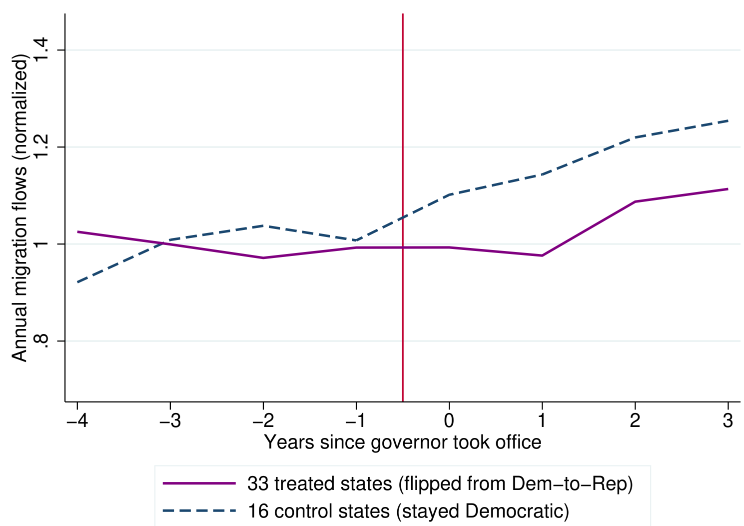

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is never used.

Figure B11: Model-implied and empirical heterogeneity in D-to-R effects

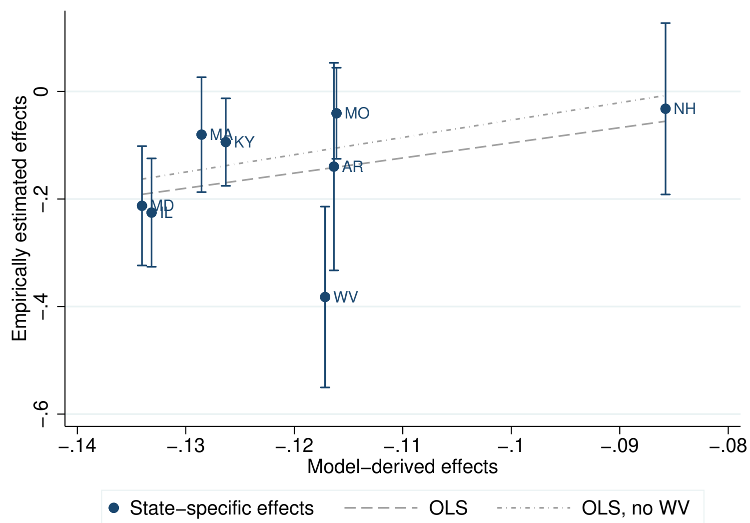

Notes: Empirical estimates (y-axis) are based on a Callaway-Sant'Anna-type estimator in which each "treated" state in individually compared to all "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). These state-specific estimates are generated for each of the eight states flipping from a Democrat to a Republican during 2015-2017 (our main sample). Model-implied effects (x-axis) are derived from our model based on simulating a counterfactual in which those eight states flipped simultaneously, and holding all other states' governors fixed as they were in 2014.

Figure B12: Assessing explanations for education polarization

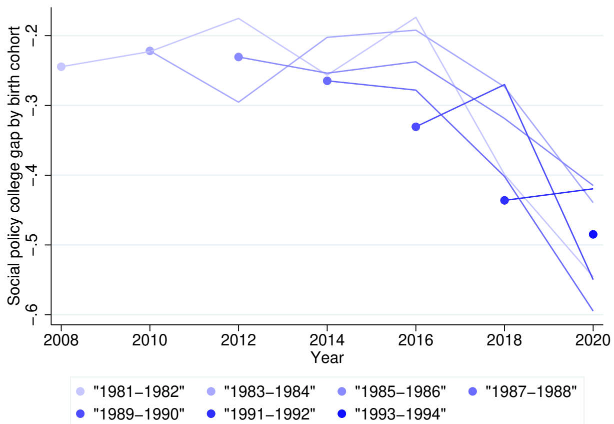

(a) Education polarization by birth cohort

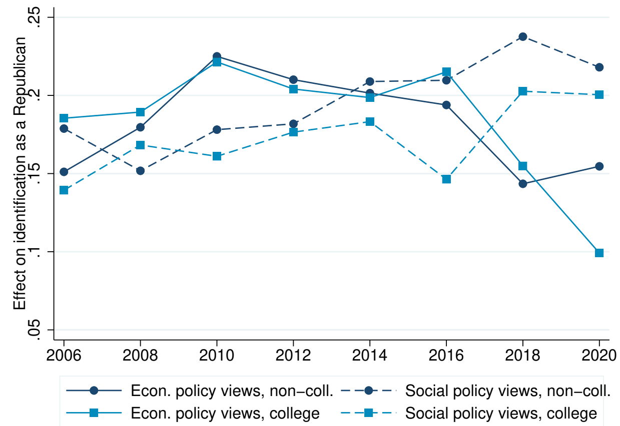

(b) Importance of social/economic policy

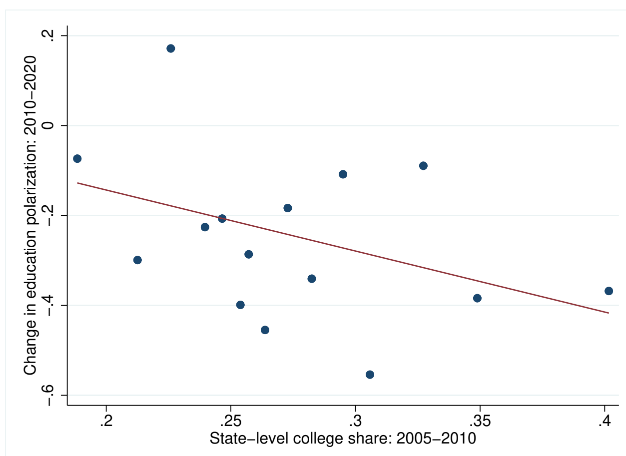

(c) Changes in polarization by state college share

Notes: All calculations based on Cooperative Election Study (CES) respondents aged 26-45. Policy view indices are based on Item Response Theory estimates using the approach developed by Imai et al. (2016). Policy questions are divided into social and economic policy issues using the scheme developed by Caughey and Warshaw (2018). Indices are normalized to have mean zero and unit standard deviation across all respondents across all years, with more positive values indicating more conservative views. All calculations use sample weights.

Figure B13: Earnings and political views by college major

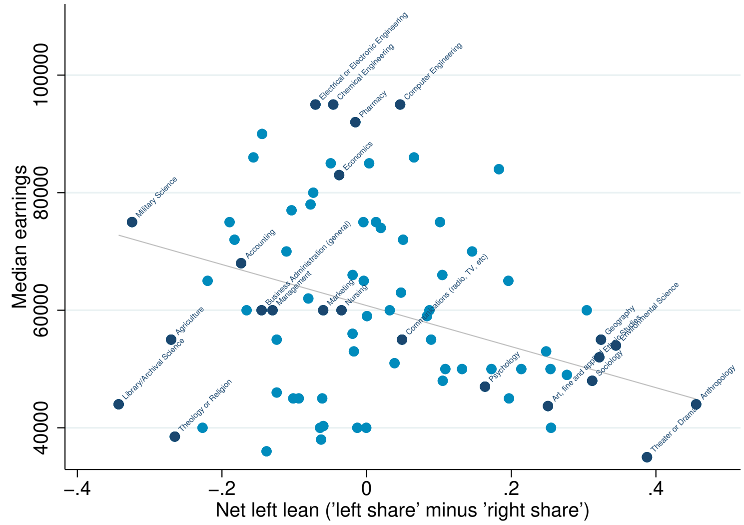

Notes: Median earnings (y-axis) are based on American Community Survey (ACS) respondents from 2009-2019. Net left lean (x-axis) is the difference between the share of Higher Education Research Institute (HERI) respondents who identify with "far left" or "liberal" political views minus the share who identify with "far right" or "conservative" political views.

Figure B14: Migration responses by political lean of major

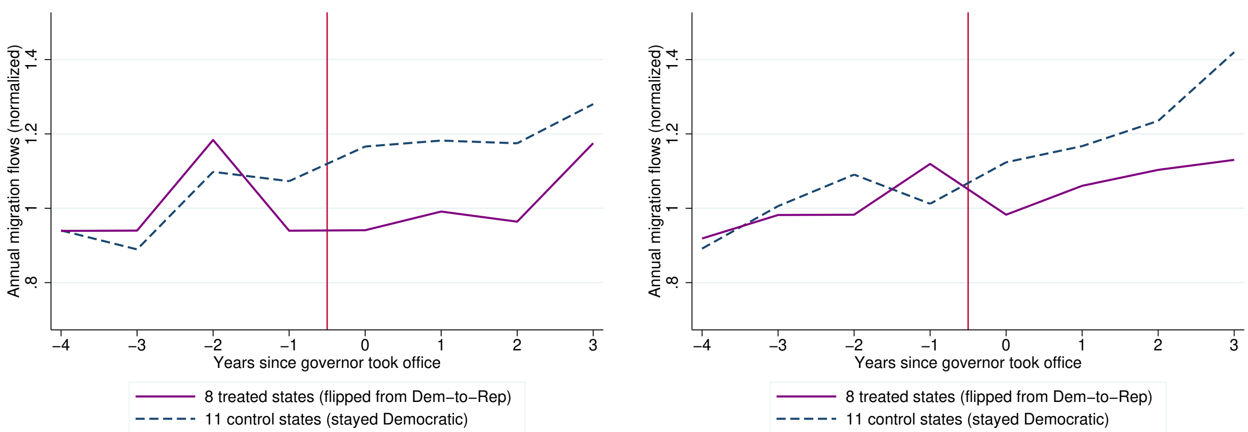

(a) More left-leaning majors

(b) Less left-leaning majors

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. In-migration is calculated separately as a function of field of degree (as recorded in the ACS). Fields are grouped based on the share of HERI respondents who identify with "far left" or "liberal" political views. They are split based on being above/below median, among the sample which migrates across state lines. Migration from 2020 onwards is never used.

Figure B15: IV intuition: National vote shares and gubernatorial election outcomes

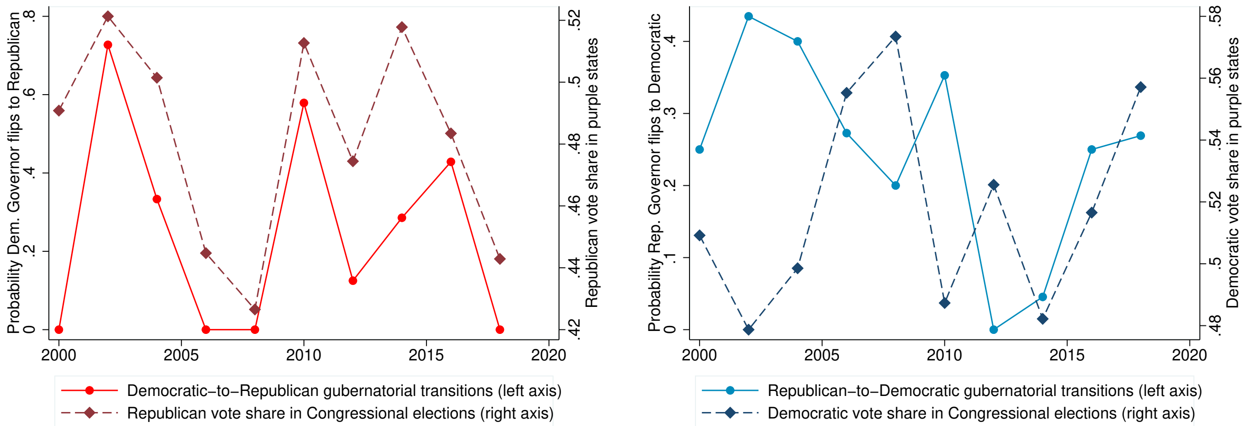

(a) D-to-R transitions

(b) R-to-D transitions

Notes: In Panel (a), the solid line displays the fraction of states i) with a Democratic governor and ii) holding a gubernatorial election, in which the Republican won; dashed line displays the Republican vote share in Congressional elections among the set of states ever observed with both Democratic and Republican governors during our sample period. In Panel (b), the solid line displays the fraction of states i) with a Republican governor and ii) holding a gubernatorial election, in which the Democrat won; dashed line displays the Democratic vote share in Congressional elections among the set of states ever observed with both Democratic and Republican governors during our sample period.

Figure B16: Alternative identification strategies by rolling 10-year window

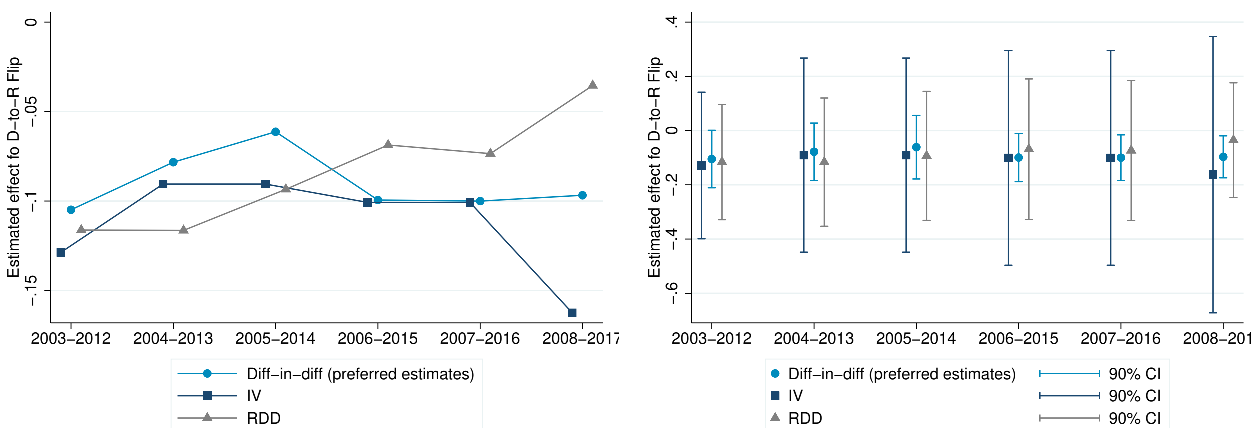

(a) No Confidence Intervals

(b) With Confidence Intervals

Notes: Diff-in-diff estimates are based on our Callaway-Sant'Anna-type estimator. IV estimates are based on the instrument discussed in equation (14) in the text based on changes in Congressional vote share. RDD estimates are based on the Calonico et al. (2014) estimate.

Figure B17: Placebo effects of Democrat-to-Republican transitions on labor demand

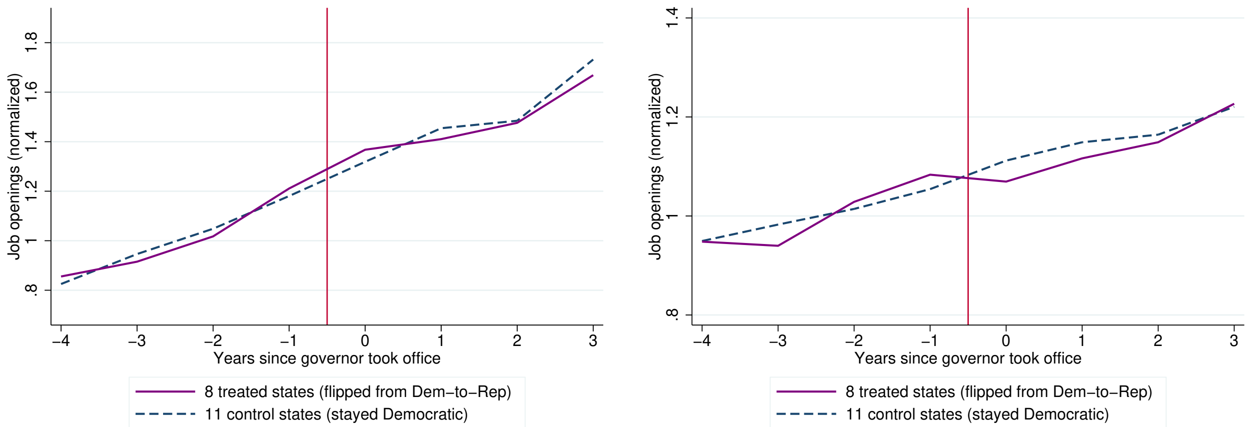

(a) Effects on job openings (b) Effects on hiring

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Democrat to Republican) are matched only with "never treated" states (states which had a Democratic governor for at least the previous and subsequent 5 years). Dependent variables are from JOLTS and are normalized to be mean-one in the pre-period. All state-years are re-weighted to ensure that years are proportional to the number of treated states. Openings and hiring from 2020 onwards are not used.

Figure B18: Placebo effects of 2017 Obama-Trump flips

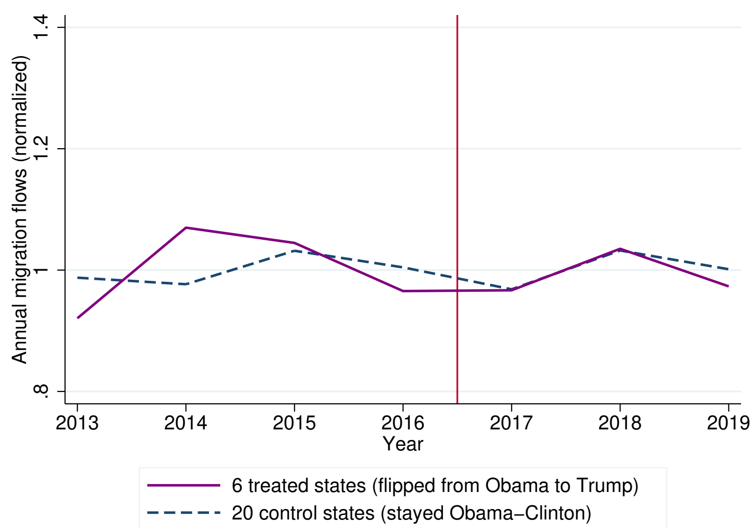

Notes: Estimates based on a Callaway-Sant'Anna-type estimator in which "treated" states (flipping from Obama in 2012 to Trump in 2016) are matched only with "never treated" states (states which voted for Obama-2012 and Clinton-2016). Dependent variable is number of college graduate in-migrants (measured in the ACS), normalized to be mean-one in the pre-period. All state-years are weighted by the number of ACS respondents, in addition to the Callaway-Sant'Anna re-weighting to ensure that years are proportional to the number of treated states. Migration from 2020 onwards is not used.

Table B1: Sources of model-implied heterogeneity

	(1)	(2)	(3)	(4)	(5)	(6)
Inflow rate	1.08***				0.66***	0.70***
	(0.12)				(0.07)	(0.05)
Migration costs		0.28**			-0.73***	-0.39***
		(0.13)			(0.25)	(0.11)
College wages			-0.58***		-0.90***	-0.58***
			(0.10)		(0.24)	(0.10)
Amenities				-0.25	-0.98***	-0.53***
				(0.26)	(0.32)	(0.13)
N	50	50	50	50	50	49
R^2	0.842	0.055	0.243	0.044	0.917	0.941
Alaska	Yes	Yes	Yes	Yes	Yes	No

* $p < .10$, ** $p < .05$, *** $p < .01$. Regressions based on the results from 50 separate counterfactuals, each of which flips only one state independently, relative to the set of governors in office in 2023. For each counterfactual, we calculate the percentage change in the college graduate workforce predicted by our model for the treated state. The absolute value of these changes is regressed on various state characteristics estimated from our model. Column 6 excludes Alaska, a clear outlier in the predicted migration effects as well as for most of the parameters we estimate.

Table B2: Decomposing effects on human capital price inequality

	(1)	(2)	(3)	(4)
Counterfactual	Small blue	Big blue	Small red	Big red
Number of states	5	9	4	8
Panel A: Human capital price, college-educated workers (w_d^C)				
Average effects	-0.012	-0.009	0.013	0.011
(min, max)	(-0.015, -0.011)	(-0.012, -0.007)	(0.012, 0.013)	(0.009, 0.012)
Panel B: Human capital price, non-college workers (w_d^N)				
Average effects	0.010	0.007	-0.009	-0.008
(min, max)	(0.009, 0.011)	(0.005, 0.009)	(-0.010, -0.009)	(-0.009, -0.006)
Panel C: Share attributable to non-college wages				
Share	.450	.437	.421	.423

Table presents changes in equilibrium outcomes for "treated" states (i.e., those changing their governors) in our four counterfactuals (Figure 8 presents the states we simulate as flipping in each counterfactual).

Table B3: Alternative identification strategies

	(1)	(2)	(3)	(4)
Identification:	Diff-in-diff	IV	RDD	RDD
Panel A: All flips (2003-2017)				
D-to-R Flip	-.117**	-.093	-.080	-.089
	(.045)	(.180)	(.098)	(.096)
N	853	213	91	91
R^2	.955	.142		.265
Year FE's	Yes	Yes		Yes
First stage F		46.2		
Estimator		2SLS	CCT-'14	OLS
Bandwidth			.049	.049
Panel B: Recent flips (2008-2017)				
D-to-R Flip	-.097**	-.162	-.035	-.115
	(.047)	(.311)	(.129)	(.109)
N	642	114	57	57
R^2	.957	.236		.258
Year FE's	Yes	Yes		Yes
First stage F		19.7		
Estimator		2SLS	CCT-'14	OLS
Bandwidth			.067	.067

* $p < .10$, ** $p < .05$, *** $p < .01$. Column (1) displays estimates from our preferred Callaway-Sant'Anna-type difference-in-difference estimator. Unit of observation is a state-year. Standard errors are clustered at the state level. Column (2) displays IV estimates using the instrument discussed in equation (14) in the text based on changes in Congressional vote share. Columns (3) and (4) are based on Regression Discontinuity estimates using a triangular kernel, no controls, and Calonico et al. (2014) estimated optimal bandwidth and standard errors (column (3)) and a rectangular kernel, year fixed effects, and the bandwidth chosen in the Calonico et al. (2014) estimate (column (4)). In columns (2)-(4), the dependent variable is the change in the log of average annual in-migration from the four pre-election years to the four post-election years.

C Model

C.1 The probability that an individual chooses one state

$$V_{ido} = \alpha_d^g(1-\gamma)w_d(1-\tau_{do}^g)\epsilon_{do}s_{id} = \tilde{w_{do}}s_{id}$$

Extreme value theory: U(.) is Frechet $\Rightarrow$ so is max u(.), which is V(.)
Without loss of generality, consider the probability that worker chooses destination 1 and denote this by π_{1o}:

$$\begin{aligned}
\pi_{1o} &= Pr\left[\tilde{w_{1o}}s_{i1} \geq \tilde{w_{so}}s_{is}\right]\forall s \neq 1 \\
&= Pr\left[s_{is} \leq \frac{\tilde{w_{1o}}s_{i1}}{\tilde{w_{so}}}\right]\forall s \neq 1 \\
&= \int F_1(s_i, T_2 s_i, ..., T_N s_i)ds_i
\end{aligned}$$

where $F_1(.)$ is the derivative of cdf with respect to its first argument and $T_l \equiv \frac{\tilde{w_{1o}}}{\tilde{w_{lo}}}$. Recall that

$$F(s_1, s_2, ..., s_N) = exp\left(-\left[\sum_{d=1}^{N} s_d^{-\theta}\right]\right)$$

Taking the derivative with respect to s_1 gives

$$F_1(s_i, T_2 s_i, ..., T_N s_i) = \theta s_i^{-\theta-1} exp(-\bar{T}_r s_i^{-\theta})$$

where $\bar{T}_r \equiv -\sum_{s=1}^{N}(\frac{\tilde{w_{ro}}}{\tilde{w_{so}}})^{-\theta}$.

$$\begin{aligned}
\pi_{1o} &= \int F_1(s_i, T_2 s_i, ..., T_N s_i)ds_i \\
&= \frac{1}{\bar{T}_r}\int \bar{T}_{ri}^{-\theta-1} exp(-\bar{T}_r s_i^{-\theta})dFs_i \\
&= \frac{1}{\bar{T}_r}\int dFs_i \\
&= \frac{1}{\bar{T}_r} \\
&= \frac{\tilde{w_{ro}}^{\theta}}{\sum_{s=1}^{N}\tilde{w_{so}}^{\theta}}
\end{aligned}$$

C.2 Average skill of workers

To calculate this conditional expectation, we use the extreme value properties of the Frechet distribution. Let $y_d = \tilde{w}_d s_d$ denote the key destination choice term.

$$y^* \equiv max_s \{y_s\} = max_s \{s_s/\zeta_s\} = s^*/\zeta^*.$$

$$\begin{aligned} Pr[y^* < z] &= \prod_{s=1}^{N} Pr[y_s < z] \\ &= \prod_{s=1}^{N} Pr[\tilde{w}_s s_s < z] \\ &= \prod_{s=1}^{N} Pr[s_s < z/\tilde{w}_s] \\ &= \prod_{s=1}^{N} exp\left[-(z/\tilde{w}_s)^{-\theta}\right] \\ &= exp\left\{-\sum_{s=1}^{N}(z/\tilde{w}_s)^{-\theta}\right\} \\ &= exp\left\{-\bar{T}z^{-\theta}\right\} \end{aligned}$$

where $\bar{T} \equiv \sum_{s=1}^{N} \tilde{w_s}^{\theta}$

The ability of people in their chosen place is also Frechet distribution:

$$F(x) \equiv Pr[s^* < x] = exp\left\{-T^* x^{-\theta}\right\} \tag{17}$$

where $T^* \equiv \sum_{s=1}^{N}(\frac{\tilde{w_s}}{\tilde{w^*}})^{\theta} = \frac{1}{\pi}$.

The expected skill of chosen d is

$$\begin{aligned} E(s_d) &= \int_0^{\infty} s dF(s) \\ &= \int_0^{\infty} \theta T^* s^{(1-\theta)} e^{-T^* s^{-\theta}} ds \end{aligned} \tag{18}$$

Recall the Gamma function is $\Gamma(\alpha) \equiv \int_0^{\infty} x^{\alpha-1}e^{-x}dx$. We replace x $= T^* s^{-\theta}$, $dx = \theta T^* s^{(-\theta-1)} ds$, $s = (\frac{T^*}{x})^{\frac{1}{\theta}}$ We can show that

$$\begin{aligned} E(s_d) =& (T^*)^{\frac{1}{\theta}} \int_0^{\infty} x^{-\frac{1}{\theta}} e^{-x} dx \\ &= (T^*)^{\frac{1}{\theta}} \Gamma(\frac{\theta-1}{\theta}) \end{aligned} \tag{19}$$

Applying this result to our equation, we have

$$
\begin{aligned}
E\left(s_d^{\theta} \mid choose\; d\; from\; o\right) &= (T^{*})^{\frac{1}{\theta}}\Gamma(\frac{\theta-1}{\theta}) \\
&= \pi_{do}^{-\frac{1}{\theta}}\Gamma(\frac{\theta-1}{\theta})
\end{aligned}
\tag{20}
$$