Partisan Effects of Religious Peer Groups*

Kurtis Gilliat
University of California, San Diego

Jacob R. Brown
Boston University

September 24, 2026

Click here for the most recent version

Abstract

How people shape each other's political affiliation is a central question in the social sciences – with broader implications for partisan sorting and polarization – but is difficult to answer because people select into peer groups. We study a natural experiment in which the Church of Jesus Christ of Latter-day Saints periodically redraws congregation assignment boundaries, reassigning members to new church peers. Linking records on every Utah registered voter from 2012-2021 to digitized congregation maps, we employ a spatial difference-in-discontinuities design that compares changes in party registration among voters who live near one another and start off in the same congregation but are assigned to different congregations. Voters conform to their new peers: Unaffiliated voters assigned to more Republican congregations become more likely to register Republican, and Republicans assigned to more Democratic congregations become more likely to leave the party. Reassigned boundaries show no relationship with prior registration changes, supporting the design's identifying assumption. The effects are not explained by correlated peer characteristics such as race or by the partisanship of congregation leaders. Effects are strongest between demographically similar peers, the members most likely to become friends, suggesting that influence operates through personal relationships among co-congregants.

*We thank Gordon Dahl, Samuel Bazzi, Eli Berman, Itzik Fadlon, Laurence Iannaccone, Saul Martinez, Michael McBride, Jared Rubin, and Daniel Hungerman for helpful feedback. We also thank seminar participants at UC San Diego, and conference participants at IRES, ASREC, SOCAE, ACLEC, APSA, and Brigham Young University. Additionally, we thank participants of the UC San Diego Undergraduate-Graduate Economic Research Lab for excellent research assistance, Jonathan Palmer for assistance in processing geographic data, and TargetSmart for providing voter file data.

1 Introduction

As the American electorate increasingly sorts by party affiliation across residential, professional, and social contexts, many worry that the loss of contact with diverse political perspectives is eroding Americans' capacity for collaborating with and understanding the other side (Bishop, 2008; Brown and Enos, 2021; Brown et al., 2025; Chinoy and Koenen, 2024; Frake, Hurst and Kagan, 2026; Huber and Malhotra, 2017; Iyengar et al., 2019; Kaplan, Spenkuch and Sullivan, 2022; Mutz, 2006). However, this concern relies on an assumption that individuals are influenced by the politics of their peers, whether by being persuaded by their views, by adopting their preferences, or at the very least by developing empathy for peers of the opposite party. If, on the other hand, individuals are unaffected (Campbell et al., 1960; Green, Palmquist and Schickler, 2002) or even harden against opposing viewpoints from peers (Bail et al., 2018), then political segregation will have muted effects on political behavior and partisan conflict. The question of how voters are affected by the politics of their peers is thus critical to understanding contemporary political polarization, yet persistently difficult to answer, because individuals largely choose where they live, work, and form social ties.

In this article, we study the effects of peers on political party affiliation by exploiting a natural experiment arising from the geographic assignment, and periodic reassignment, of members of The Church of Jesus Christ of Latter-day Saints (the LDS Church) to local congregations in Utah. LDS members are administratively assigned to small, geographically defined congregations called wards, whose boundaries are intermittently redrawn by regional church leaders responding to local changes in church attendance. Interviews with several of these leaders, together with official church guidelines, indicate that they redraw lines to keep congregations within a workable size for an all-volunteer lay ministry, to balance the number of youth and children across wards, and to follow natural geographic divides. They do so using a GIS application that reports membership counts by age and sex for each proposed configuration, and that reports nothing about members' politics. These redistricting events plausibly exogenously assign members to new congregations, and thus to new peer groups, while leaving their homes, neighborhoods, and jobs unchanged. Using records for every registered voter in Utah, we ask whether individuals reassigned to a more Republican or more Democratic congregation change their own party registration in the direction of their new peers.

This setting offers a rare combination of features that allow us to both identify and measure the effect of peer-group exposure on party affiliation. Isolating the causal effect of peers is notoriously difficult. The ideal experiment would randomly assign voters to peer groups with different partisan compositions, while holding all other variables fixed. Furthermore, the assigned peer groups must actually be relevant to a person's life. They need to be a group of people who individuals will treat as real peers, otherwise peer effects will be null by construction. It is unlikely that such relationships could be generated in a lab or field experiment ex nihilo. Ward redistricting addresses these concerns directly. Boundary changes reassign peers while holding residential location fixed, separating congregation exposure from the neighborhood, employment, and other contextual changes that typically accompany a change in peers, and doing so at a fine geographic scale. LDS wards are also an important peer group, as members spend substantial time together each week through worship services, lay-ministry assignments, and social activities. Beyond identification, the geographic nature of ward assignments also provides us with a way to measure peer groups at a large scale. By linking address-level voter records to digitized ward-boundary maps, we are able to identify the ward that each Utah voter is assigned to.

Two features of the data shape our empirical strategy. First, the voter file does not report religion, so we cannot observe which residents are LDS and therefore subject to ward assignment. We address this by studying Utah, where LDS membership is widespread—about 69% of residents in 2010 and 61% in 2020 (Mason, Toney and Cho, 2011)—and by treating every Utah voter, LDS or not, as "assigned" to the ward containing their address. Our estimates are thus intent-to-treat effects of ward partisan composition on the full population of registered voters rather than on LDS members alone. Including non-members effects our estimates in two ways: non-members do not respond to ward assignment, which dilutes the estimated effect, and measuring peer composition over all voters rather than members alone mismeasures the partisan contrast members actually experience. We derive the combined bias in closed form and show that, for the ward configurations in our data, it attenuates our estimates, which therefore understate the effect of peers on LDS members. Second, we exploit the panel structure and address-level detail of the voter file to implement a spatial difference-in-discontinuities design, comparing changes in party registration between the 2012 and 2020 elections for voters living just on either side of a redrawn boundary. Ward assignment is measured at each voter's 2012 address, so voters who move after a boundary change keep their original assignment, and a flexible function of distance to the boundary absorbs the smooth partisan sorting across space that would otherwise confound a simple comparison of congregations.

Our central finding is that voters conform to the partisan leanings of their new peers. For voters registered as Republican in 2012, the largest effects come from Democratic exposure: assignment to the more Democratic side of a redrawn boundary lowers the probability of remaining Republican by about 0.8 percentage points on average, with a corresponding 0.4

percentage point rise in the probability of becoming Democrat and 0.4 percentage point rise in becoming Unaffiliated. For 2012 Unaffiliated voters,[1] the largest effects come from Republican exposure: assignment to the more Republican side raises the probability of registering Republican by about 1.4 percentage points, with corresponding reductions in the probability of staying Unaffiliated or becoming Democrat. For context, between 2012 and 2020, 12.7% of Utah Republicans left their party and 27.9% of Unaffiliated voters joined it, so assignment across a typical boundary accounts for roughly 5% of each flow. We find no significant effects for baseline Democrats, who constitute a much smaller share of the Utah electorate, and a smaller share still of the Utah LDS population.

We support the validity of this design with two complementary exercises. First, a placebo analysis using boundary changes that occurred after our 2012–2020 outcome period yields consistently null estimates, indicating that our results are not driven by differential trends correlated with where new boundaries are drawn. Second, a simulation analysis comparing observed ward configurations to thousands of feasible alternatives shows that observed wards are, if anything, slightly more balanced in partisan composition than the sampling distribution, so there is no sign that wards are drawn to segregate partisans. The modest balancing we do find is in levels of partisanship, which our first-differenced design absorbs, and to the extent it extends to partisan trends it would tend to attenuate our estimates. Finally, congregation boundary lines are drawn along sensible geographic divides such as major roads, neighborhood edges, and school boundaries. These divides may coincide with differences in housing stock, on which Republicans and Democrats may sort differently, and with residential sorting into school zones. We therefore show that our estimates are unchanged when we allow trends to differ by demographics, by property characteristics from linked CoreLogic records, and by school catchment.

In principle, the effect of assignment to a more Republican or more Democratic ward could reflect not only peers' partisanship but also peer attributes correlated with it, or the partisanship of the ward's leader. We test for these alternative channels and find no evidence that they drive our results. Wards that differ in partisanship also differ in racial composition, turnout, homeownership, and demographics, but adding each of these peer characteristics to the regression alongside peer partisanship leaves the partisan coefficient essentially unchanged. What matters about a voter's new peers appears to be their politics, not the other characteristics that come with them.

We next ask whether the influence comes from fellow congregants or from the congrega-

[1] We group Utah voters into three registration categories: Republicans (45%), Democrats (9%), and Unaffiliated (46%), with the latter pooling all third-party and unaffiliated registrants. Percentages indicate the percent of all 2012 registered Utah voters belonging to each category.

tion's leader, and find little role for the latter. Each ward is led by a bishop, a lay member drawn from the ward's own membership who serves for about five years. The bishop sets the program for Sunday services, including who gives sermons, meets individually with members about their spiritual and financial needs, and works closely with the ward's youth. To test whether bishops move members' politics, we assemble a new dataset recording the bishop of nearly every Utah ward each quarter since 2023, link each bishop to his own voter record, and exploit the routine rotation of bishops in a difference-in-differences design. Acquiring a Republican rather than a non-Republican bishop has no detectable effect on members' registration, and the confidence intervals are tight. The bishop channel can contribute to the peer effect only if a more Republican ward is also more likely to have a Republican bishop. Even if that relationship were one-for-one, so that a ward's Republican share translated fully into the probability of a Republican bishop, the largest bishop effect our confidence intervals allow would account for at most 3.1% of the peer effect we estimate among Unaffiliated voters. Such a small bishop effect is consistent with survey evidence that LDS members hear political messaging from the pulpit far less often than members of any other major U.S. religious tradition.[2] We stress that this is not a general test of whether leaders matter less than peers. LDS church leaders are expected by norm to keep their politics to themselves, so the null likely reflects a setting in which leaders are unusually apolitical.

Several patterns point to influence through personal relationships among members. Friendships form disproportionately among people who share demographic characteristics (Currarini, Jackson and Pin, 2009; McPherson, Smith-Lovin and Cook, 2001), and within wards this is reinforced by the Church's separate men's and women's organizations. When we decompose peer exposure into same- and opposite-demographic components, the partisan effect loads on same-demographic peers, most sharply by gender. Influence broadcast from the pulpit would reach men and women alike, so a gender gradient points to relationships rather than to the shared Sunday experience. If influence travels through conversation, it should also be strongest among voters who discuss politics and who live close enough to their ward-mates for frequent contact, and it is: effects are larger among voters who turned out in the 2010 election ($p = 0.005$) and among non-rural residents ($p = 0.040$).

Peers appear to change not only the party a voter registers with but also the partisan identity behind it. Utah's Republican primary is closed and usually decides the election, which gives non-Republicans a strategic reason to register Republican. Part of the movement we estimate is consistent with such a strategic motive: Unaffiliated voters assigned

[2]In the Faith Matters Survey (Putnam and Campbell, 2010), 2% of LDS respondents report that political issues are discussed from the pulpit at least monthly, the lowest share of any tradition surveyed (Appendix A.2).

to more Republican wards become more likely both to register Republican and to vote in the Republican primary. A Republican, by contrast, gains no primary access by leaving the party, yet Republicans exposed to more Democratic peers become more likely to defect, and they do so even when they cast no primary ballot at all. At least part of our effect therefore reflects genuine change in ideology or identity.

Together, these results contribute to three literatures. The first is the study of peer effects on political behavior and, through it, the debate over partisan sorting. The cleanest existing evidence assigns peers directly, to college roommates (Strother et al., 2021) and classmates (Campos, Hargreaves Heap and Leite Lopez De Leon, 2016), but relies on institutions populated by young adults still forming their political identities. In contrast, studies of neighborhoods, states, and workplaces (Brown, 2025; Brown et al., 2023; Cantoni and Pons, 2022; Chinoy and Koenen, 2024; Chyn and Haggag, 2023; Hurst et al., 2025; Perez-Truglia and Cruces, 2017) reach a broad adult population, but cannot isolate the effects of peers from other place-based attributes. Ward boundary changes combine the strengths of both by cleanly isolating the effects of peers in a broad population of adults.[3] We find effects across the full age distribution, which matters for sorting: people move throughout the life cycle, so peer influence confined to early adulthood could do little to amplify the effects of sorting. We also find that voters conform to peers across the partisan divide, and that exposure to the out-party moves them more than reinforcement from co-partisans does. Sorting therefore does not merely reflect polarization but feeds it. As social contexts homogenize, voters lose the cross-party contact that would otherwise pull them toward the other side.

Our second contribution is to the study of how political preferences are transmitted between people. Models of cultural transmission distinguish socialization within the family, by peers, and by institutions such as schools, churches, and media (Bisin and Verdier, 2001, 2011). Causal evidence on political preferences has concentrated on the first and third of these, particularly on media (DellaVigna and Kaplan, 2007; Enikolopov, Petrova and Zhuravskaya, 2011), while transmission between adult peers has been harder to isolate. A congregation combines an institution, with a leader who speaks for it, and a set of peers, and our setting separates the two. The rotation of bishops varies a ward's leader while holding its members fixed, and the leader's partisanship has no detectable effect. Because LDS leaders are encouraged to keep their politics private, we do not take this as evidence that leaders never matter, but where the institution strives to be nonpartisan, what a congregation transmits is transmitted between members.

[3]Relative to U.S. adults, Utah Latter-day Saints are younger, whiter, more Republican, more religiously observant, and more likely to be married with children, and are similar on education and income (Appendix A.4).

The third contribution is to the study of religious congregations as political environments. Congregations are a setting of sustained interaction for many Americans (Djupe and Gilbert, 2006, 2009; Lim and Putnam, 2010; Putnam and Campbell, 2010; Wald, Owen and Hill, 1988), and one that is increasingly politicized as political orientation shapes religious involvement and vice versa (Margolis, 2018). Previous studies have shown that churches play important roles as political organizers (Campbell, 2004; Verba, Schlozman and Brady, 1995), that declines in church attendance can demobilize voters (Gerber, Gruber and Hungerman, 2016; Gruber and Hungerman, 2008), and that voters may choose their level of religious involvement and their congregation based on the perceived politics of the pastor or congregation (Audette and Weaver, 2016; Hafner and Audette, 2023; Margolis, 2018). Few studies, however, have tested whether the people voters attend church with influence them (see the review by Ladam, Shapiro and Sokhey, 2019). We show that they do, and that the influence comes from fellow congregants rather than from the pulpit. Congregations can therefore be settings of political persuasion, where individuals conform toward, rather than react against, co-congregants who differ from them, even where clergy stay out of politics. This gives the growing alignment between religion and partisanship a mechanism that does not run through politicized leaders: a congregation can pull its members' politics toward the local majority without anyone preaching it.

The rest of the paper proceeds as follows. Section 2 describes the organization of LDS congregations in Utah, the process by which ward boundaries are redrawn, and the Utah political institutions that shape what party registration measures. Section 3 describes our data. Section 4 explains our empirical design, estimating equation, and the interpretation of the linear-in-means estimand. Section 5 presents the analyses supporting the identification assumptions of our difference-in-discontinuities design. Section 6 presents our core results on partisan registration, together with their magnitude, robustness, and sensitivity to selective attrition. Section 7 analyzes mechanisms, asking whether effects are driven by peers or by church services, which voters respond, and what an induced change in registration represents. Section 8 concludes.

2 LDS Congregations and Utah Politics

Our natural experiment is built on the way the Church of Jesus Christ of Latter-day Saints (LDS Church) organizes its members into local congregations called wards.[4] The interpretation of our principal outcome, party registration, depends in addition on features of Utah's

[4]Despite the shared name, LDS wards are purely ecclesiastical units with no connection to the administrative wards used in some U.S. municipalities.

political landscape, including its electoral rules. This section describes each in turn.

Wards are small, geographically defined congregations, each typically comprising several hundred individuals who attend weekly worship services and participate in social and religious activities together. Three institutional features make wards a useful laboratory for studying peer influence: (1) members spend substantial time together, (2) ward assignment is geographically determined with high compliance, and (3) ward boundaries are periodically redrawn in ways that are plausibly exogenous to political outcomes.

2.1 Wards are a Consequential Peer Group

Members of the LDS Church have one of the highest rates of weekly church attendance among major U.S. religious groups, with 77% of LDS respondents in the 2014 Pew Religious Landscape Study reporting weekly or more frequent attendance (Pew Research Center, 2015). In addition, many members hold volunteer lay-ministry assignments, known as "callings" (e.g., Sunday School teacher, choir director, even bishop[5]—the ward's presiding lay minister), which often involve serving on committees with other members. Cnaan, Evans and Curtis (2012) estimate that the average LDS member spends about 8 hours per week on church-related volunteering, and the Church's own materials describe lay-ministry assignments as requiring anywhere from roughly 5 to 30 hours per week (The Church of Jesus Christ of Latter-day Saints, 2026*b*). Beyond weekly meetings and church assignments, many members also participate in social activities including service projects, potlucks, sports leagues, choir, youth activities, and summer camps.

This intensity of contact shows up in members' reported social networks. In Robert Putnam and David Campbell's Faith Matters Survey—a nationally representative survey of U.S. adults with a uniquely detailed battery on respondents' religious social ties (Putnam and Campbell, 2010)—LDS attenders are more likely than attenders of any other major U.S. religious tradition to report six or more close friends in their congregation. Among LDS respondents, congregation friends also rank above workplace friends on every tie-strength measure the survey asks about (Appendix A.1). A member's particular mix of ward members therefore strongly shapes their set of friends, day-to-day social contacts, and broader support network. Appendix A.1 presents additional survey evidence on each of these dimensions of ward life.

[5]An LDS bishop leads a single ward, making the role closer to a parish pastor than to a bishop in Catholic or episcopal traditions. The nearer analogue of a diocesan bishop is the stake president introduced below. That office, too, is an unpaid lay position.

2.2 Ward Assignment is Geographically Determined

Each member is assigned to the ward whose boundaries contain their home address. Thus, short of moving, members of the LDS Church, unlike congregants in most American denominations, have little ability to sort into like-minded congregations. In Utah, where we focus our attention, wards are geographically small (the median ward is 0.5 square kilometers in area), meaning that members of neighboring congregations typically face similar local environments, which is exactly the comparison our empirical design exploits.

Compliance with ward assignment is high for two reasons. First, full participation in a ward—including the ability to serve in a lay ministry role and to participate in key religious rites (e.g., baptism, marriage)—is contingent on having one's official membership records assigned to that ward. Second, moving membership records to a ward outside one's residential area requires approval from the Church's highest governing body in Salt Lake City and is rarely granted. This policy reflects the Church's belief that geographically defined wards encourage bridging social ties. As official Church materials put it: "Wards are chosen for us based on a reasonable geographic alignment, and we learn to live with, serve, and love people who might well be different in background, preferences, and opinions" (The Church of Jesus Christ of Latter-day Saints, 2023).

2.3 Ward Boundaries are Occasionally Changed

Ward boundaries are occasionally redrawn, which we exploit as a plausibly exogenous shock to a person's peer group. Any adjustment to ward boundary lines, including the creation of new wards and the closing of old wards, is initiated by the stake president, the lay leader of a group of 5–10 contiguous wards called a stake[6] (see Figure 1a for an example of one stake). To learn how these decisions are made in practice, we surveyed former Utah stake presidents who served during our 2013–2017 study window, contacting 492 by email and receiving substantive responses, in writing or by phone, from 19 who together led stakes in seven Utah counties.[7] What follows draws on those responses and on the Church's published guidelines, which agree with one another.

Boundary changes are, first and foremost, a response to changes in membership. Every respondent who had overseen a change described it as driven by growth, decline, or an imbalance across neighboring wards, and the reason ward size matters is that wards are run

[6] Changes that move wards between stakes are coordinated across the affected stake presidencies by a regional Church officer. In our data, 17% of boundary-change clusters cross a stake line (Appendix B.5).

[7] Appendix B describes the sample frame and response, reproduces the recruitment email, and summarizes the responses by topic.

entirely by unpaid lay members.[8] In a ward that is too small, the same few active members must staff every position, while in one that is too large, the bishop is overwhelmed and new members are easily lost. In declining areas, some wards may need to be closed and absorbed into neighbors; in growing areas, new wards are created by splitting existing ones; other times, attendance can be rebalanced by simply shifting a boundary line between neighboring wards.

In addition to rebalancing overall attendance, stake presidents in our survey described taking three other factors into account. First, they try to ensure that each ward has enough children under 12, teenage boys, and teenage girls so that members of these groups can make friends. Second, they prefer balance across wards to similarity within them: asked whether they would divide an area containing a wealthier and a poorer neighborhood so that each new ward drew from both, or so that each drew mostly from one, respondents overwhelmingly chose the former. Third, they try to follow sensible geographic divides (major roads, rivers, neighborhood edges) when drawing lines, and where a neighborhood must be divided, to run the line along the backs of lots rather than down the middle of a street. We return to the potential confounds introduced by these secondary considerations in Section 4 when we discuss the controls and fixed effects used to address them.

Beyond these factors, stake presidents do not account for individual characteristics, both because most are unobservable to them and because anchoring boundaries to the particular individuals living there at the time would be shortsighted given normal residential turnover. The one exception that came up in several conversations concerns where the sitting bishop of each existing ward lives. When a ward is split, some stake presidents consider which side of the new line the bishop falls on, to avoid disrupting leadership continuity, while others deliberately ignore it on the grounds that a bishop's tenure is short relative to a boundary's. Even this consideration does not appear to operate on a broad scale. If boundaries were drawn to keep sitting bishops on a particular side, bishops would be disproportionately clustered near new lines, and Appendix Figure B1 shows they are not. Politics, in particular, never entered the conversations we had: respondents uniformly reported that they neither observed nor considered members' partisanship, and that it never arose in deliberations.

Mechanically, a boundary change proceeds as follows. The stake president, sometimes in consultation with each ward's bishop, draws up a proposed set of new lines using an online mapping tool provided by the Church that displays the membership counts above for each

[8]The Church discourages wards that are too small or too large, preferring wards of several hundred members; respondents described a target range of roughly 300–500. In creating a new ward, the Church requires a minimum of 250 members on the roll, with at least 100 actively participating adults, and at least 20 members eligible to serve in key leadership positions (essentially, actively attending men who pay a full 10% tithe) (The Church of Jesus Christ of Latter-day Saints, 2026*a*).

ward under the proposed configuration, and nothing about members' politics. Respondents described iterating on the map repeatedly, one recalling seven or eight configurations before finding one that satisfied his criteria for every ward at once. Once satisfied, the stake president submits the proposal to Church headquarters, then waits, sometimes months, for approval, which is not automatic. Only if and after the proposal is approved is anything announced to ward members, who are expected to comply with their new assignments the following week. Compliance is high: most respondents could recall only one or two households in their entire stake attending somewhere other than their assigned ward, each under unusual circumstances and with headquarters' approval.

Between 2013 and 2017, the years for which we have ward boundary data, the number of wards in our data rose from 3,975 to 4,115, as 239 new wards were created and 99 old wards were dissolved. Often, boundaries for multiple wards were adjusted simultaneously since any change to one ward's boundary tended to set off compensating adjustments in neighboring wards. Figure 1a shows an example, using a stake which grew from eleven to twelve wards between 2014 and 2015. In order to accommodate growth, this stake reorganized four wards into five, while leaving the other seven wards in the stake unaffected. Altogether, of the 3,975 wards we observe in 2013, 1,334 (33.6%) experienced some form of boundary adjustment by 2017. At this rate, the average ward is redrawn about once every 11.9 years.

Consistent with attendance rebalancing, new wards were created where the population grew the fastest. Over our study window, the population grew by 108.0% in areas that added a ward, compared with 38.5% where the boundaries were realigned but the number of wards held constant, and only 14.3% where a ward closed. Additionally, new-ward areas saw pronounced single-family home construction booms in the years immediately preceding the change (Appendix B.6).

How disruptive are boundary changes to ward social networks? As Figure 1a illustrates, a reassigned voter's new ward often combines members of several formerly distinct congregations. Among individuals who experienced a ward boundary change, the median person found themselves in a post-change ward where 16.7% of their fellow ward members had not been in their ward prior to the change. Exposure is unevenly distributed across voters (Figure 1b): roughly 41.6% of individuals experienced a change in which 5% or fewer of their assigned congregation were new, while 26.0% found themselves in a ward where at least half of their fellow ward members were new. Most of the identifying variation therefore comes from a minority of substantially disrupted voters. Stake presidents confirmed that the disruption is real. Several volunteered that a new line attenuates friendships across it, even between households a few dozen feet apart, because the weekly and midweek activities that sustained those friendships no longer coincide. For our design, what matters is not disruption itself but the difference in composition between the two new wards a former congregation is divided into. Voters on both sides of a new line are disrupted, but they differ in whom they are now assigned to. Section 7.5 uses the distinction between retained and newly assigned ward-mates directly, investigating whether old or new ties are more important drivers of peer effects.

2.4 The Utah Political Landscape

Utah leans heavily Republican. In 2012, 45% of registered Utah voters were Republicans, 9% were Democrats, and 46% were Unaffiliated—a category that pools voters registered with no party (the large majority) with registrants of minor parties. By 2020, those shares had shifted to 52% Republican, 14% Democratic, and 34% Unaffiliated. Much of this shift came from party switching rather than from entry into and exit from the electorate. Appendix Figure A3 traces registration flows among all Utah voters observed in both 2012 and 2020. Over this period, 27.9% of 2012 Unaffiliated voters moved to the Republican Party and 12.7% to the Democratic Party, while 12.7% of 2012 Republicans and 28.3% of 2012 Democrats left their party.

Unaffiliated voters thus make up a large share of the Utah electorate, and they feature in several of our main findings, so it is worth understanding how they lean. We characterize their presidential vote choice using self-identified Independents in the Cooperative Election Study (CES), which is the closest available survey proxy for the Unaffiliated registrants we study. Utah Independents have generally leaned left in recent presidential elections—50% voted for Obama versus 37% for McCain in 2008, and 52% for Biden versus 41% for Trump in 2020—with 2012 the exception, when Mitt Romney's candidacy, unusually popular in Utah given his LDS faith, pulled Utah Independents sharply rightward (58% Romney vs. 35% Obama). In 2016, a large share defected from both major-party candidates to the independent candidacy of Evan McMullin, an LDS conservative (Appendix Figure A4).

Utah's primary elections give party registration a second, strategic dimension. Throughout our study period, the Republican primary was closed; only registered Republicans could participate. The Democratic primary, by contrast, was open to both Democrats and Unaffiliated voters.[9] Because Republicans dominate Utah politics, the Republican primary is

[9]Utah law lets each party decide which voters may participate in its primary. Before each primary election, a party files a statement with the lieutenant governor declaring which parties' members, and whether Unaffiliated voters, may vote in its primary (Utah Code §20A-9-403). The Republican Party has restricted its primary to registered Republicans since 2002—before which Utah did not register voters by party at all (Bernick, 2002)—while the Democratic Party has allowed Unaffiliated voters to participate throughout our study period (Hutson, 2020; Roche, 2020). Unaffiliated voters who wish to vote in the Republican primary may also affiliate with the party at the polls on election day.

frequently the decisive election, giving non-Republicans a strategic incentive to register Republican. This incentive was especially visible in 2020, when a competitive four-way Republican gubernatorial primary prompted organized public campaigns encouraging Democrats and Unaffiliated voters to register as Republicans in order to participate. The campaign of Jon Huntsman, Jr., a moderate former governor, ran social-media ads urging supporters to register Republican, and the former chair of the Utah Democratic Party publicly switched parties and encouraged fellow Democrats to do the same (Roche, 2020). In the eight weeks before that primary, roughly 73,000 Utahns registered as Republicans, an increase of about 10% in the party's registration (Hutson, 2020). Episodes of statewide strategic switching like these are absorbed by our design, which compares registration trends across newly drawn ward boundaries within the same neighborhood at the same time. However, understanding these strategic incentives still matters for interpreting what a party switch means in this setting.

These institutions clarify what party registration measures. A voter's registration can reflect internal partisan identity—beliefs, ideology, and attachment to a party—as well as strategic considerations, most importantly access to a closed primary. The two motives, however, map asymmetrically onto the possible switches. Strategic motives can explain why an Unaffiliated voter or a Democrat would register Republican. They cannot explain why an Unaffiliated voter would register as a Democrat, nor why a Democrat would become Unaffiliated. Neither move changes the set of primaries in which the voter may participate. Likewise, a Republican has essentially no strategic reason to leave the party—and, in particular, no strategic reason to register as a Democrat rather than becoming Unaffiliated, since either allows participation in the Democratic primary. Movements toward the Democratic Party, which feature prominently in our results, are therefore difficult to attribute to strategic registration. Moreover, even registration changes that begin strategically appear to acquire genuine political content. Gerber, Huber and Washington (2010) randomly encouraged unaffiliated voters to register with a party ahead of a closed primary and found that the induced registrations subsequently shifted partisan identification and political attitudes. We return to this asymmetry when interpreting our main results in Section 6.

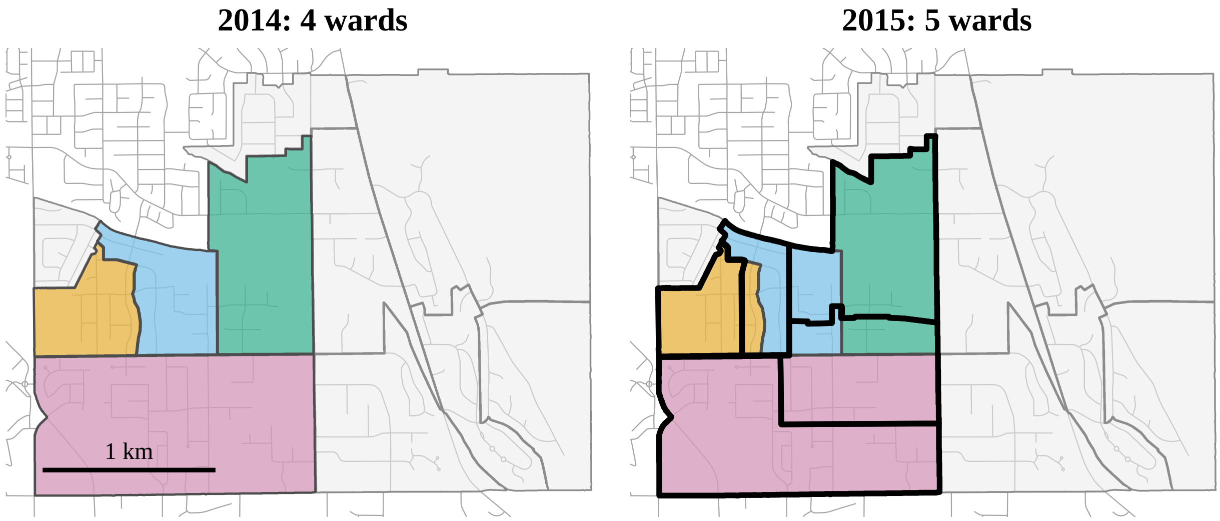

(a) Example of a ward boundary change

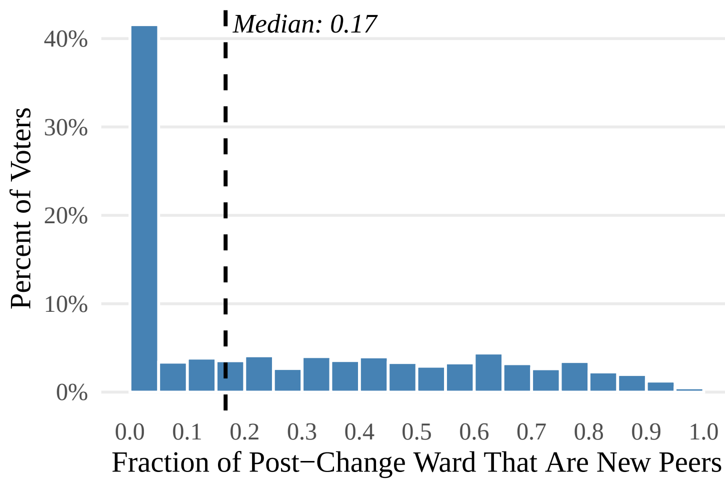

(b) Share of new ward-mates among affected voters

Figure 1: Ward boundary changes and the peer disruption they cause. Panel (a): the left map shows a stake as of 2014, with gray wards unaffected by the change and each subsequently redrawn ward shaded its own color; the right map superimposes the new (2015) boundaries in bold black over the same stake, so that differently colored areas within each bold outline are groups of members who previously attended different congregations. Roads are drawn faintly for scale. Panel (b): the voter-weighted distribution of the fraction of each affected voter's post-change ward roster that was not in their pre-change ward. Each voter's exposure is computed from their own old-ward-to-new-ward transition, pooling all 2013–2017 changes; a ward adjusted in more than one year contributes one observation per change. The dashed line marks the median (16.7%).

3 Data

We assemble a panel of Utah voters by linking four sources: (1) TargetSmart voter files (2012–2021), which provide voter addresses and political outcomes; (2) LDS ward boundary maps (2013–2017), which define each voter's peer group and the boundary changes we exploit; (3) CoreLogic property records, which provide property-level controls; and (4) school boundary shapefiles, which capture school-driven residential sorting. Our main analyses compare each voter's party affiliation at two endpoints, as of the 2012 and 2020 presidential elections, which bracket the 2013–2017 boundary changes we study.[10] Our main analysis sample comprises 121,876 registered voters observed in both 2012 and either 2020 or 2021 who lived within 400 meters of a ward boundary newly drawn between 2013 and 2017, with voters on both sides of the line. Voters near more than one new boundary enter the sample once per boundary, as we describe below.

Two supplemental analyses draw on more recent ward boundary data. Since 2023 we have scraped ward boundaries each quarter from the Church's online meetinghouse locator, which also lists each ward's bishop. We use these data in two ways. First, as a validation exercise, we test whether boundaries drawn in 2023–2026 predict 2012–2020 party switching among our TargetSmart voters. Second, to study the effect of bishops on their members, we link the quarterly bishop listings to the L2 voter file (2015–2025), which covers those years. For each voter in the 2023 L2 file we observe their party registration from 2015 to 2025 and the party affiliation of each bishop they were assigned between 2023 and 2026.

Neither voter file records religion, so we cannot identify LDS members directly. We address this by studying Utah, where most counties are over 70% LDS (Appendix E.5 describes how we estimate LDS density within counties, at the ward level), and by treating every voter as assigned to the ward containing their address. The resulting estimates are intent-to-treat effects on the full electorate, which understate the effect on members for reasons we quantify in Section 6.1. The remainder of this section describes our main data sources and sample construction, and Appendix Table C4 reports summary statistics for the analysis sample by baseline party.[11]

[10]Most registration changes occur during presidential election years (Appendix Figure C1), making 2012 and 2020 natural measurement points. Because of the low rates of party switching in non-election years, we do not include intermediate snapshots in our main analyses. The 2016 election, the one intermediate presidential election, warrants separate comment. We cannot cleanly order it relative to our observed boundary changes; it falls after 2013–2014 and 2014–2015 boundary changes, but, because the exact timing of boundary changes is unobserved, it is unknown whether 2015–2016 and 2016–2017 boundary changes precede or postdate it. Changes measured against 2016 would therefore mix pre- and post-treatment exposure.

[11]We supplement the Targetsmart voter data with additional data from a similar vendor, L2, covering 2014-2025. We use these data for secondary analyses of church leadership changes from the 2020 to 2024 presidential election. We rely on the TargertSmart panel in our main analyses because it spans the time

3.1 Voter Panel

To measure political behavior, we use individual-level voter registration lists from TargetSmart, a commercial vendor that cleans and standardizes publicly available state registration files. The data consist of annual snapshots of Utah's registered voters, 2012–2021, with party registration, turnout history, name, residential address, date of birth, sex, and race for each voter-year.[12] Each snapshot records registration as of its pull date, while validated turnout for a November election enters the following year's snapshot. Consistent with this timing, year-over-year party switching spikes in the election-year files themselves (Appendix Figure C1). TargetSmart geocodes voters' street addresses, which allows us to assign each voter to a contemporaneous ward and to measure distance to ward boundaries. We construct our primary outcome, a three-way party-registration category—Republican, Democratic, or Unaffiliated—from TargetSmart's party-registration field, pooling unaffiliated and third-party registrants into the Unaffiliated category.[13]

TargetSmart links voters across years using state registration records, national change-of-address databases, and proprietary linking methods. We build on this vendor linkage by further deduplicating records and identifying movers based on name, date of birth, and voting history. Appendix C.1 describes these cleaning steps in detail. From this process we construct a longitudinal panel of 753,451 registered Utah voters observed in both 2012 and either 2020 or 2021—69.0% of the 1,091,892 registrants in our cleaned 2012 cross section.[14] Voters who move remain in the panel and their treatment is always defined by their baseline (2012) residence. Section 6.4 examines attrition from the panel directly, and Appendix C.2 validates our cleaned panel against published state totals for registration, turnout, and party shares.

3.2 Ward Assignments

To measure church peer groups, we digitized LDS ward boundaries for 2013–2017, approximately 4,000 ward polygons per year, from boundary maps published by the Church, using

period over which we have LDS ward data.

[12]TargetSmart imputes race from surname and local demographics, following methods akin to Imai and Khanna (2016), and sex from name and commercial records.

[13]Minor-party registrants (Libertarian, Green, Constitution, Independent American, and a residual "other" category) are 1.8% of 2012 registrants and 3.9% of the pooled Unaffiliated category, so the pooling has little bearing on our results.

[14]We pool 2020 and 2021 as our endline because some voters appear in 2021 but not 2020, and 2021 registrations reflect post-2020-election affiliation (party switching between the 2020 and 2021 files is minimal; see Appendix Figure C1). This increases our effective panel size without changing the interpretation of the outcome. When voters appear in both 2020 and 2021 we use their outcomes from the 2021 file. We do not pool 2012 with 2013 at the start of the panel because boundary changes begin in 2013, so pooling would risk treating post-treatment observations as baseline.

a combination of automated image processing and manual review. Appendix C.3 describes the procedure.[15] We then assign each voter-year to a contemporaneous ward based on the voter's geocoded address. The digitized maps assign a ward to 98.5% of geocoded 2012 registrants (Appendix Table C4), with the remainder living in areas without digitized coverage. Across wards, the median land area is 0.5 square kilometers, the median size is 247 registered voters, and the median partisan composition is 50.3% Republican and 6.4% Democratic.[16]

We identify boundary changes by comparing the full statewide set of ward polygons in each consecutive pair of years. A new boundary line is defined as any line dividing two neighboring wards in one year that bisects a ward from the previous year. Imperfections in digitization could cause small discrepancies in a ward's polygon across years to be misclassified as new boundary lines, so after programmatically identifying candidate lines we manually checked each against the source maps and removed false positives. This left us with a set of 2,205 new boundary lines (443 from 2013–2014, 468 from 2014–2015, 658 from 2015–2016, and 636 from 2016–2017).

3.3 Property and School Controls

CoreLogic property records. To capture neighborhood-level differences in housing stock that may coincide with ward boundaries, we add property-level controls from CoreLogic, a commercial vendor that aggregates publicly available property records from county assessors. CoreLogic provides annual records for 2012–2021 with address, assessed value, square footage, lot size, year built, and bedroom and bath counts, all of which we fix at their baseline (2012) values. Because these controls are meant to describe the home a voter actually lives in, we restrict property records to owner-occupied single-family homes (487,060 homes in 2012), for which the owner and the resident are most likely the same person. We link properties to TargetSmart voters by address and owner name, first at the exact address and then within 100 meters for owners still unmatched (Appendix C.5 details the procedure). This produces matches for 73.3% of homeowners, consistent with near-complete linkage of registered owners given that registration is higher among homeowners than in the general population,[17] and adds property characteristics for 40.4% of the registered voters in our 2012 cross section. Linked voters are older, more likely to have voted in 2012, and more Republi-

[15] We collected the maps from a repository on the Church's website that was publicly accessible at the time but is no longer. The repository carried maps only between 2013 and 2017. Maps for later years require a separate per-map application to the Church. Our placebo boundaries therefore come from quarterly scrapes of the Church's Meetinghouse Locator (Section 3.5).

[16] Interquartile ranges are [0.3, 1.8] km^2, [183, 338] voters, [40.0%, 59.0%] Republican, and [4.0%, 9.9%] Democratic. Figure C2 plots the full weighted and unweighted distributions.

[17] Utah's 2012 registration rate was 59.4% of the voting-age population (U.S. Census Bureau, 2013).

can than unlinked voters, which is consistent with the latter being mostly renters (Appendix Tables C6 and C7).[18]

School boundaries. To capture school-related residential sorting that could confound ward-level comparisons, we asked each Utah school district for elementary school catchment shapefiles. We obtained detailed boundaries from 22 of 41 districts, with rural districts being the most likely to lack digital records. For voters in these districts we assign an elementary catchment, covering 94.9% of the registered voters in our 2012 cross section, and for voters elsewhere we assign the broader school district as a fallback. When a specific year's shapefile is unavailable, we use the closest available year. Appendix Table C3 lists districts and availability of elementary school boundaries.

3.4 Analysis Sample

Our main analysis focuses on voters who, in 2012, lived in wards that were split between 2013 and 2017; that is, wards for which at least one 2012 resident ended up on each side of a newly drawn boundary.[19] We further restrict attention to voters living within 400 meters (about four city blocks) of a newly drawn line, so that we compare voters in the same immediate neighborhood who differ mainly in which side of the line they fall on. We chose 400 meters as the widest bandwidth at which that comparison remains credible, and our estimates are stable across bandwidths from 50 to 600 meters (Section 4.1 and Appendix Figure F4). Of the 2,205 newly drawn boundary lines, 1,516 have panel voters living within 400 meters on both sides. Of the 753,451 voters in our cleaned panel, 121,876 live within 400 meters of at least one such line, and roughly four in five are near exactly one (98,613 of 121,876). Voters near several lines enter the sample once per line, which our clustering accounts for (Section 4.1).

Two features of this sample are worth noting. First, because our panel requires a voter to be registered across an eight-year window, our sample under-represents temporary residents such as college students, short-term renters, and recent movers (who are more difficult to link across voter-file snapshots). As a result, a higher share of our analysis-sample voters are linked to CoreLogic homeowner records (46.3%) than of voters in the full 2012 cross section (40.5%), and analysis-sample voters are also somewhat more likely to have voted in 2012 and to be registered as Republicans. Second, because ward boundary changes are often triggered by local population growth from new home construction, the median home in our analysis

[18] We also link CoreLogic mortgage records to HMDA loan-application records to recover applicant income. Because the resulting match rate is low (roughly 27% of 2012 originations), we omit income from our main analyses. Appendix C.6 describes the linkage procedure.

[19] Ward assignment is fixed at 2012 residence so that treatment is not affected by moves made in response to boundary changes.

sample is about a decade newer than in the full cross section. Appendix Table C5 reports the full comparison.

Appendix Table C4 reports summary statistics for each baseline-party subsample, which are the samples our regressions use. Two features matter for what follows. First, the Democratic subsample is less than a seventh the size of either of the others, which is why its estimates are far less precise and appear in the appendix. Second, between a quarter and a third of voters near an eligible boundary are not observed at endline, and over half of those who are had moved by 2020. Treatment is defined by 2012 residence, so moving does not change a voter's assignment, but selective attrition could, and Section 6.4 bounds its effect. Party switching in the sample broadly mirrors the statewide flows (Appendix Figure A3).

3.5 Placebo Sample

For our main validation exercise, we construct a parallel placebo sample by applying the same design to voters living near ward boundaries that were drawn after the 2012–2020 study window. Because these boundaries did not yet exist during our outcome period, any estimated "effect" on 2012–2020 political behavior necessarily reflects a problem with our design rather than the causal impact of a boundary change. It therefore serves as a direct test of the identifying assumption that boundary placement is unrelated to differential trends in political outcomes. We present the placebo estimates in Section 5.1.

Since 2023, we have scraped ward boundaries each quarter from the LDS Church's Meetinghouse Locator tool, a Google Maps-style interface that returns a ward assignment for any Utah address submitted to it.[20] Our current placebo sample uses boundary changes observed between Spring 2023 and Spring 2026, yielding 2,173 newly drawn boundary lines, of which 1,101 have panel voters within 400 meters on both sides; the resulting placebo sample contains 103,335 voters (Appendix Table C8). Appendix Table C8 compares the placebo sample to our main analysis sample. Placebo-sample voters are somewhat less Republican, own less expensive and older homes, and are less likely to be rural. These differences are as expected since construction booms often draw more Republican-leaning voters into new homes in previously rural places, and the booms that would later trigger the 2023–2026 changes had not yet occurred in 2012.

[20] Appendix C.9 details the scraping procedure and address list.

4 Empirical Design

We use a geographic difference-in-discontinuities design to estimate the causal effect of peer partisan composition on party affiliation and voter turnout. Our design compares voters who belonged to the same ward at baseline (2012) but who, through a subsequent boundary change, were reassigned to neighboring wards with different partisan compositions. We ask whether voters assigned to wards with a higher share of Republicans (or Democrats) experience different 2012–2020 trends in these outcomes than voters assigned to the neighboring ward.

The design combines four features, each addressing a different threat to identification. A spatial discontinuity strips out smooth partisan sorting across space. First-differencing nets out cross-boundary differences in levels. Reassignment to congregations postdates residential choices, which rules out sorting across the boundaries at baseline. Boundary fixed effects absorb area-level differences in political trends. Together these features form a difference-in-discontinuities framework that is substantially more credible than any one component alone (Butts, 2023).

Figure 2 illustrates the design with an example from our data. Panel (a) shows a single ward at baseline (2012), labeled Ward A, with each point representing a registered voter colored by party. The dashed lines preview a boundary change that, from the perspective of 2012, still lies in the future—a new line will split Ward A, combining its northern half with neighboring territory to the northwest to form Ward B, and its southern half with neighboring territory to the southwest to form Ward C. Crucially, voters could not have sorted across a boundary that did not yet exist. Panel (b) shows the post-change configuration. Our design compares the trends in political outcomes of Ward A voters on opposite sides of the new line, while controlling for smooth spatial gradients via the running variable. In this example, Ward A voters assigned to Ward C found themselves in a 69% Republican ward, while their former ward-mates assigned to Ward B found themselves in a 54% Republican ward. Our treatment variable, described in more detail below, is a function of this difference, with peer composition and residential locations measured at baseline.[21] Note that ward partisan composition is computed over each new ward's full roster, including the voters beyond 400 meters from the new line who appear as faded dots in the figure. Only the voters near the line, the saturated dots, enter the estimation sample.

We discuss each feature in turn, beginning with the spatial discontinuity. Because Republicans and Democrats tend to cluster in different neighborhoods, voters living in similar

[21] The quoted shares are overall ward means across all 2012 registered voters, as reported below the figure. In estimation we use leave-one-out means, which differ slightly across individuals by their own baseline party (Section 4.2).

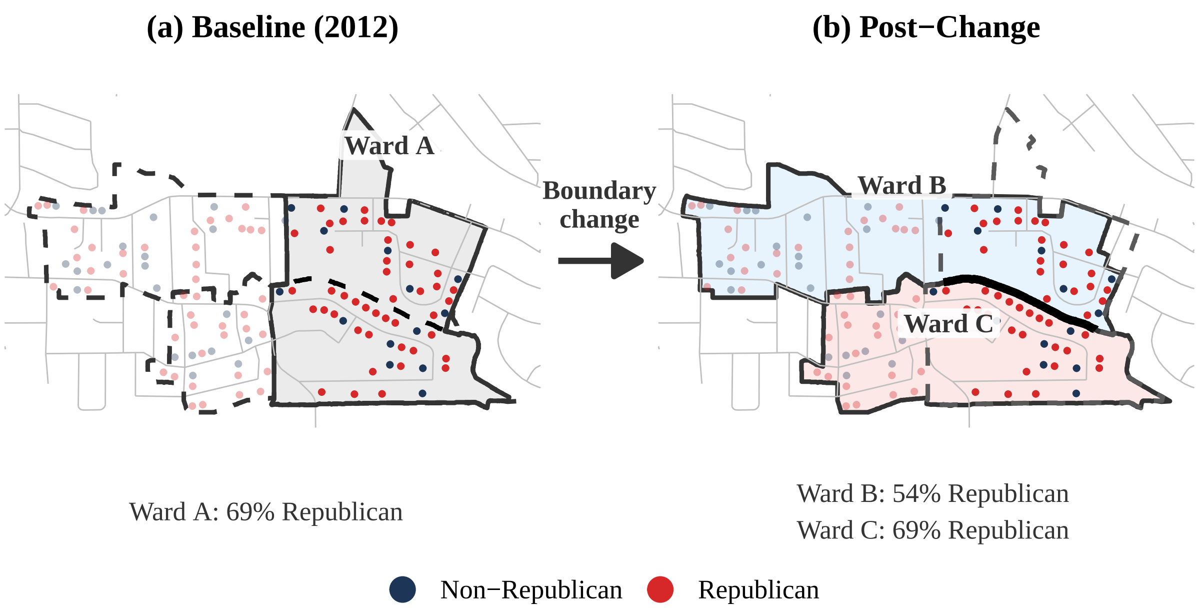

Figure 2: Research Design Illustration. Panel (a) shows the baseline (2012) ward, Ward A, with each point a registered voter colored by party registration; to keep the figure legible, one voter from each single-family household is plotted. The future boundaries are superimposed in dashed lines. Voters who are both in Ward A and within 400 meters of the new boundary line are shown in saturated color, indicating that they make up the analysis sample; faded voters are used only to construct peer exposure measures. Panel (b) shows the post-change configuration: Wards B and C, shaded by partisan lean (pink = more Republican, blue = less Republican), with the new boundary in bold and the old ward boundaries dashed. The text row below the panels reports each ward's Republican share among all 2012 registered voters.

localities but different neighborhoods may differ systematically in political behavior absent any peer effect. We therefore compare voters who live close to each other but just across a new ward boundary. That is, voters who made similar decisions about where to live but face different congregation partisan compositions purely because of where the line falls (Keele and Titiunik, 2015; Keele, Titiunik and Zubizarreta, 2015). In our estimation, a smooth function of distance to the boundary absorbs the underlying spatial gradient in partisan sorting, so that identification comes strictly from the discontinuous change in peer composition at the boundary. This is an advantage relative to a simple difference-in-differences design, which would attribute all cross-boundary variation in trends to treatment, including variation driven by smooth spatial sorting.

Second, by differencing outcomes between a baseline period (2012) and an endline period (2020), we relax the identifying assumption from one about levels to one about trends. We do not require that covariates be balanced or smooth across the new boundary, only that trends in political outcomes be smooth through it. This is a much weaker assumption than that required for a static spatial regression discontinuity, allowing for any imbalances in levels across the boundary that are constant over time (e.g., differences in the housing stock or school assignment). Differencing also reduces concerns about intentional gerrymandering on partisanship, since it is far easier to observe who is Republican than who is trending Republican.

Third, boundary lines are drawn and announced only after administrative approval (Section 3), years after the voters in our sample chose where to live, so voters at baseline could not have sorted across boundaries that did not yet exist. Had we instead studied boundaries that persisted throughout the study window, both components of the design would inherit a sorting problem, because the line would already have existed at baseline to sort across. A static spatial regression discontinuity would need to rule out sorting on levels. With first-differencing, the concern is instead sorting on trends, which is weaker but not innocuous. An outside observer may have a difficult time discerning who is trending Republican, but voters know their own trajectories and may sort accordingly. A Democrat drifting rightward may be open to living in a more Republican ward in a way that a committed Democrat is not. Because the boundaries we study did not exist when residential choices were made, neither form of sorting can operate. The identifying assumption of smooth potential-outcome trends across the boundary is therefore easier to satisfy.

Fourth, our estimation includes a fixed effect for each newly drawn boundary, ensuring that all comparisons are between voters near the same new line who shared a ward before it was redrawn. Without them, the design would implicitly compare voters near boundaries in, say, downtown Salt Lake City to voters near boundaries in rural Utah County, which are areas that differ dramatically in both the level and trajectory of partisanship for reasons entirely unrelated to peer effects. Boundary fixed effects restrict identification to within-boundary contrasts, netting out any area-level differences in political trends.

Finally, we discuss the additional control variables. The secondary considerations stake presidents describe (Section 2.3), balancing children and youth across wards, mixing wealthier and poorer neighborhoods within wards, and following roads and other natural dividers, are, from the design's perspective, level characteristics. Where a new line places more young families or a wealthier neighborhood on one side, that difference is constant over the study window and absorbed by first-differencing our outcome variables. Our preferred specification therefore includes no individual-level controls. One remaining concern is that these considerations could proxy for differential trends. To address this, robustness specifications control for differential trends by demographics, property characteristics, and school catchment zones, and their inclusion does not affect our estimates.

4.1 Sample Restrictions

We restrict the sample in two ways. First, as described in Section 3.4, we focus on voters living within 400 meters of a newly drawn boundary. Approximately 18% of these voters live within 400 meters of more than one new boundary.[22] We include such voters once for each nearby boundary and cluster standard errors by voter, as well as by boundary side, to account for the duplication. Peer composition itself is measured over each ward's full roster of registered, street-geocoded voters in our panel, not just those near the boundary.

Second, we estimate all specifications separately by baseline (2012) party affiliation: Republicans, Democrats, and Unaffiliated voters. This serves two purposes. First, because our outcomes are binary indicators measured in changes, splitting the sample yields coefficients with a clean transition interpretation: for 2012 Republicans, the coefficient on Republican peer exposure captures the effect of peer composition on remaining Republican; for 2012 Democrats or Unaffiliated voters, it captures the effect on becoming Republican. Second, splitting the sample allows us to measure peer composition using a simple leave-one-out ward mean without introducing exclusion bias. Exclusion bias is a mechanical, finite-sample bias that arises because a person cannot be their own peer. This causes own outcomes to be negatively correlated with leave-one-out peer means, which can contaminate peer-effects regressions even under random assignment (Angrist, 2014; Caeyers and Fafchamps, 2024; Guryan, Kroft and Notowidigdo, 2009). Splitting by baseline party eliminates this channel entirely, because the leave-one-out measure is then constant within each ward-by-party cell—no individual's own affiliation moves their own regressor. Appendix Table C4 reports summary statistics for the three resulting subsamples.

4.2 Estimating Equation

For each new boundary, we compare changes in party registration between the 2012 and 2020 elections for voters living just on either side of a redrawn boundary, all of whom belonged to the same ward in 2012. Because the two new wards differ in partisan composition, voters on one side of the line are assigned to a more Republican (or more Democratic) set of peers than their former ward-mates on the other side, and we ask whether their registration changes differ in proportion to that gap.

Formally, let b index a newly drawn ward boundary separating two new wards, w and w' (Wards B and C in Figure 2). For each individual i near boundary b, define $\bar{Z}_w$ as the leave-one-out share of Republicans (or Democrats, depending on the specification) in ward

[22]Specifically, 18% of Republicans, 15% of Democrats, and 18% of Unaffiliated voters live near multiple boundaries kept in their sample-specific analyses.

w, with partisanship and addresses both measured at baseline (2012).[23]

Define the running variable r_{ib} as signed distance to the boundary, normalized so that $r_{ib} \geq 0$ places individual i in whichever of the two wards has the higher partisan share.[24] Let

$$D_{ib} = \mathbf{1}[r_{ib} \geq 0]$$

indicate assignment to the high-exposure side, and define

$$I_b = \left|\bar{Z}_{B(b),2012} - \bar{Z}_{C(b),2012}\right|$$

as the absolute difference in baseline peer composition across boundary b.

In the example of Figure 2, with treatment defined by exposure to Republican peers, Ward B is 54% Republican and Ward C is 69% Republican, so Ward C is the high-exposure side: voters assigned to Ward C have $D_{ib} = 1$, voters assigned to Ward B have $D_{ib} = 0$, and $I_b \approx 0.15$.

We estimate the following first-differenced specification separately by baseline party affiliation:

$$\Delta Y_{ib} = \beta\,(D_{ib} \times I_b) + \delta D_{ib} + \mu_b + f(r_{ib}) + \mathbf{X}_i'\gamma + \varepsilon_{ib}, \tag{1}$$

where $\Delta Y_{ib} \equiv Y_{i,2020} - Y_{i,2012}$ is the change in a binary political outcome (Republican registration, Democratic registration, or Unaffiliated registration; Section 7.2 applies the same specification to changes in primary- and general-election turnout).[25] The interaction $D_{ib} \times I_b$ is the treatment variable equaling the cross-boundary difference in baseline peer composition for voters on the high-exposure side and zero for voters on the low-exposure side. The coefficient β therefore measures how a unit increase in peer partisan exposure (a 0-to-1 change in peer share) affects the change in the individual's political outcome. Note that voters on both sides of a new line experience the boundary change; the two sides shared a ward at baseline, and both end up in reorganized wards. The design therefore compares two treated groups that differ in the size of the induced change in peer composition, rather than a treated group against unaffected controls.

The specification includes boundary fixed effects μ_b, so all comparisons are within the same boundary pair. It also includes a flexible function of signed distance, $f(r_{ib})$, which we parameterize as quadratic in distance with coefficients estimated separately on each side of

[23]We drop the standard $-i$ subscript on the leave-one-out share for notational simplicity and because within a baseline-party subsample, $\bar{Z}_w$ is identical for all voters in ward w (Section 4.1).

[24]If the two sides have exactly equal shares, $I_b = 0$ and the treatment is zero regardless of which side is labeled high-exposure; such ties are rare and contribute nothing to identification.

[25]Endline registration is taken from the 2020 voter file, or from the 2021 snapshot for voters not observed in 2020 (Section 3.1).

the boundary.[26] We use distance to the boundary rather than a polynomial in latitude and longitude because distance is defined the same way at every boundary, so a single function controls for the tendency of partisans to live near one another wherever the line falls. A global polynomial in coordinates would capture only coarse statewide gradients—in Utah the fact that Democrats cluster in Salt Lake City—and boundary-specific polynomials overfit, since many boundaries have few voters within the bandwidth (see Keele and Titiunik, 2015, for a discussion of the alternatives). Estimates are similar with a linear distance control and with a donut that drops voters within 5 meters of the line (Appendix Figure F5). The vector $\mathbf{X}_i$ includes controls (age, sex, race, property characteristics, school-zone indicators) measured at baseline, although, as discussed above, our preferred specification omits these. The stand-alone regressor D_{ib} ensures that β is identified from how outcome differences scale with the dose I_b rather than from any level difference between sides, and it provides a validation check on the design: as I_b approaches zero, there should be no difference between being assigned to the high- and low-exposure sides. Our estimates of δ are small, within three percentage points of zero, and mostly statistically insignificant (Appendix Table F1). Standard errors are clustered two-way by ward-boundary side (each side of each new boundary is one cluster) and voter, which treats the two sides of a line as separate clusters. Our results are robust to using Conley spatial standard errors, where we allow correlation among all voters within 5 kilometers regardless of the line (Appendix Figure F5).

The continuous treatment intensity I_b improves both power and interpretability relative to a binary treatment indicator. Many adjacent wards have similar partisan compositions, so a simple indicator for crossing a boundary would pool together large and small changes in peer exposure. By interacting the boundary-crossing indicator with the size of the induced change, we allow the estimated effect to scale with treatment dose. This specification embeds a linearity assumption—that the marginal effect of peer composition is constant across the distribution—which we assess in Section 6.2. Figure 3 shows the distribution of I_b across boundary pairs. The median cross-boundary difference in Republican share is about 5.2 percentage points, with the top decile of boundaries exceeding 14.2 percentage points.

Because a substantial fraction of Utah voters are unaffiliated (46% in 2012), exposure to Republican peers is far from the inverse of exposure to Democratic peers. In fact, the high-Republican side is also the high-Democratic side for 35% of boundaries. Appendix Figure C3 quantifies the trade-offs. A cross-boundary increase in Republican share comes mostly at the expense of Unaffiliated rather than Democratic share, while an increase in Democratic share primarily comes at the expense of the Republican share. To investigate which partisan peer measure best explains party switching, we also estimate a specification that includes

[26] In our baseline specification, we include r_{ib} and r_{ib}^2, each interacted with D_{ib}.

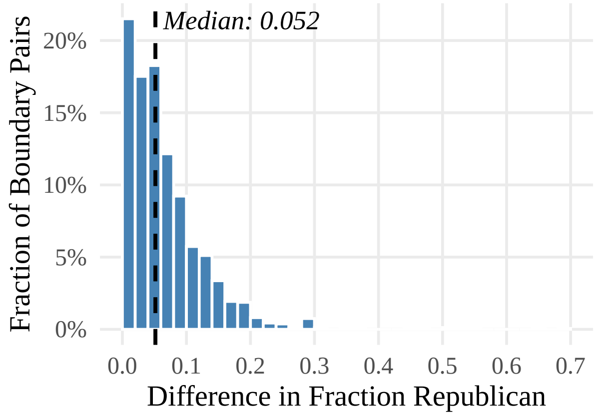

(a) Republican-share treatment

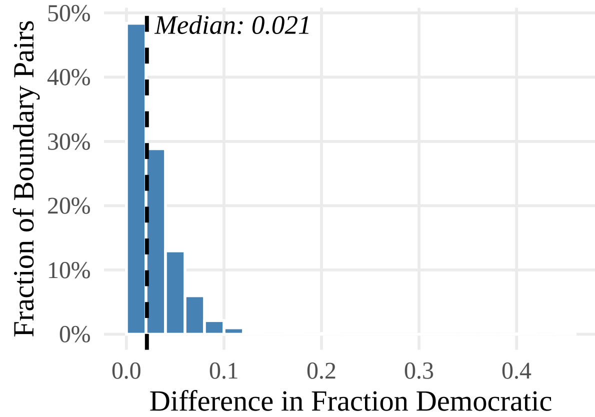

(b) Democratic-share treatment

Figure 3: Distribution of Treatment Intensity Across Boundary Pairs. Each observation is a newly drawn ward boundary. The horizontal axis measures the absolute difference in baseline (2012) fraction Republican (panel a) or Democratic (panel b) between the two wards adjacent to the boundary. The dashed line marks the median.

both Republican-share and Democratic-share exposures jointly:

$$\Delta Y_{ib} = \beta_r \left(D_{ibr} \times I_{br}\right) + \delta_r D_{ibr} + \beta_d \left(D_{ibd} \times I_{bd}\right) + \delta_d D_{ibd} + \mu_b + f\left(r_{ib}\right) + \mathbf{X}_i' \gamma + \varepsilon_{ib}, \tag{2}$$

where D_{ibr} and D_{ibd} are indicators for being on the high-exposure side of the boundary for the Republican-share and Democratic-share treatments, respectively, and I_{br} and I_{bd} are the corresponding treatment intensities. Equation (2) estimates the effect of being assigned to a more Democratic ward holding the Republican share constant (and vice versa), thereby measuring the effect of exposure to additional peers from each party on its own, independent of the extent to which such peers substitute for opposite-party peers. We report these estimates in the "Both treatments" block of our main results tables.

Before turning to identification, we note two considerations that govern how to interpret β. First, because the voter file does not record religion, we can restrict neither the sample nor the peer measure to LDS members. Every registered voter near a boundary enters the sample, and peer composition is measured over the assigned ward's full electorate (Section 4.1). Our estimates are therefore intent-to-treat effects of assigned peer composition on the full near-boundary electorate. We are interested, however, in the causal effect of LDS wards on their members. We show both that our main ITT effects are conservative[27] and that they can be

[27]The inclusion of non-LDS voters in our sample attenuates our estimates because, while non-LDS voters may be influenced by their LDS neighbors and vice-versa, their interactions are governed by residential proximity rather than by ward assignment. Thus, for a non-LDS voter, there should be no causal effect of being directly on one side of a ward boundary as opposed to the other.

rescaled to recover our effects of interest. Section 6.1 and Appendix E quantify the implied attenuation and adjust our estimates for it.

Second, we map β to standard peer-effects channels. Manski (1993) distinguishes three reasons that members of a group behave alike: *endogenous* effects, in which individuals respond to their peers' behavior; *contextual* effects, in which individuals respond to their peers' predetermined characteristics; and *correlated* effects, in which group members share environments or have selected into the group on similar unobservables. The reflection problem arises when an outcome is regressed on the contemporaneous mean outcome of the group, so that each member's behavior is both a cause and effect of the group mean. Our regressor is instead the predetermined 2012 composition of the ward a voter is later assigned to, so there is no simultaneity. Persistence in a voter's own preferences cannot reintroduce it since the first difference removes any fixed component of the voter's partisanship, and estimating within baseline party fixes the voter's own 2012 affiliation, so the only variation in the regressor is across the new line. Furthermore, because boundary changes quasi-exogenously assign voters to new wards, while leaving the rest of their environment unaffected, our estimates of β are separately identified from any correlated effects (Section 5 presents the supporting evidence)

The coefficient should therefore be read as the total effect of a shock to peer composition, combining Manski's endogenous and contextual channels and inclusive of any equilibrium feedback among reassigned voters as they influence one another over the study window. We view this composite as the policy-relevant object, since any real-world reassignment of peer groups, a redrawn boundary, a rezoned school, a reorganized unit, triggers the same equilibrium adjustment (Moffitt, 2001).[28]

Appendix D gives Equation (1) a structural interpretation. Suppose each member forms a circle of friends by drawing at random from their ward and conforms toward the average affiliation of those friends with weight ρ. Although the researcher never observes these friendship ties, under random friending, the expected composition of a voter's friend circle equals the leave-one-out ward mean, regardless of the number of friends drawn. In this case, the linear-in-means coefficient recovers ρ directly. If members instead befriend copartisans disproportionately, the expected friend circle tilts toward the voter's own type and β understates ρ, so the estimate is a lower bound on the influence of a member's actual friends (Appendix D.1). The same appendix shows that influence broadcast to the whole congregation, rather than transmitted between friends, loads on the same regressor, so that

[28] In Section 7.1 we present evidence that our estimates of β are driven by peers' political behavior rather than other characteristics, such as race. Hence, while the composite is policy relevant in our setting, we argue that our estimates likely reflect primarily endogenous, rather than contextual, peer effects.

β recovers the sum of the two channels; Section 7.4 uses the demographic composition of peer exposure to separate them. The framework, extended in Appendix D.2, also lets us separate the influence of retained from newly assigned ward-mates, which we exploit in Section 7.5.

5 Identification and Validation

Our design requires that, absent the boundary change, voters on the two sides of a new line would have followed the same trends in party registration. As discussed in Section 2.3, stake presidents do not draw boundaries at random. They respond to membership imbalances, follow roads and other natural dividers, and balance wards on the number of children and youth, the number of eligible leaders, and the mix of wealthier and poorer neighborhoods. Where these features are correlated with partisanship, new lines may coincide with differences in the levels of political variables across the line. The design allows for such level differences. What it does not allow is for new lines to coincide with differences in trends. We present three pieces of evidence to test for this potential threat. The first is a placebo test using boundaries drawn after the study window. The second is a simulation exercise that tests whether the redrawing process sorts voters by partisanship. The third is a balance test in baseline levels across newly drawn boundaries, paired with a robustness check that allows trends in partisanship to differ by those characteristics.

5.1 Placebo Test: Future Boundaries

Our identifying assumption requires that the placement of new ward boundary lines be uncorrelated with trends in partisanship. We test this directly using a placebo design based on ward boundaries that were drawn after the study window. Specifically, we apply the same estimating equation (1) to voters living near boundaries created between 2023 and 2026, examining whether assignment to the high-exposure side of these future boundaries predicts changes in political outcomes between 2012 and 2020—a period during which these boundaries did not yet exist. If stake presidents tend to draw boundary lines along pre-existing partisan fault lines—places where political trends were already diverging—then future boundaries should also appear to "predict" political change in the preceding period. A null result, by contrast, would indicate that boundary placement is unrelated to differential trends in political outcomes, directly supporting our identifying assumption.[29]

[29] The placebo also serves as a specification test. If the functional form of the geographic running variable were misspecified in a way that generates spurious treatment effects, we would expect to see significant placebo coefficients as well.

Table 1 reports the placebo estimates for 2012 Republicans, Unaffiliated voters, and Democrats. No coefficient is statistically significant at the 5% level, and the joint tests that all coefficients in a block are zero have p-values between 0.16 and 0.99. Two features of the table deserve comment. First, the placebo sample is about a quarter smaller than the main sample, so the placebo is more informative for some cells than for others. For the effect of Republican exposure on Unaffiliated voters' Republican registration, our most precisely estimated main effect, the placebo estimate is a precise zero. For the effects of Democratic exposure on Republicans, the placebo estimates are imprecise, with confidence intervals that include zero and also include effects of the size we estimate in Section 6; the placebo therefore bounds rather than rules out confounding trends in those cells. Second, the two cells that reach significance at the 10% level all involve Republican exposure among Republicans, and each is opposite in sign to the corresponding main estimate, so they cannot be producing our results. Overall, the placebo results are consistent with the assumption that new ward boundary lines are placed in a way that is uncorrelated with trends in partisanship.

5.2 Simulated Boundary Exercise

The placebo test above directly tests our identifying assumption by asking whether boundary placement predicts differential political trends. A complementary question is whether the redistricting process itself tends to sort voters by partisanship—even inadvertently—by balancing wards on dimensions correlated with party affiliation. As discussed in Section 2.3, stake presidents report that their primary goal is to equalize attendance across wards, and that when they account for individual characteristics at all, it is to balance rather than segregate groups. If this rebalancing motive extends to characteristics correlated with partisanship, we would expect post-redistricting wards to be more similar in partisan composition than pre-redistricting wards, relative to a benchmark of feasible alternative configurations.

To test this, we conduct a simulation exercise analogous to those used in the gerrymandering literature, using graph partition algorithms to draw counterfactual wards sampled from the target distribution of potential ward assignment definitions (DeFord, Duchin and Solomon, 2021; McCartan and Imai, 2023). For each year 2013–2017, we simulate alternative ward boundaries within each stake, since wards are typically redrawn such that their original stake is preserved. We do this using the Merge-Split algorithm (Carter et al., 2019) implemented through `redist` software (Kenny et al., 2022), redrawing wards using Census blocks. The Merge-Split algorithm generates alternative districting or aggregation plans by repeatedly merging adjacent geographic units into larger regions and then probabilistically splitting them back apart while preserving constraints like population balance and contigu-

Table 1: Placebo Test: Effect of Peer Partisan Exposure Using Future Boundaries (2023–2026)

	Democratic-share treatment			Republican-share treatment			Both treatments		
	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem
Panel A: 2012 Republicans									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	−0.122	0.003	0.119				−0.073	−0.043	0.117
	(0.176)	(0.148)	(0.074)				(0.181)	(0.155)	(0.078)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				−0.112	0.131*	−0.019	−0.130*	0.121	0.009
				(0.073)	(0.072)	(0.025)	(0.076)	(0.075)	(0.026)
Joint test p-value		0.273			0.160			0.223	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.333	0.678	0.148
Dep. Var. Mean	-0.146	0.124	0.022	-0.146	0.124	0.022	-0.146	0.124	0.022
# Voters	44,324	44,324	44,324	44,324	44,324	44,324	44,324	44,324	44,324
# Boundaries	1,029	1,029	1,029	1,029	1,029	1,029	1,029	1,029	1,029
Panel B: 2012 Unaffiliated voters									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	0.189	−0.087	−0.102				0.216	−0.074	−0.142
	(0.198)	(0.253)	(0.160)				(0.218)	(0.274)	(0.171)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.022	0.037	−0.059	0.035	0.001	−0.037
				(0.094)	(0.101)	(0.069)	(0.105)	(0.107)	(0.073)
Joint test p-value		0.524			0.688			0.784	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.360	0.821	0.371
Dep. Var. Mean	0.263	-0.397	0.134	0.263	-0.397	0.134	0.263	-0.397	0.134
# Voters	46,835	46,835	46,835	46,835	46,835	46,835	46,835	46,835	46,835
# Boundaries	1,044	1,044	1,044	1,044	1,044	1,044	1,044	1,044	1,044
Panel C: 2012 Democrats									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	−0.509	0.316	0.193				−0.617	0.551	0.066
	(0.425)	(0.342)	(0.441)				(0.431)	(0.384)	(0.477)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.009	0.023	−0.031	−0.141	0.259	−0.118
				(0.173)	(0.169)	(0.220)	(0.173)	(0.190)	(0.242)
Joint test p-value		0.423			0.988			0.370	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.126	0.100	0.931
Dep. Var. Mean	0.118	0.159	-0.276	0.118	0.159	-0.276	0.118	0.159	-0.276
# Voters	9,986	9,986	9,986	9,986	9,986	9,986	9,986	9,986	9,986
# Boundaries	711	711	711	711	711	711	711	711	711
Boundary FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Distance Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Specification, outcome definition, clustering, and joint test as in Table 3.

ity. The results is a distribution of alternative plans that on average capture the partisan segregation across wards we expect to observe due to the geographic distribution of partisans across Utah. This distribution of alternative ward configurations allows us to compare the actual ward map to a benchmark of what kind of partisan distribution across wards we would observe due to partisan geography alone, absent any rebalancing by stake presidents. This is akin to identifying a gerrymandered district plan by simulating potential plans, calculating partisan balance statistics and comparing the distribution of these statistics in simulated plans to the observed plans. In our case, we are testing whether wards are (intentionally or unintentionally) gerrymandered.

Following best practice in the algorithmic gerrymandering literature, for each year-stake, we run four simulation chains for 50,000–400,000 simulations each[30]. We discard the first one-fifth of simulations as warmup and store 4,000 systematically thinned draws in total, 1,000 from each chain. With these data, we compare the partisan balance across actual wards to the partisan balance in the counterfactual distribution.

Figure 4 plots the distribution of across-ward standard deviation in partisan composition (proportion Republican and proportion Democrat) for all 4,000 sets of simulated wards for each year 2013–2017.[31] We compare this distribution to the standard deviation of partisan composition across actual wards in each year. We find that, compared to the counterfactual distribution, wards are drawn such that they are more balanced than what we would expect from partisan geography alone. For each year and for both proportion Democrat and proportion Republican, the standard deviation for the observed ward map falls in the lower tail of the distribution. However, the magnitude of the difference between observed wards and the simulated distribution is small, with differences between the observed standard deviation and the simulation average ranging from -0.0026 to -0.0035 for proportion Democrat and -0.0034 to -0.0046 for proportion Republican. Importantly, these differences increase only slightly across years, indicating that the redrawing of boundaries occurring during our time period of study is not dramatically altering the extent to which observed maps are balancing partisans across wards.

This finding is consistent with the rebalancing motive described by stake presidents and suggests that redistricting compresses rather than amplifies cross-ward partisan differences. It is important to note that this exercise tests for rebalancing in levels of partisan composition across wards, not for correlation with partisan trends locally around the boundary—the latter is what our identifying assumption requires and is tested directly by the placebo exercise above. This exercise complements our placebo test by showing that the attendance and demographic rebalancing motive of stake presidents may also translate into an implicit rebalancing on partisanship. To the extent that wards are rebalanced on characteristics correlated with partisan trends, the resulting bias works against finding peer effects of either

[30]We increase iterations as need in wards where initial convergence statistics are not satisfactory at lower numbers of iterations. We assess convergence using split R-hat diagnostics (Gelman and Rubin, 1992) that compare variation of map characteristics (e.g., population, racial demographics, compactness) within and across simulation chains to verify that chains converge on stable values of these variables. Appendix Table F2 contains summary of convergence statistics.

[31]The distribution of the standard deviation across simulated plans is narrow, reflective of the limited capacity to draw wards with dramatically different partisan compositions on average given current geographic distributions of partisans in Utah, as well as the relatively narrow set of redrawing possibilities for groups of wards within stakes, which are themselves fairly small geographies. The simulated distributions of Democratic share are bimodal due to related constraints in feasible maps.

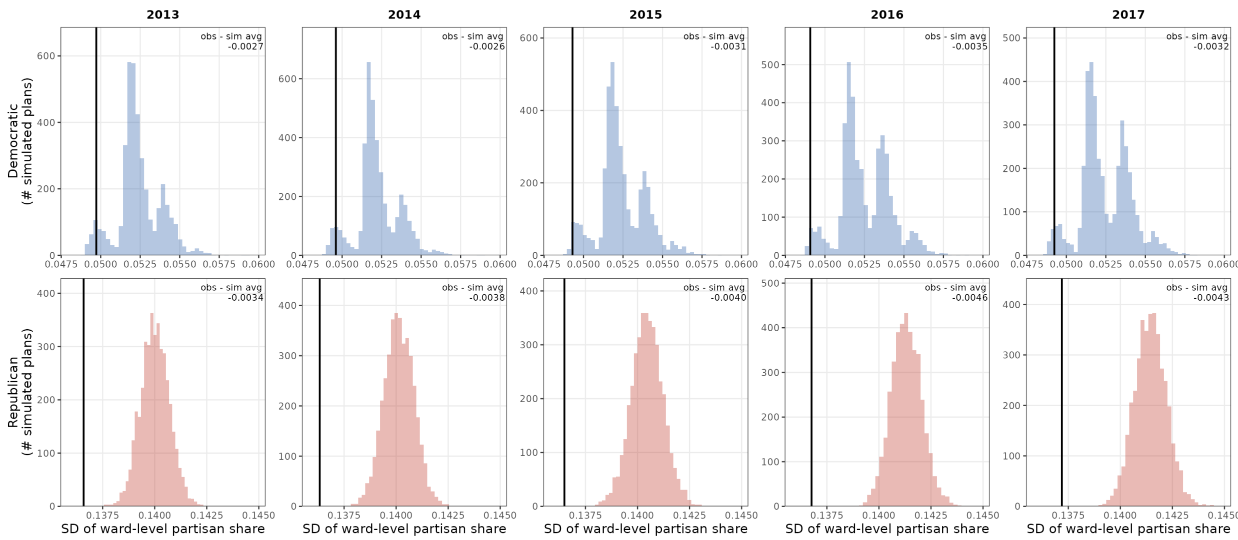

Figure 4: Distribution of partisan balance across counterfactual wards, 2013–2017. Histograms show the distribution of the standard deviation of proportion Democrat (top row) and proportion Republican (bottom row) across simulated ward partitions for each set of wards 2013–2017. The vertical black line marks the standard deviation for the observed ward map. The difference between the observed standard deviation and the simulation average is shown in the top right corner.

sign as it would attenuate estimates toward zero. Our results are therefore, if anything, conservative in magnitude relative to the true peer effect.

5.3 Cross-Boundary Balance in Levels

Our design assumes parallel counterfactual trends in political outcomes between the two sides of each boundary. Unfortunately, our data do not allow us to implement a conventional pre-trend test since our panel begins in 2012, boundaries are drawn in 2013–2017, and political outcomes vary mostly across presidential elections. The only candidate pre-period observation, 2016, is already post-treatment for the 2013–2015 cohorts and, because each map vintage reflects a change occurring at any point since the prior year's map (Section 3.2), not clearly pre-treatment even for the later ones.

What we can do is characterize the sorting in levels that new boundaries pick up. Table 2 presents results mirroring our main regression specifications, but replacing first-differenced political outcomes with baseline demographic, property, and turnout characteristics. Results are shown for 2012 Republicans and Unaffiliated voters. To make magnitudes comparable across characteristics, cells report the coefficient in standard deviations of the characteristic per 0-to-1 change in peer share. Multiplying by the average realized cross-boundary contrast (0.059 for Republican exposure, 0.026 for Democratic exposure) converts these to realized differences.

We find that new boundaries coincide with discontinuities in several baseline variables. Property values, home sizes, and construction dates are higher on the more-Republican side of a line and lower on the more-Democratic side, consistent with lines following the edges of subdivisions and other breaks in the housing stock, and prior turnout is higher on the more-Republican side among Republicans and lower on the more-Democratic side. Scaled by the average realized contrast, the housing differences are on the order of a tenth of a standard deviation. Because these are differences in levels, the first difference absorbs them, and they threaten identification only if they also predict 2012–2020 trends in partisanship. We test that directly by re-estimating our specifications with the following controls: demographics (age, sex, race), property characteristics (assessed value, square footage, year built), and school-catchment fixed effects, all of which absorb correlations between baseline characteristics and subsequent partisan trends in the first-differenced specification. In the first-differenced equation these controls enter additively, allowing the trend in registration to differ by each characteristic. Estimates are stable across specifications with and without these controls (Section 6.3).

Table 2: Cross-Boundary Differences in Baseline Characteristics (Levels)

	Republican exposure		Democratic exposure	
	2012 Rep.	2012 Unaff.	2012 Rep.	2012 Unaff.
Age (years)	0.171	−0.027	−0.515	−0.995**
	(0.208)	(0.228)	(0.454)	(0.460)
Male	−0.099	0.155	0.287	−0.694**
	(0.138)	(0.153)	(0.281)	(0.340)
White	0.141	0.434**	0.140	−0.512
	(0.223)	(0.215)	(0.411)	(0.513)
Homeowner (CoreLogic match)	−0.292	0.097	−0.652*	−0.744*
	(0.196)	(0.214)	(0.386)	(0.439)
Log property value	0.837***	1.292***	−2.776***	−2.434***
	(0.308)	(0.350)	(0.685)	(0.705)
Log building sqft	0.819***	0.609*	−3.722***	−2.438***
	(0.311)	(0.319)	(0.753)	(0.729)
Year home built	0.881***	1.174***	−0.763	−0.454
	(0.309)	(0.348)	(0.734)	(0.895)
Household income (HMDA)	−0.749	0.433	−5.026*	−3.321
	(1.167)	(1.497)	(2.563)	(2.529)
Rural	0.047	0.029	−0.266	−0.360
	(0.284)	(0.236)	(0.546)	(0.557)
Voted 2010 general	0.163	−0.006	−1.262***	−0.712*
	(0.202)	(0.188)	(0.399)	(0.413)
Voted 2012 general	0.363**	−0.246	−0.712*	−0.115
	(0.180)	(0.186)	(0.376)	(0.408)
Voted 2012 primary	0.563***	−0.048	−0.963**	−0.337
	(0.167)	(0.167)	(0.378)	(0.391)

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each cell is a separate regression of the 2012 level of the row characteristic on the continuous peer exposure, with boundary fixed effects and the main specification's distance controls (a quadratic in distance, with each term also interacted with the high-exposure side), for the indicated 2012-party subsample at the 400m bandwidth. Coefficients are standardized as standard deviations of the characteristic per 0-to-1 change in peer share; multiply by the average cross-boundary contrast (0.059 Republican, 0.026 Democratic exposure) for realized differences. Standard errors two-way clustered by boundary side and voter. Because outcomes in the main design are first-differenced, level balance is not required for identification.

6 Main Results

We now present our main estimates of the effect of ward peer composition on party affiliation. All specifications use the preferred bandwidth of 400 meters and include boundary fixed effects and quadratic distance controls, but omit individual-level controls. We show robustness to the inclusion of controls and to alternative bandwidths in Section 6.3.

Table 3 reports the effects of peer partisan exposure on party registration, with a panel for each 2012 baseline party: Republicans, Unaffiliated voters, and Democrats. Each panel reports the effect of Democratic peer exposure, of Republican peer exposure, and of both included jointly. We focus our discussion on Republicans and Unaffiliated voters for two related reasons. First, they make up a much larger proportion of the Utah electorate than Democrats, and thus estimates for the Republican and Unaffiliated subsamples are far more precise than those for Democrats. Second, they are much more likely to be LDS than Democrats, making them less subject to the attenuation induced by unobserved LDS status (quantified in Section 6.1).

The results for party affiliation are consistent with conformity. Across the table, each statistically significant coefficient is signed as conformity predicts—individuals exposed to more members of a party are more likely to switch to or stay with that party—and most point estimates follow suit. Throughout, we quote coefficients as the effect of a 100 percentage-point change in peer share (Section 6.1 rescales them to realized cross-boundary differences). For 2012 Republicans (Panel A), Democratic peer exposure lowers the probability of remaining Republican by 32.2 percentage points ($p = 0.016$) and raises the probability of switching to Democrat by 16.1 percentage points ($p = 0.001$); the probability of switching to Unaffiliated also rises, though not significantly ($p = 0.157$). For 2012 Unaffiliated voters (Panel B), Republican peer exposure raises the probability of registering Republican by 22.8 percentage points ($p = 0.004$) and lowers the probability of registering Democrat by 10.8 percentage points ($p = 0.029$). The remaining coefficients are smaller and mostly statistically insignificant.

The results suggest an asymmetry between own-party and out-party exposure, best seen by focusing on the "Both treatments" specification. Because Unaffiliated voters make up nearly half of the Utah electorate, exposure to Republicans and exposure to Democrats are distinct margins rather than inverses of one another (Appendix Figure C3 quantifies these trade-offs). For baseline Republicans, we see that exposure to Democrats, holding the Republican share constant, leads to a 29.1 percentage point reduction in their likelihood of remaining Republican with a 14.8 percentage point increase in their likelihood of switching to Democrat. Notice that because the Republican share is held constant, the Democratic share rises at the expense of the Unaffiliated share. In contrast, exposure to Republicans, holding the Democratic share constant, has much smaller and largely statistically insignificant effects. One interpretation is that co-partisans do little to reinforce an existing affiliation, while exposure to the other party does move voters. The evidence for Unaffiliated voters is less

Table 3: Effect of Peer Partisan Exposure on Party Registration, by 2012 Baseline Party

	Democratic-share treatment			Republican-share treatment			Both treatments		
	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem
Panel A: 2012 Republicans									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	-0.322^{**}	0.161	0.161^{***}				-0.291^{**}	0.143	0.148^{***}
	(0.133)	(0.114)	(0.048)				(0.139)	(0.119)	(0.050)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.079	-0.033	-0.047^{**}	0.044	-0.011	-0.034^{*}
				(0.057)	(0.053)	(0.020)	(0.058)	(0.055)	(0.019)
Joint test p-value		0.003			0.046			0.014	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.128	0.353	0.036
Dep. Var. Mean	-0.139	0.121	0.018	-0.139	0.121	0.018	-0.139	0.121	0.018
# Voters	60,639	60,639	60,639	60,639	60,639	60,639	60,639	60,639	60,639
# Boundaries	1,428	1,428	1,428	1,428	1,428	1,428	1,428	1,428	1,428
Panel B: 2012 Unaffiliated voters									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	-0.182	0.117	0.064				-0.029	0.051	-0.023
	(0.176)	(0.179)	(0.117)				(0.185)	(0.195)	(0.123)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.228^{***}	-0.120	-0.108^{**}	0.213^{**}	-0.121	-0.092^{*}
				(0.080)	(0.079)	(0.050)	(0.085)	(0.086)	(0.053)
Joint test p-value		0.568			0.008			0.088	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.414	0.770	0.442
Dep. Var. Mean	0.295	-0.401	0.106	0.295	-0.401	0.106	0.295	-0.401	0.106
# Voters	49,600	49,600	49,600	49,600	49,600	49,600	49,600	49,600	49,600
# Boundaries	1,440	1,440	1,440	1,440	1,440	1,440	1,440	1,440	1,440
Panel C: 2012 Democrats									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	0.225	-0.082	-0.143				0.231	-0.401	0.169
	(0.323)	(0.329)	(0.371)				(0.347)	(0.367)	(0.418)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				-0.026	-0.151	0.176	0.004	-0.302^{*}	0.298
				(0.128)	(0.127)	(0.158)	(0.139)	(0.156)	(0.185)
Joint test p-value		0.785			0.439			0.353	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.569	0.129	0.379
Dep. Var. Mean	0.129	0.165	-0.294	0.129	0.165	-0.294	0.129	0.165	-0.294
# Voters	8,308	8,308	8,308	8,308	8,308	8,308	8,308	8,308	8,308
# Boundaries	875	875	875	875	875	875	875	875	875
Boundary FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Distance Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each column is a separate estimation of Equation (1) on the indicated 2012-party subsample at the preferred 400-meter bandwidth, with boundary fixed effects, a quadratic in signed distance interacted with the high-exposure side, and no individual-level controls. Outcomes are first differences in a registration indicator between 2012 and 2020/2021, so a coefficient is the change in the probability of the indicated 2020 registration per 0-to-1 change in assigned-ward peer share; Table 4 rescales these to the average realized cross-boundary difference. $D_{ib}^{X} \times I_b^{X}$ is the treatment of Equation (1) for exposure X: the cross-boundary difference in baseline party-X peer share, switched on for voters on the higher-exposure side. The first two blocks enter each exposure alone; the "Both treatments" block enters them jointly as in Equation (2). All three blocks are estimated on the same sample, the voters near boundaries with a nonzero cross-boundary contrast in both Republican and Democratic share. Standard errors, in parentheses, are two-way clustered by ward-boundary side and voter. "Joint test p-value" is from a Wald test that all $D \times I$ coefficients in the block are zero, estimated by stacking the ΔRep and ΔDem regressions (ΔUnaff is their negative within voter) with the same clustering, so that the test accounts for the correlation between columns. "p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$" tests, for each outcome in the "Both treatments" block, that a unit of Democratic exposure moves the outcome by the same amount as a unit of Republican exposure moves it in the opposite direction.

clean, but still broadly consistent with this story. While both major parties are out-groups for an Unaffiliated voter, as discussed in Section 2.4, Unaffiliated voters in Utah tend to vote more similarly to Democrats. Thus, it is plausible that Republicans are a more salient out-group from their perspective. In this case, our finding that Republican exposure matters much more than Democratic exposure for Unaffiliated voters would support a story of out-partisans mattering more in shaping partisan choices. The caveat is that Democrats are the group least likely to be LDS, so Democratic exposure suffers most from measurement-error attenuation, and its estimates for this subsample are also the least precise; the insignificant cells could therefore reflect attenuation and imprecision rather than truly null effects.

6.1 Contextualizing Effect Size

The raw coefficients in Table 3 are scaled per 100 percentage-point change in peer share, while the actual cross-boundary differences are much smaller. The median difference in Republican share is about 5.2 percentage points (Figure 3). To gauge the magnitude that matters in practice, we rescale the estimates by the average cross-boundary difference in peer composition:

$$\tilde{\beta} = \hat{\beta} \cdot \frac{\sum_b N_b \cdot |p_b - q_b|}{\sum_b N_b},$$

where N_b is the number of voters near boundary b, and p_b and q_b are the ward-level peer composition means on each side.

Table 4 reports these rescaled effects for every party-registration margin in the 2012 Republican and Unaffiliated samples, expressed as the average percentage-point difference in outcomes between voters on the high- and low-exposure sides of a boundary.

Because our treatment variation is modest, our rescaled estimates are likewise also modest. For 2012 Republicans, assignment to the more Democratic side of a typical boundary lowers Republican retention by about 0.8 percentage points and raises the probability of switching to Democrat by 0.4 percentage points. For 2012 Unaffiliated voters, assignment to the more Republican side of a typical boundary raises the probability of registering Republican by 1.4 percentage points and lowers the probability of registering Democrat by 0.6 percentage points.

How do these magnitudes compare to the literature? We place our estimates and those of similar studies on a common scale of percentage-point change in an individual's partisan outcome per 10 percentage-point change in peer partisan composition. On this scale, a 10 percentage-point increase in assigned-peer Republican share raises a 2012 Unaffiliated voter's probability of registering Republican by 2.3 percentage points over eight years, and a

Table 4: Rescaled Effects: Average Impact of Ward Assignment

	$\hat{\beta}$	Avg. $\lvert p_b - q_b \rvert$	Rescaled effect	LDS adj. mult.	LDS-adjusted effect	Mean ΔY
2012 Rep. → 2020 Rep. (Rep. exposure)	0.079	0.059	0.5 pp	1.28	0.6 pp	-13.9 pp
2012 Rep. → 2020 Unaff. (Rep. exposure)	-0.033	0.059	-0.2 pp	1.28	-0.2 pp	12.1 pp
2012 Rep. → 2020 Dem. (Rep. exposure)	-0.047^{**}	0.059	-0.3^{**} pp	1.28	-0.4^{**} pp	1.8 pp
2012 Rep. → 2020 Rep. (Dem. exposure)	-0.322^{**}	0.026	-0.8^{**} pp	2.55	-2.2^{**} pp	-13.9 pp
2012 Rep. → 2020 Unaff. (Dem. exposure)	0.161	0.026	0.4 pp	2.55	1.1 pp	12.1 pp
2012 Rep. → 2020 Dem. (Dem. exposure)	0.161^{***}	0.026	0.4^{***} pp	2.55	1.1^{***} pp	1.8 pp
2012 Unaff. → 2020 Rep. (Rep. exposure)	0.228^{***}	0.059	1.4^{***} pp	2.17	2.9^{***} pp	29.5 pp
2012 Unaff. → 2020 Unaff. (Rep. exposure)	-0.120	0.059	-0.7 pp	2.17	-1.6 pp	-40.1 pp
2012 Unaff. → 2020 Dem. (Rep. exposure)	-0.108^{**}	0.059	-0.6^{**} pp	2.17	-1.4^{**} pp	10.6 pp
2012 Unaff. → 2020 Rep. (Dem. exposure)	-0.182	0.026	-0.5 pp	4.34	-2.1 pp	29.5 pp
2012 Unaff. → 2020 Unaff. (Dem. exposure)	0.117	0.026	0.3 pp	4.34	1.3 pp	-40.1 pp
2012 Unaff. → 2020 Dem. (Dem. exposure)	0.064	0.026	0.2 pp	4.34	0.7 pp	10.6 pp

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Rescaled effects are $\hat{\beta}\times$ the voter-weighted average absolute cross-boundary difference in peer composition, the average realized change in exposure for a voter assigned to the high-exposure side. "LDS adj. mult." is the adjustment multiplier $1/k$ from Appendix E for the row's baseline party and exposure, which undoes the attenuation from treating non-LDS voters as ward members, and "LDS-adjusted effect" multiplies the rescaled effect by it to recover the implied effect on LDS members. Because each rescaled and adjusted effect is $\hat{\beta}$ times a constant, statistical significance remains the same. "Mean ΔY" is the average first-differenced outcome in the subsample, in percentage points (positive = net adoption into the outcome category over 2012–2020; negative = net loss).

10 percentage-point increase in Democratic peer share raises a 2012 Republican's probability of registering Democrat by 1.6 percentage points. The closest benchmarks are larger, but of the same order of magnitude. Cantoni and Pons (2022) estimate that adults who move to a state with a 10 percentage-point higher Republican-affiliation share become 2.8 percentage points more likely to affiliate Republican at their first post-move election, and Brown et al. (2023) estimate that a childhood spent in a 10 percentage-point more Republican county raises Republican registration at first eligibility by 4.7 percentage points (about 0.52 points per year of teenage exposure). Both are context effects that bundle peers with institutions, media markets, and place, while our estimate isolates the peer-composition component of such bundles and recovers a substantial share of their magnitude. A second comparison uses the persuasion rate of DellaVigna and Gentzkow (2010), which expresses an effect as the share of those exposed who change behavior among those who would not have done so otherwise. Here the assigned ward's peer share is the exposed share, so the rate is the coefficient divided by one minus the baseline transition rate: 0.32 for Unaffiliated voters' Republican registration and 0.37 for Republicans leaving their party under Democratic exposure. Compared with typical persuasion rates for media exposure or campaign contact, our estimates of church peer exposure over an eight-year horizon are, as might be expected, substantially higher.

Our rescaled effects are conservative estimates of the underlying causal effect on LDS members. The regression cannot distinguish LDS from non-LDS voters, treating every Utah voter both as a subject of ward peer influence and as a contributor to the leave-one-out peer composition, so the estimates blend the response of LDS members with the null response of their non-LDS neighbors and are thus attenuated relative to the true LDS-only effect. Appendix E derives the implied attenuation in closed form, validates it via Monte Carlo simulation, and verifies it against the actual Utah ward-boundary configurations in our analysis sample. Under the following survey-estimated assumptions about probability LDS by party—$p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, $p_{\text{LDS}}(D) = 0.25$ (Appendix E.1)—the implied adjustment multipliers range from roughly $1.3\times$ for baseline Republicans exposed to Republican wards to $8.7\times$ for baseline Democrats exposed to Democratic wards. Applied to Table 4, these multipliers would scale every reported effect upward, with the smallest correction ($\approx 1.3\times$) applying to the Republican-baseline cells.

6.2 Assessing the Linear-in-Means Approximation

Equation (1) is a linear-in-means specification—an individual's change in party registration is linear in the (leave-one-out) partisan share of their ward. Linear-in-means is the workhorse model of the empirical peer effects literature (Blume et al., 2015; Manski, 1993; Sacerdote, 2011). Linear-in-means involves two distinct assumptions, which we consider in turn: that the group mean is the feature of peer composition that matters, and that the individual response is linear in that mean.

The first assumption, mean sufficiency, is restrictive in many applications. In the economics of education, for example, peer effects appear to load on features of the ability distribution beyond its mean, where disruptive "bad apple" peers near the bottom and exceptional peers near the top can matter more than the average (Hoxby and Weingarth, 2005; Sacerdote, 2011). These concerns have limited force in our setting, however, for the simple fact that our peer attribute is binary. The within-ward distribution of a binary characteristic is Bernoulli, so the partisan share is a sufficient statistic for the entire composition distribution.

The second assumption, linearity in the share, is more substantive. It requires that the marginal effect of peer composition be constant, implying that moving a voter from a 10% to a 20% Republican ward should shift behavior by the same amount as moving them from a 70% to an 80% Republican ward, for example. A number of models predict otherwise, each implying a different departure: threshold and tipping models imply a dose-response curve that is flat until a critical mass and steep thereafter (Brock and Durlauf, 2001; Card, Mas and Rothstein, 2008; Granovetter, 1978; Schelling, 1971); complex contagion implies a superadditive response, since switching requires reinforcement from several co-partisan ties (Centola and Macy, 2007); bounded-confidence models imply that peers whose views are too distant exert no pull at all (Hegselmann and Krause, 2002); and backlash models allow out-party exposure to deepen rather than soften partisanship, reversing the sign (Bail et al., 2018). Whether the linear form is a good approximation is therefore an empirical question specific to the setting and the outcome.

We show that in our setting, over our observed support, the peer-response function appears to be approximately linear-in-means. We bin cross-boundary partisan exposure, our measure of treatment intensity, into quartiles and replace I_b with bin-specific dummies to estimate the effects of being assigned to the high-exposure side within each bin. Figure 5 plots our results, showing that the binned estimates sit on the linear fit throughout the observed dose range, with no detectable curvature or flat-then-steep threshold pattern, suggesting that the linear specification is a good approximation. We also test whether effects vary by origin ward's baseline composition and find no evidence for heterogeneity (Appendix F.3), which provides further support for our linear assumption.

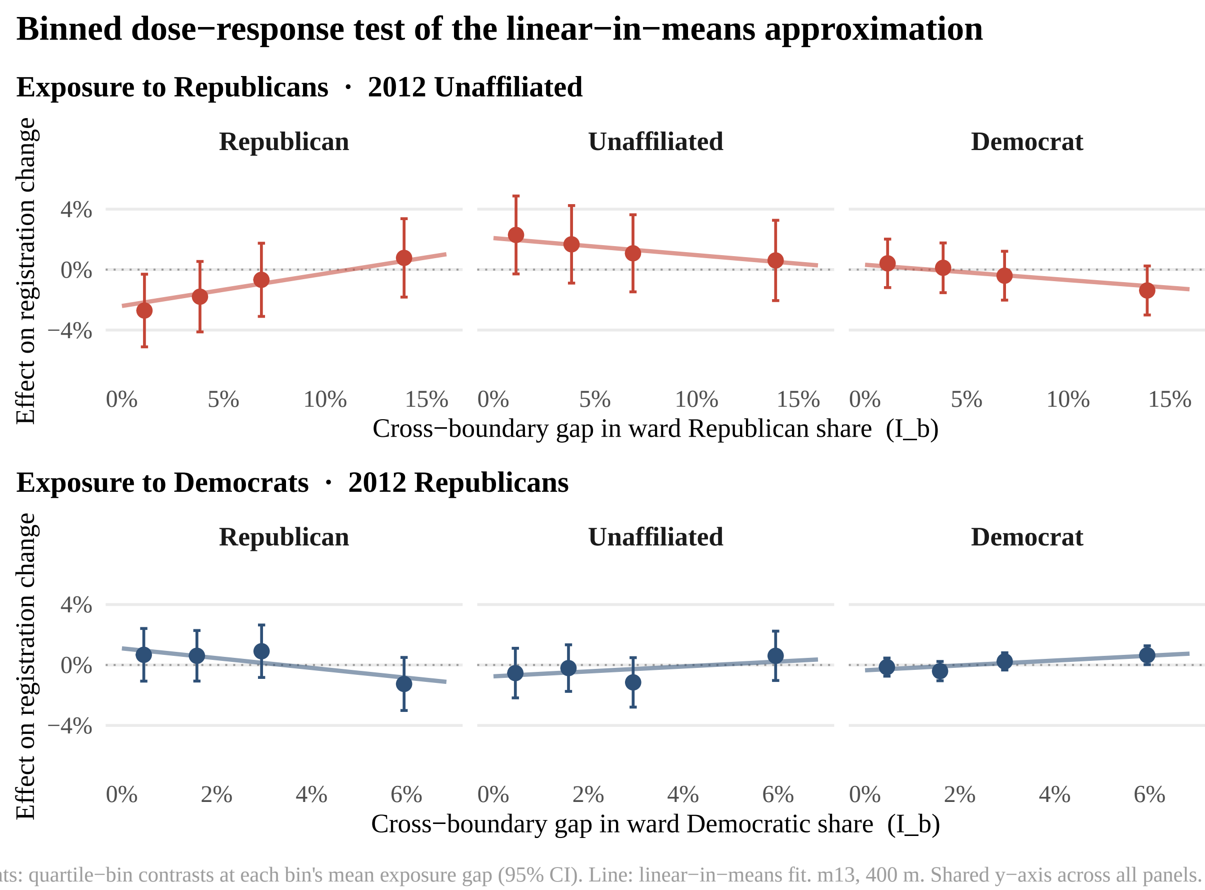

Figure 5: Binned Dose-Response Test of the Linear-in-Means Approximation. Points show quartile-bin contrasts of the cross-boundary exposure gap (γ_k from a binned version of Equation (1)) plotted at each bin's mean exposure gap, with 95% confidence intervals; the line shows the linear fit. Top row: Republican exposure among 2012 Unaffiliated voters. Bottom row: Democratic exposure among 2012 Republicans. The remaining exposure-sample combinations appear in Appendix F.

6.3 Robustness

Our preferred specification includes boundary fixed effects and quadratic distance controls but omits individual-level controls. Appendix Figure F3 shows that the partisanship estimates are stable across specifications that progressively add controls for demographics, property characteristics, and school-catchment fixed effects. Appendix Figure F4 shows that the six main estimates (the five significant effects in Table 4 plus Republicans' switching to Unaffiliated) are stable across bandwidths from 50 to 600 meters, and Appendix Figure F5 shows they are likewise stable under a boundary donut, linear distance controls, dropping the most extreme boundary, and restricting to once-changed boundaries.

Standard errors for our main estimates are two-way clustered by boundary side and voter. In Appendix Table F3 we present p-values using alternative standard error definitions. Clustering by boundary pair instead, which lets errors correlate across a line, and Conley spatial standard errors at cutoffs from 1 to 10 kilometers widen the intervals but do not change the picture: Republicans' switching to Democrat under Democratic exposure is significant at the 1% level under every method, and the two largest effects, Unaffiliated voters' Republican registration under Republican exposure and Republicans' retention under Democratic exposure, remain significant at the 5% level under most alternatives and at the 10% level under all of them. The two smaller effects, Unaffiliated voters' Democratic registration and Republicans' Democratic registration under Republican exposure, are less robust, with p-values between 0.04 and 0.17 across methods. We also report p-values from randomization inference, which uses the design itself, rather than a large-sample approximation, to judge how unusual our estimates are. Across 1,000 draws we randomly relabel which side of each boundary is the high-exposure side, which is the assignment our design treats as as-good-as-random, holding geography, exposure contrasts, and outcomes fixed, and re-estimate the main specification. If peer composition does not matter, the relabeling is inconsequential, so the draws trace out the distribution of the t-statistic under that null, and we judge the observed t-statistic against it. The permutation distribution is wider than the standard normal, so these p-values are larger than the clustered ones: $p = 0.008$ for Republicans' switching to Democrat, 0.030 for Unaffiliated voters' Republican registration, 0.059 for Republicans' retention, and between 0.064 and 0.278 for the remaining estimates (Appendix F.7).

6.4 Attrition

Because our registration estimates condition on observing a voter in both 2012 and 2020, selective attrition threatens their interpretation. Voters disappear from our panel for two reasons: either they choose to become unregistered, or we fail to link them to their 2020 outcomes, often because of a move. Both are possible responses to changes in peer partisanship. It is possible that out-party exposure induces the least conformist voters to unregister or move away, which would cause the stayers we observe to over-represent conformists, thereby inflating our estimates.

We test whether partisan exposure predicts sample attrition, estimating the probability that a voter disappears from the panel entirely on the full 2012 baseline sample, and find that exposure matters for one subsample. A unit increase in Republican exposure lowers the probability that a 2012 Unaffiliated voter disappears by 16.3 percentage points ($p = 0.011$), and Democratic exposure raises it by 23.0 percentage points, though only marginally significantly ($p = 0.089$). Neither exposure predicts attrition among 2012 Republicans (-9.1 percentage points, $p = 0.510$, for Democratic exposure; -4.4 percentage points, $p = 0.497$, for Republican exposure).

We address this threat by bounding the registration estimates using the trimming approach of Lee (2009). The idea is to make the two sides of a boundary comparable again by discarding, from the side that retained more of its voters, the excess retained voters with the most extreme outcome changes. On each boundary's higher-retention side we drop a share of retained voters equal to that boundary's retention differential, trimming from the top of the outcome distribution for the lower bound and from the bottom for the upper bound, and re-estimate Equation (1) on what remains. Table 5 reports the bounds for the twelve registration estimates of Section 6 for 2012 Republicans and Unaffiliated voters. The bounds are tight because differential attrition is small. In every exposure–sample block the retention differential between the high- and low-exposure sides is under half a percentage point, compared with a retention rate near 70%. Every statistically significant estimate keeps its sign at both bounds, and all but one remain significant at the 5% level: the lower bound on Unaffiliated voters' Republican registration (0.166) is significant only at the 10% level. Selective attrition cannot account for the sign or the bulk of the magnitude of our results. Because most panel loss is associated with a move, Appendix F.8 additionally re-estimates the main specification separately for voters who did and did not move after baseline; the two sets of estimates are statistically indistinguishable.

Table 5: Lee Bounds on Registration Effects Under Selective Attrition

	Point	Lower bound	Upper bound	Ret. diff.	Trim share
2012 Rep. → 2020 Rep. (Rep. exposure)	0.079	0.074	0.103	−0.005	0.043
	(0.057)	(0.060)	(0.064)		
2012 Rep. → 2020 Unaff. (Rep. exposure)	−0.033	−0.052	−0.025		
	(0.053)	(0.059)	(0.057)		
2012 Rep. → 2020 Dem. (Rep. exposure)	−0.047**	−0.049**	−0.047**		
	(0.020)	(0.020)	(0.019)		
2012 Rep. → 2020 Rep. (Dem. exposure)	−0.322**	−0.330**	−0.279**	−0.005	0.051
	(0.133)	(0.141)	(0.135)		
2012 Rep. → 2020 Unaff. (Dem. exposure)	0.161	0.122	0.158		
	(0.114)	(0.114)	(0.119)		
2012 Rep. → 2020 Dem. (Dem. exposure)	0.161***	0.129***	0.171***		
	(0.048)	(0.045)	(0.051)		
2012 Unaff. → 2020 Rep. (Rep. exposure)	0.228***	0.166*	0.329***	+0.001	0.057
	(0.080)	(0.100)	(0.082)		
2012 Unaff. → 2020 Unaff. (Rep. exposure)	−0.120	−0.217**	−0.044		
	(0.079)	(0.085)	(0.097)		
2012 Unaff. → 2020 Dem. (Rep. exposure)	−0.108**	−0.112**	−0.108**		
	(0.050)	(0.054)	(0.050)		
2012 Unaff. → 2020 Rep. (Dem. exposure)	−0.182	−0.253	−0.094	−0.003	0.052
	(0.176)	(0.188)	(0.198)		
2012 Unaff. → 2020 Unaff. (Dem. exposure)	0.117	−0.021	0.202		
	(0.179)	(0.205)	(0.196)		
2012 Unaff. → 2020 Dem. (Dem. exposure)	0.064	0.053	0.154		
	(0.117)	(0.123)	(0.125)		

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Trimming bounds following Lee (2009), with trimming fractions computed within boundary: on each boundary's higher-retention side, the retained first-differenced outcome distribution is trimmed from below (upper bound) or above (lower bound) by that boundary's retention differential. "Ret. diff." is the retention differential between high- and low-exposure sides and "Trim share" is the average share of the retained sample trimmed; both are constant within a subsample × exposure block and reported on its first row. Standard errors two-way clustered by boundary side and voter.

7 Mechanisms

This section asks what channels underlie the effects documented in Section 6. It is useful to begin by noting what does and does not vary across wards. The LDS Church is highly centralized: doctrine, lesson manuals, and organizational policies are set by church headquarters in Salt Lake City, and meetinghouse buildings are standardized. LDS congregations are also institutionally apolitical, in that LDS members are far less likely than members of other major U.S. religious traditions to hear politics discussed from the pulpit or to see their congregation organize voter mobilization (Appendix Figures A2a and A2b). What varies across wards is the membership itself, and with it the social interactions the ward generates.

We highlight two channels through which ward composition may shape political behavior.

First, wards define the pool from which individuals draw friends, social contacts, and support. Through close, sustained relationships with these friends, often involving contact outside of church meetings, they may adopt some of their views, policy preferences, and possibly their political identity (Carrell, Sacerdote and West, 2013; Dellavigna et al., 2017; Nickerson, 2008; Sacerdote, 2001). Second, a ward's church meetings serve as a kind of public square where members give sermons, bear testimony in open-mic meetings, and participate in discussions during Sunday School classes. Influence through this channel operates through broader exposure to the views and norms expressed in the congregation, even among individuals who are not close friends (DellaVigna and Kaplan, 2007; Enikolopov, Petrova and Zhuravskaya, 2011; Mutz and Mondak, 2006; Paluck, 2009; Wald, Owen and Hill, 1988). We take up each channel in turn, after first establishing what the treatment is and what a change in registration represents.

7.1 Endogenous versus Contextual Peer Effects

Wards that differ in partisan composition also differ in other compositional attributes including racial makeup, turnout propensity, and types of housing. Our reduced-form estimates measure the effect of being assigned to the type of peers who are more Republican or more Democratic, with everything that goes with that. For many policy questions this composite is exactly the relevant object since any real-world reassignment of peer groups delivers the bundle, not a single attribute (Moffitt, 2001). Even so, it matters which characteristic drives the effects, because the correlations among peer attributes are setting-specific. Suppose our estimates in fact reflected exposure to non-White peers moving voters from Republican to Democrat, a channel with its own literature (Enos, 2016), rather than exposure to Democrats. Then in a setting where race and party are correlated differently than in Utah, our estimates would be a misleading guide to the effects of peer partisanship. The two readings also imply different mechanisms: contact reshaping racial attitudes, and through them policy preferences and ultimately party identity, is a different process from voters being persuaded by their Democratic peers' views and coming to share their political identity. In Manski's terms, we ask whether our estimates reflect contextual effects driven by peers' fixed characteristics or endogenous responses to peers' political behavior.

We implement this comparison one competing characteristic at a time. For each of eight competing peer characteristics—the ward's White share; its shares who voted in the 2010 general, 2012 general, and 2012 primary elections; and its homeowner, male, older, and rural shares—we re-estimate Equation (1) with the competing peer exposure added alongside the partisan exposure, and ask how much the partisan coefficient moves: $\hat{\Delta} = \hat{\beta}_{\text{joint}} - \hat{\beta}_{\text{alone}}$. Because a voter cannot be their own peer, a leave-one-out measure of a competing characteristic would be mechanically correlated with the voter's own value of it, the exclusion bias discussed in Section 4.1. We therefore estimate each comparison separately within subsamples defined by the voter's own value of the competing characteristic (for example, White and non-White voters for the White-share comparison), with voters missing the characteristic excluded. Figure 6 plots $\hat{\Delta}$ for the two main exposure–sample pairs, Democratic exposure among 2012 Republicans and Republican exposure among 2012 Unaffiliated voters, with confidence intervals from a stacked estimation that accounts for the covariance between the two coefficients. The changes are tightly centered on zero. Across the 28 subsamples with at least 5,000 voters, the median absolute change in the partisan coefficient is 0.016, an order of magnitude smaller than the main estimates. The one exception is among 2012 Unaffiliated voters who voted in 2010 where adding peer turnout reduces the Republican-exposure coefficient by 0.15 ($p = 0.002$), yet the coefficient remains large and statistically significant (0.34, $p = 0.032$). Peer race, peer turnout propensity, and peer demographic composition therefore do not account for the partisan effects we estimate, which is consistent with party affiliation being the operative peer characteristic for party switching.

7.2 Strategic Versus Ideological Switching

We next ask what an induced change in registration represents. Party registration changes may reflect adjustments in a voter's ideology, beliefs, and identity, or they may reflect a response to strategic incentives. As discussed in Section 2.4, features of Utah's political landscape and electoral rules create strategic incentives for a voter to register as a Republican, since the Republican primary is closed and the winner of the Republican primary almost always wins the general election. One interpretation of part of our main effects is that Unaffiliated voters realize the stakes of voting in the Republican primary after being exposed to Republican ward members with whom they disagree. Rather than being persuaded by their peers and shifting right, they are reacting against their peers and switching parties in order to vote against them. Alternatively, perhaps Republican voters discuss the primary at church, highlighting its importance and persuading Unaffiliated voters to register so they can have a voice. Both channels imply strategic registration, though they carry opposite implications for how these voters intend to vote.

While such strategic considerations may explain Unaffiliated or Democratic voters switching to the Republican party, it is harder to explain Republicans switching parties in this way. In a state dominated by the Republican party, there is little reason to vote in the Democratic primary (as evidenced by 2020 primary turnout of 18.8% among registered Democrats, com-

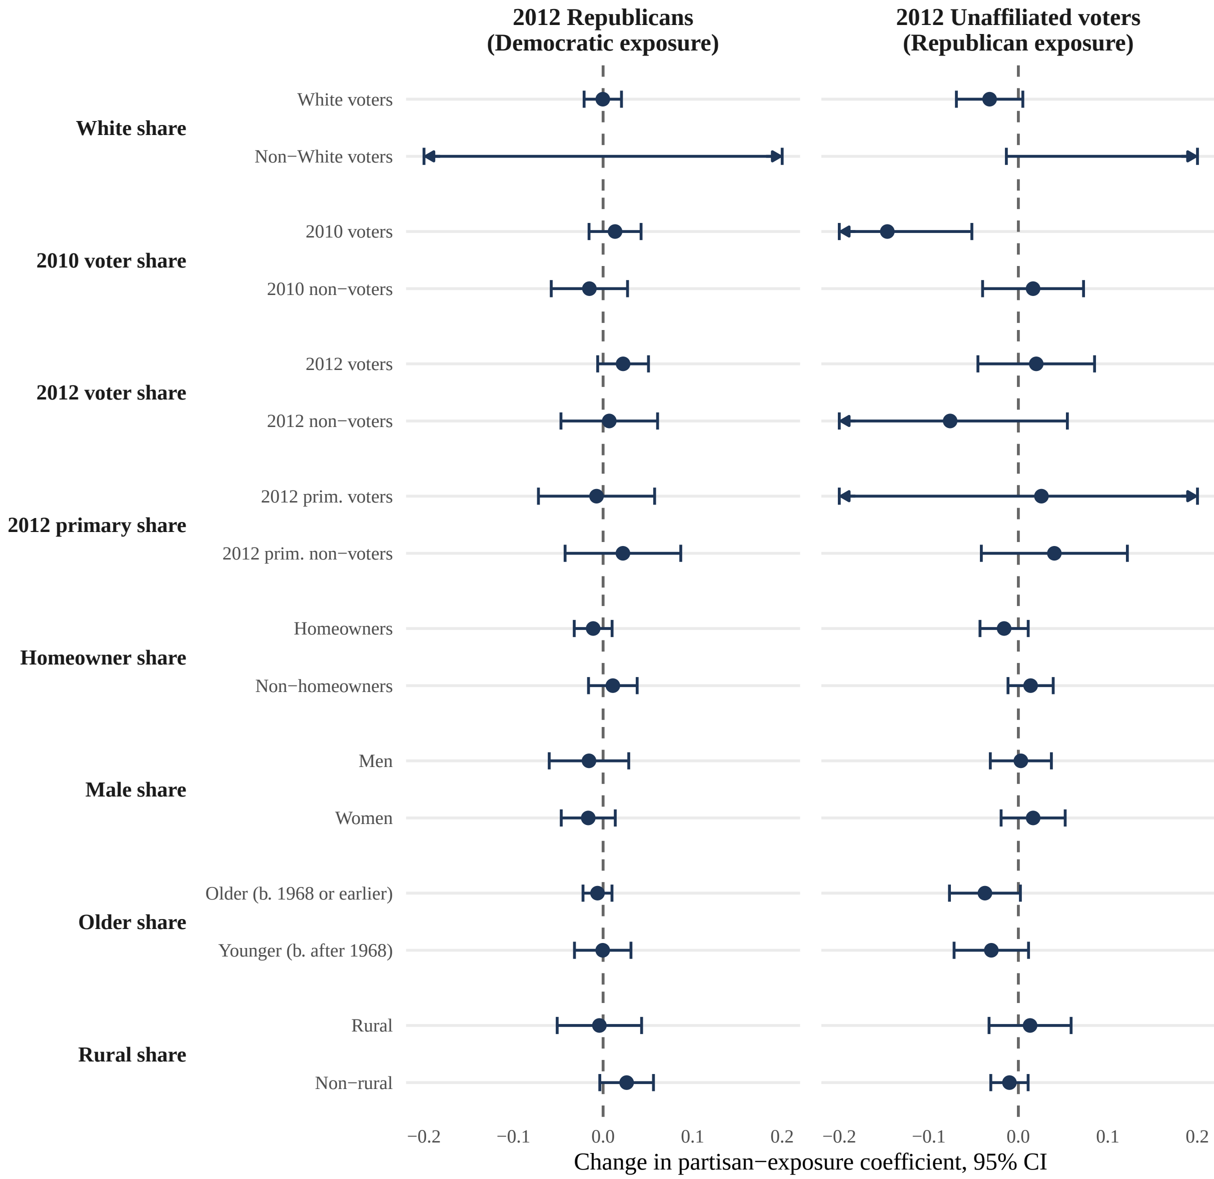

Figure 6: Adding a Competing Peer Exposure Does Not Move the Partisan Coefficient. Each point is the change in the partisan peer-exposure coefficient (on the change in Republican registration) when the competing peer exposure named at left is added to the regression: $\hat{\Delta} = \hat{\beta}_{\text{joint}} - \hat{\beta}_{\text{alone}}$, with 95% confidence intervals from a stacked coefficient-comparison estimator clustered by ward-boundary side and voter. Left panel: Democratic exposure among 2012 Republicans; right panel: Republican exposure among 2012 Unaffiliated voters. Each row's sample conditions on the voter's own value of the competing characteristic, so that a voter's own characteristic cannot mechanically move the competing exposure (Section 4.1). The axis is truncated at ± 0.2; arrowheads mark confidence bounds extending beyond the plotted range (these occur only in small subsamples).

pared with 54.9% among registered Republicans). Furthermore, insofar as party labels enter a person's utility function, someone with Republican preferences may get more disutility from registering as a Democrat than registering as Unaffiliated, and since the Democratic primary is open to Unaffiliated voters, there is no strategic reason for a Republican to re-register as a Democrat.

To further explore this, we re-run our main specifications with interacted outcomes as dependent variables. Specifically, we interact our three party affiliation categories with an indicator of whether a person voted in a 2020 primary election. Each interacted indicator is first-differenced from its 2012 analog, like our main outcomes. Switching to the Republican party and voting in its primary could indicate strategic incentives at play. In contrast, switching to a party but not voting in its primary is more consistent with ideological or identity-based motivations.

Figure 7 shows our effects on these interacted outcomes. For 2012 Unaffiliated voters, we find that exposure to more Republican peers led to an increase in people who both switched to the Republican party and voted in its primary (23.1 percentage points, $p = 0.003$; the switched-but-did-not-vote cell is unchanged, -1.3 percentage points, $p = 0.816$). Because the six cells partition the sample, this increase must come at the expense of the remaining cells; the offsetting declines are spread across remaining Unaffiliated and the Democratic-without-voting cell, though only the latter is individually significant (-10.3 percentage points, $p = 0.016$). This is consistent with strategic switching, but could also reflect a two-step process in which voters switch parties for ideological reasons and then follow their party-specific primary turnout incentives. For 2012 Republicans, we find that exposure to Democratic peers leads to an increase in Democratic registration among both primary voters and non-voters, and the increase is, if anything, concentrated among those who did not vote in a 2020 primary (12.0 percentage points, $p = 0.003$, versus 6.3 percentage points, $p = 0.019$, among primary voters). This result is hard to square with pure strategic switching and likely reflects real ideological change. Appendix H.4 reports the corresponding effects on marginal primary- and general-election turnout for all three baseline-party subsamples.

7.3 Bishop Partisanship and Member Registration

We now turn to the channel itself. The leading alternative to peer-to-peer influence is influence from the congregation's leadership. Each ward is led by a bishop, a lay member with unusual visibility and moral authority who serves for about five years, before being replaced by another ward member. The bishop determines who holds each position in the congregation and who addresses it in meetings, and meets privately with individual members

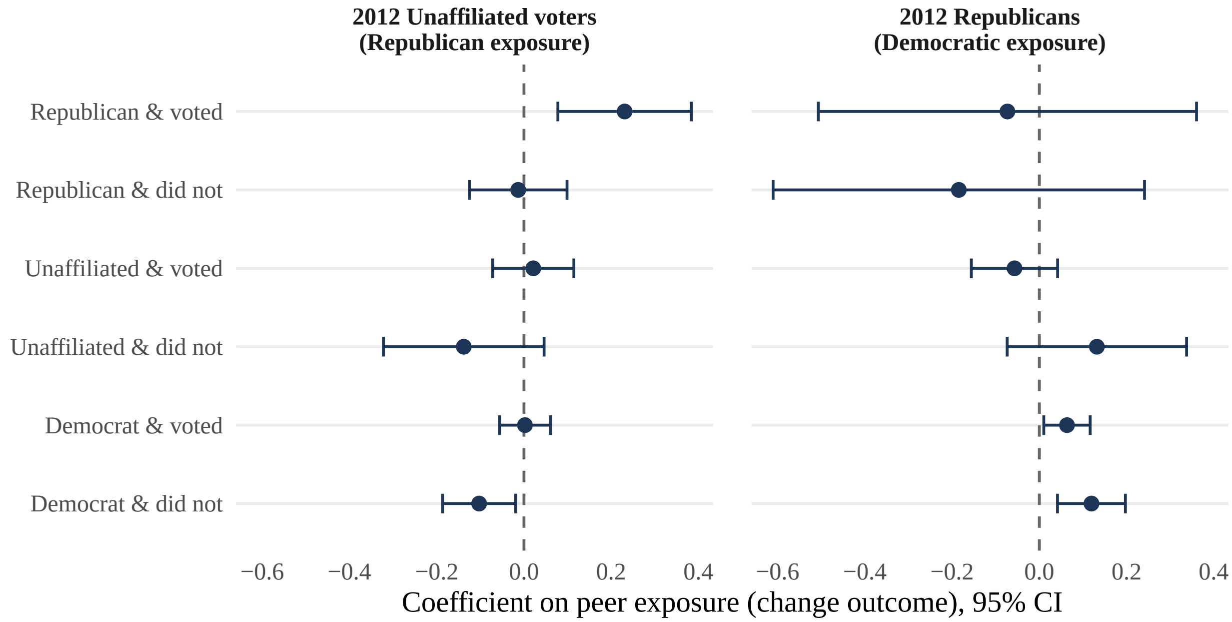

Figure 7: Effect of Peer Exposure on 2020 Party × Primary-Turnout Outcomes (preferred specification, 400m bandwidth). Points are the coefficient on peer exposure with 95% confidence intervals. Left panel: Republican peer exposure among 2012 Unaffiliated voters; right panel: Democratic peer exposure among 2012 Republicans. The six outcomes are the exhaustive 2020 party × 2020 primary-turnout cells (each 2020 party by voted / did not vote in a 2020 primary), each first-differenced from its 2012 analog. All outcomes are in changes.

to counsel and assist with spiritual and financial needs. If the ward effects we document reflect the Sunday church experience rather than social ties among members, the bishop should matter more than any individual congregant, since he has the largest role in shaping church services.

We test this directly with a separate analysis of bishops and their congregations, which uses newly collected data on each ward's bishop from 2023 onward. Because our TargetSmart voter files end in 2021, before our bishop data begins, we instead use L2 voter files covering 2015–2025 for this analysis. Additionally, because bishop turnover is not tied to boundary changes, the sample is every Utah ward rather than only the voters near a new line. We link each bishop to his own voter record to measure his party, and estimate the effect of 2023–2024 bishop changes on 2020–2024 changes in members' party registration. About 24.9% of wards received a new bishop in the five quarters before the 2024 election. A difference-in-differences with individual and election-year fixed effects then compares members' registration changes between the two elections as a function of the incoming bishop's partisanship, so the estimate comes only from wards whose bishop turned over. The identifying assumption is that, within a ward, the incoming bishop's party is unrelated to members' registration trends. Bishops are not random draws from their wards: Republican men are about three times as likely to be called as Unaffiliated men from the same ward, and wealthier men more likely still. However, the party of a ward's incoming bishop is unrelated to the party of the outgoing one, and stake presidents who are not Republican call Republican bishops at the same rate as those who are. Bishops' Republican tilt therefore reflects who is eligible for the calling rather than callers choosing on party, so within a ward the incoming bishop's party looks like a draw from the eligible pool rather than a choice by the stake president (Appendix G.3). The sample is roughly one million voters observed in both elections, split by 2020 party. Appendix G gives the data, the matching procedure, and the validation: applying the same treatment to the 2016–2020 registration change shows no pre-trend, and bishop changes that occurred only after the 2024 election have no effect (the second and third rows of Table 6). The design has the power to detect effects an order of magnitude smaller than our peer-composition estimates.

We find no evidence that bishops' partisanship moves members' registration. Table 6 reports the estimates for 2020 Republicans: acquiring a Republican rather than a non-Republican bishop, or a Democratic rather than a non-Democratic bishop, produces precisely estimated null effects on every registration margin, with 95% confidence intervals ruling out effects larger than 0.6 percentage points. Results for 2020 Unaffiliated voters and Democrats are similarly null (Appendix G). Using the upper end of each of these results' confidence intervals, we estimate the maximum proportion of our main peer effects estimates that can be attributed to bishop effects. Our main effects in Table 3 are scaled to a change in partisan share from zero to one. At most, such a change in partisan share could raise the probability of having a Republican bishop from zero to one, and even that is generous, since bishops' partisanship rises only modestly with their wards' (Appendix Figure G1c). Acquiring a Republican bishop changes an Unaffiliated member's probability of registering Republican by at most 0.7 percentage points, the upper end of the confidence interval, against a peer-composition effect of 22.8 points on the same outcome. Bishops can therefore account for at most 3.1% of the peer effect for Unaffiliated voters. The analogous bound for Republicans' retention is 7.9%.

Combined with the survey evidence that political messaging from the pulpit is rarer in LDS congregations than in any other major U.S. religious tradition (Appendix A.2), this null effect of bishop partisanship suggests that ward effects operate through fellow congregants rather than through the Sunday church experience. The remaining subsections examine the friendship channel directly.

Table 6: Bishop Partisanship and Member Registration: Null Effects (2020 Republicans)

	Republican bishop			Democratic bishop		
	ΔRep	ΔDem	ΔUnaff	ΔRep	ΔDem	ΔUnaff
Bishop party (2020 → 2024)	0.002	−0.001	−0.001	0.004	−0.005	0.000
	(0.002)	(0.001)	(0.001)	(0.010)	(0.008)	(0.005)
Pre-trend (2016 → 2020)	0.001	−0.001	0.000	0.009	−0.006*	−0.004
	(0.002)	(0.001)	(0.002)	(0.006)	(0.003)	(0.005)
Placebo (post-election change)	−0.003	0.001	0.002	0.006	0.000	−0.007
	(0.002)	(0.001)	(0.002)	(0.005)	(0.002)	(0.005)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep. Var. Mean	0.969	0.005	0.026	0.969	0.005	0.026
# Voters	515,828	515,828	515,828	515,828	515,828	515,828
# Wards	3,364	3,364	3,364	3,364	3,364	3,364

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each cell is a separate member-level difference-in-differences with individual and election-year fixed effects (Equation (24)), comparing the change in the registration outcome named in the column header between the 2020 and 2024 elections across wards that did and did not acquire a bishop of the party named in the block header. Bishop party is measured from the bishop's own 2020 voter record. The first row is the main estimate. The second row applies the same 2023–2024 bishop change to the 2016–2020 registration change, which it cannot have caused, as a test of pre-trends. The third row replaces the treatment with bishop changes that occurred only after the 2024 election (2025q1–2026q2), which likewise cannot have caused the 2020–2024 outcome. Standard errors clustered by ward in parentheses. The sample is voters registered with the party named in the caption in 2020 and observed in both elections; bishops themselves are excluded.

7.4 Same- and Opposite-Demographic Peers

If influence operates through friendships, it should flow disproportionately from the peers most likely to become friends. A large literature documents that friendships form more readily between demographically similar people (Feld, 1981; McPherson, Smith-Lovin and Cook, 2001). Friendship networks within wards are partly segregated by gender, both for this reason and because the church organizes separate men's and women's groups (Elders Quorum and Relief Society), which each operate separate home-visiting programs. Additionally, as with any group, ties may form more readily between people of a similar age (Marsden, 1988), race (Currarini, Jackson and Pin, 2009; Moody, 2001), and socioeconomic status (DiPasquale and Glaeser, 1999). If party switching is driven by close, sustained relationships, then the partisan composition of same-demographic peers should matter more than that of opposite-demographic peers in explaining party switching. This comparison is also what separates the two channels the estimating equation pools. As Appendix D.1 shows, dyadic influence among friends and influence broadcast to the whole congregation load on the same ward-mean regressor, so β recovers their sum; a pulpit broadcast is demographically blind, whereas influence flowing through homophilous friendships is not, so a demographic gradient in the effect is evidence that the dyadic channel is active.

We decompose the ward-level peer composition measure into same- and opposite-demographic components and include both in the regression. We define same-demographic peers as those who share the same demographic category as the voter and opposite-demographic peers as all others with a non-missing value for that category. Categories are defined as: male or female for gender, white or non-white for race, 18–29, 30–44, 45–64, and 65 or older for age, and whether the individual links to the CoreLogic property data for homeownership status, which serves as our proxy for socioeconomic status since the voter file records neither education nor income.

We estimate separate models for each demographic, again subsetting by 2012 partisanship, and separate models based on Democratic or Republican exposure. In our regression models, we include terms for both same- and opposite-demographic peer exposure, so that the coefficient on same-demographic peer exposure captures the effect of additional peers of the same demographic group holding constant the number of opposite-demographic peers, and vice versa. As such, the estimation for these specifications takes the form:

$$\Delta Y_{ib} = \beta\left(D_{ibs} \times I_{bs}\right) + \delta D_{ibs} + \theta\left(D_{ibo} \times I_{bo}\right) + \lambda D_{ibo} + \mu_b + f(r_{ib}) + \mathbf{X}_i'\gamma + \varepsilon_{ib}, \qquad (3)$$

where D_{ibs} and D_{ibo} are indicators for being on the high-exposure side of the boundary in the same- and opposite-demographic peer specifications, and I_{bs} and I_{bo} are the corresponding differences in exposure across the boundary. The high-exposure side is defined separately for each component, so that $D_{ibs} \times I_{bs}$ is the signed cross-boundary difference in same-demographic exposure and $D_{ibo} \times I_{bo}$ the signed difference in opposite-demographic exposure; where the two shares are higher on opposite sides of a line, the two indicators differ, which is what allows the regression to separate them. The distance controls use the baseline side convention of Section 4.2.

We present the results for same- and opposite-demographic peer exposure side-by-side in Figure 8. The gender decomposition is the cleanest test, given the Church's formal gender-segregated organizations. Among 2012 Republicans, exposure to same-gender Democratic peers reduces the probability of remaining Republican by 65.0 percentage points ($p = 0.005$), while the corresponding opposite-gender coefficient is small and insignificant ($p = 0.394$). Among 2012 Unaffiliated voters, same-gender Republican exposure raises the probability of registering Republican by 43.2 percentage points ($p = 0.002$), while the opposite-gender analogue is insignificant and opposite-signed. Across the gender panels of Figure 8, the statistically significant effects load on same-gender exposure.

We also test the difference between the same- and opposite-gender coefficients formally using a Wald test. Among 2012 Unaffiliated voters under Republican exposure, all three same-minus-opposite differences carry the sign the friendship account predicts, though they vary in precision: the difference is statistically significant for Republican adoption ($p = 0.016$), marginally significant for Democratic adoption ($p = 0.063$), and insignificant for remaining Unaffiliated ($p = 0.198$). Among 2012 Republicans, the differences under Democratic exposure point the same direction but are individually weaker: $p = 0.058$ for Republican retention, $p = 0.106$ for switching to Unaffiliated, and $p = 0.165$ for switching to Democrat.

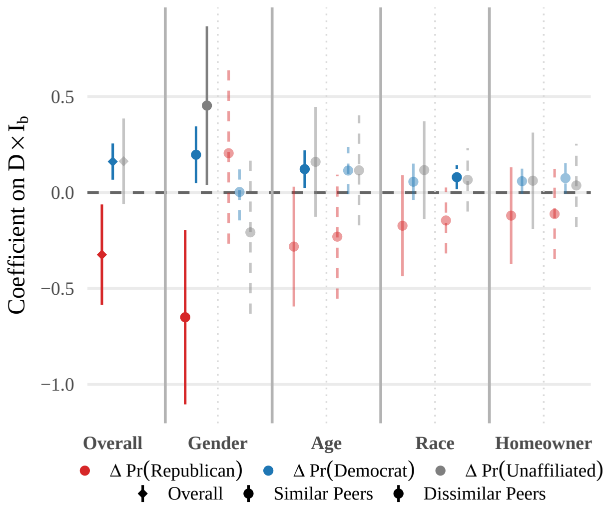

(a) Democratic Exposure for 2012 Republicans

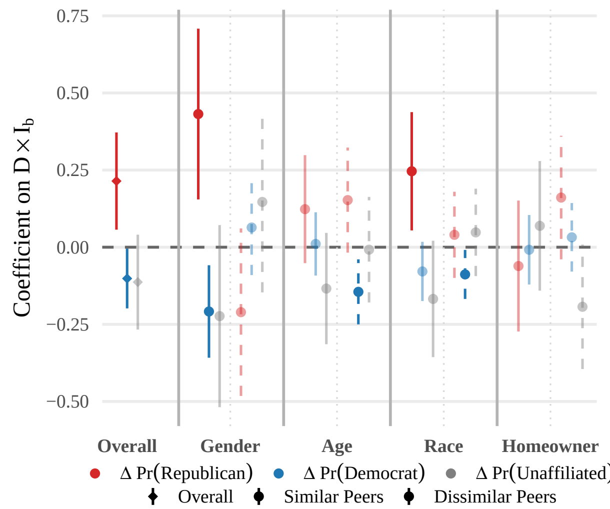

(b) Republican Exposure for 2012 Unaffiliated voters

Figure 8: Same- and Opposite-Demographic Peer Partisan Exposure Effects. Each panel plots the overall peer-exposure coefficient alongside its decomposition into same- and opposite-demographic components (Equation (3)) by gender, age, race, and homeownership, with 95% confidence intervals; solid whiskers mark same-demographic exposure and dashed whiskers opposite-demographic exposure, and faded points are insignificant at the 5% level. The remaining exposure–sample combinations appear in Appendix Figure H1.

The age, race, and homeownership decompositions don't show a consistent pattern, possibly a result of statistical imprecision and possibly an indication that, within wards, friendship networks are not particularly segregated along these three demographic categories. The full estimates appear in Appendix H.

Overall, these patterns, particularly the results for gender, are consistent with a friendship channel. The party-affiliation effects load on same-demographic peers, who are the most likely to become close friends, rather than on opposite-demographic peers who share the ward's public spaces but may form weaker individual ties. We emphasize that this evidence is suggestive. The same-demographic peer measure is mechanically noisier than the pooled measure because each ward's within-demographic sample is smaller, and we lack direct measures of friendship networks that would let us pin down the proposed mechanism.

7.5 Retained versus Newly Assigned Peers

A boundary change does two things at once: it changes the composition of a voter's assigned peer group, and it disrupts existing relationships while creating opportunities for new ones. This duality lets us ask which ties carry political influence. Long-standing friends might matter more because trust accumulates over time and because political events are interpreted jointly with people one already knows. Newly assigned peers might matter more because they bring perspectives and information a voter has not already encountered and discounted (Bursztyn, González and Yanagizawa-Drott, 2020; Golub and Jackson, 2010). The friendship-disruption model of Appendix D formalizes this decomposition of peer exposure into a component from retained peers (former ward-mates assigned to the same new ward) and a component from newly assigned peers, with distinct conformity parameters ρ_{old} and ρ_{new}.

Figure 9 reports the two exposure coefficients for the main estimates. Both channels operate, with strikingly similar magnitudes: for 2012 Unaffiliated voters, Republican exposure through retained peers and through new peers carry nearly identical coefficients (0.16 and 0.17, $p = 0.037$ and $p = 0.035$, on Republican registration); for 2012 Republicans, Democratic exposure through retained peers is somewhat larger than through new peers (-0.31, $p = 0.044$, versus -0.24, $p = 0.130$, on Republican retention), but the two are statistically indistinguishable. It should be noted that the two components separate cleanly only where boundary changes were large and newcomers compositionally distinct, so these estimates are less precise than the pooled ones. With that caveat, the evidence suggests peer influence is neither exclusively a slow product of accumulated trust nor purely a novelty effect as both inherited and newly formed ties transmit it.

7.6 Which Voters Respond

The preceding subsections suggest our effects are driven by direct peer influence between members. If that is correct, effects should be largest among the voters with the most opportunity for political conversation with their ward-mates, namely, those who are more politically engaged, and those living close enough to fellow members for frequent contact. We test for heterogeneous effects by interacting treatment with eight baseline characteristics: prior turnout and rurality, which speak to this prediction, together with age group, gender, race, homeownership, home size, and name commonness, a proxy for individualism following

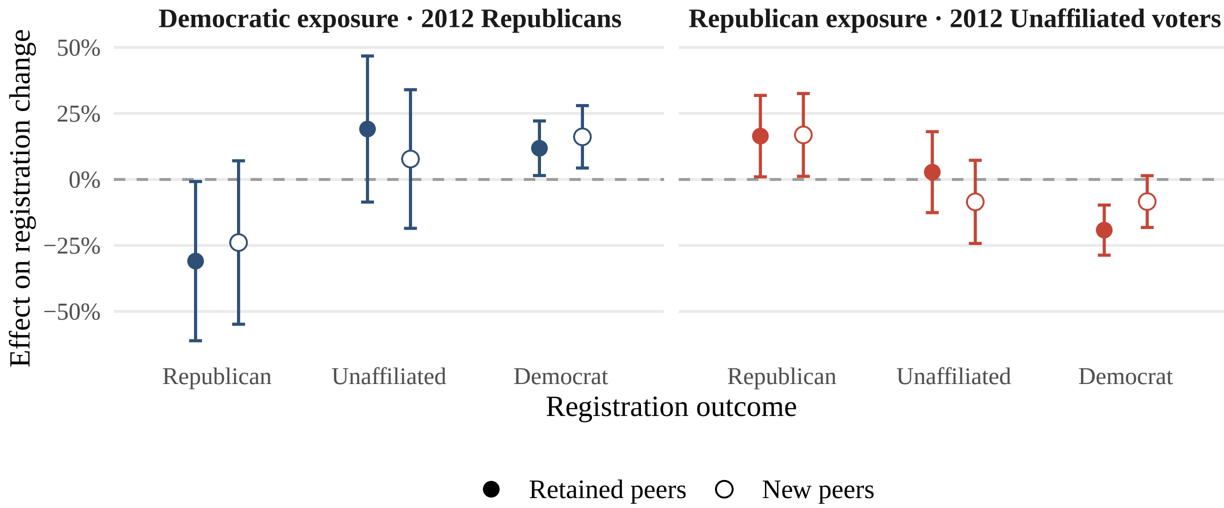

Figure 9: Peer Exposure Through Retained vs. Newly Assigned Peers. Coefficients from the disruption specification of Appendix D, which splits each voter's post-change ward exposure into the component from retained former ward-mates (solid) and newly assigned peers (open), with 95% confidence intervals. Left: Democratic exposure among 2012 Republicans. Right: Republican exposure among 2012 Unaffiliated voters. Full estimates in Appendix H.

Bazzi, Fiszbein and Gebresilasse (2020), measured as the rarity of the voter's first name in the Utah voter file and split at the median. Results for Republican registration appear in Figure 10; the accompanying results for Democratic and Unaffiliated registration are in Appendix H (Figures H2 and H3). Subgroup effects are imprecisely estimated and most differences between subgroups are statistically indistinguishable, so we view the pattern as suggestive.

We find the strongest differences along voter turnout and rurality. Voters who participated in the 2010 general election respond more strongly than non-voters ($p = 0.005$), and non-rural residents respond more strongly than rural ones ($p = 0.040$). The turnout contrast is the sharper of the two and the only one that survives a Bonferroni correction across the eight characteristics we test.

Because the cleanest existing evidence on political peer effects comes from college students, a population still forming its political identity, an important question is whether peer influence extends across the life course. We find that it does. Among 2012 Unaffiliated voters, Republican exposure raises Republican registration for voters aged 18–29, 45–64, and 65 or older, with the largest and most precise estimate in the 45–64 group, while the 30–44 group shows no response. Among 2012 Republicans under Democratic exposure, the 30–44 group is instead the one that responds most. No age group is consistently more responsive than the others across the two comparisons, and in neither is the youngest group the most responsive. Peer influence on partisanship thus does not appear to be confined to early adulthood.

We also investigate whether effects are stronger in regions where a larger share of the population is LDS, which are the regions where our estimates are likely less attenuated (Appendix E). Interacting treatment with quartiles of local LDS density we find, counterintuitively, that effects are concentrated where the LDS share is lowest and are close to zero where it is highest. The gradient is statistically significant for Republicans under Democratic exposure on the ward measure ($p = 0.000$ between the top and bottom quartiles) and for Unaffiliated voters under Republican exposure on the county measure ($p = 0.011$). Two readings are available and we cannot separate them with these data: LDS share is correlated with other regional characteristics, notably the ward's Republican share (0.44), which leaves less scope for out-party exposure; or the true effect really is larger where density is lower, because there the ward is a member's main source of co-religionist ties, whereas in a heavily LDS neighborhood a member can find them next door regardless of which side of the line they fall on. Appendix E.6 reports the full set of estimates and checks.

Heterogeneity in these results can reflect genuine differences in peer responsiveness across individual characteristics, or differences in the share of each subgroup that is LDS and therefore exposed to ward peers at all. Gender illustrates the ambiguity. Treatment effects are slightly larger for women than for men, though statistically indistinguishable. A larger true effect on women could reflect greater responsiveness to social cues (Croson and Gneezy, 2009; Ellingsen et al., 2013), or simply that women tend to be more religiously engaged than men (Becker, Bentzen and Kok, 2026; Miller and Stark, 2002).

Taken together, the mechanism evidence points in a consistent direction. The operative peer characteristic is party affiliation rather than the attributes bundled with it, and the registrations it induces are not merely strategic as movement away from the Republican Party is concentrated among voters who cast no primary ballot at all. The bishop, the one member who addresses the whole congregation, has no detectable partisan influence, and survey evidence indicates that the pulpit is largely apolitical. Influence concentrates among demographically similar peers, who are the most likely to become friends, and it operates through both long-standing and newly formed ties. It is strongest among the voters with the most opportunity for political conversation. This pattern is more consistent with political influence arising from ordinary friendship within the congregation than from the religious institution itself.

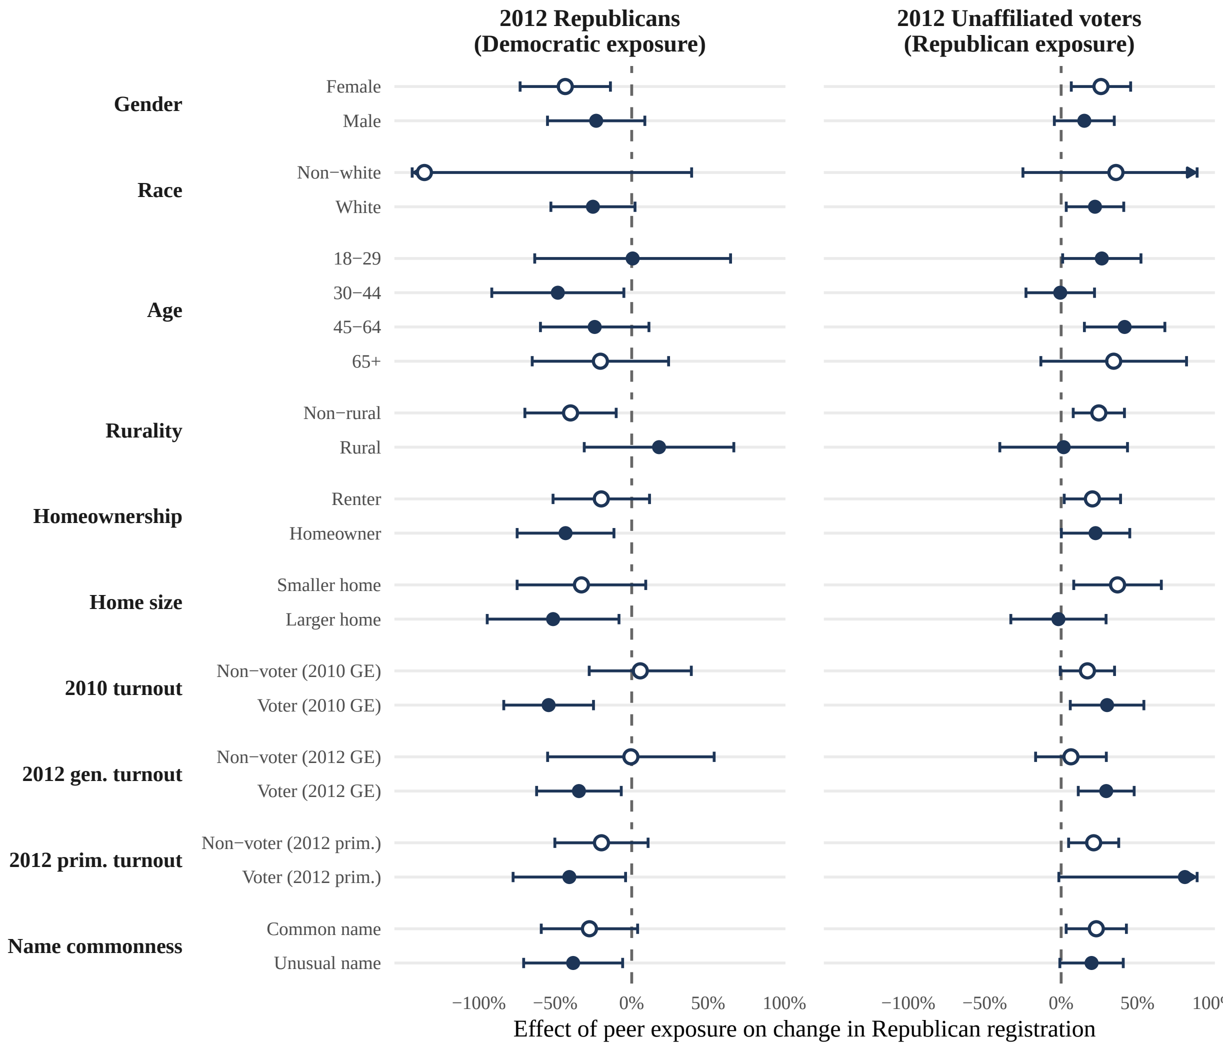

Figure 10: Heterogeneity in the Peer Effect by Individual Characteristics. Each point is the subgroup-specific effect of peer exposure on the change in Republican registration—the peer-exposure coefficient from the preferred specification (400m bandwidth) interacted with the voter's baseline characteristic—with 95% confidence intervals, clustered by ward-boundary side and voter. Open points mark each characteristic's reference subgroup. Left panel: Democratic exposure among 2012 Republicans; right panel: Republican exposure among 2012 Unaffiliated voters. Arrowheads mark confidence bounds extending beyond the plotted range. Full results for all three registration outcomes appear in Appendix Figures H2 and H3.

8 Conclusion

Much of the literature on peer and contextual effects either bundles these two components together, or makes a tradeoff where peer effects can be identified but in limited settings. We provide quasi-experimental evidence that peers from voters' religious congregations, a common and routine social context for millions of American voters, shape partisan affiliation. We do so by combining data on geographic LDS ward reassignments with individual-level data on every Utah registered voter from 2012–2021. With these data, we use a difference-in-discontinuities design to make progress on substantial identification challenges, comparing changes in partisan registration for voters living across redrawn boundaries who differ in which voter was assigned to wards with higher shares of Democratic or Republican registered peers.

This design provides estimates of the effect of religious congregation peers on partisan registration. We find that voters conform to the partisan registration of their peers: unaffiliated voters become more likely to register as Republicans when reassigned to a more Republican ward, and Republican voters become more likely to defect from the Republican party when assigned to a more Democratic ward. These effects are substantively meaningful: assignment to the more Republican side of a typical boundary raises an Unaffiliated voter's probability of registering Republican by 1.4 percentage points, and assignment to the more Democratic side lowers a Republican's probability of remaining Republican by 0.8 points.

We complement these analyses by exploring mechanisms, first testing whether partisan composition, rather than other characteristics of peers, uniquely influences voters, finding strong evidence to support direct influence of Republican or Democratic peers. We next distinguish peer effects from church leadership effects, finding no evidence that acquiring a new bishop of a different party influences members' registration. Finally, we demonstrate that voters are most influenced by peers with similar demographic characteristics to them, suggesting that interpersonal relationships drive our effects.

Beyond its implications for peer-to-peer transmission of political behavior, our results inform understandings of the political impact of religious institutions in modern American politics, and the trajectory of political segregation in social contexts. To the first, in addition to whatever influence the pulpit may have on congregation members' politics, religious congregation peers influence each other to perpetuate political norms. These effects may also reinforce political segregation across congregations, across religions, and across religious versus non religious divides, as homogeneity within congregation begets more homogeneity, as out-party members come to match the partisanship of their peers. More broadly, partisan differences across religious and other social contexts are in part due to behavioral influence of peers in those contexts, rather than being purely a result of sorting processes across contexts.

References

Angrist, Joshua D. 2014. "The perils of peer effects." *Labour Economics*, 30: 98–108.

Audette, Andre P., and Christopher L. Weaver. 2016. "Filling Pews and Voting Booths: The Role of Politicization in Congregational Growth." *Political Research Quarterly*, 69(2): 245–257.

Bail, Christopher A., Lisa P. Argyle, Taylor W. Brown, John P. Bumpus, Haohan Chen, M. B. Fallin Hunzaker, Jaemin Lee, Marcus Mann, Friedolin Merhout, and Alexander Volfovsky. 2018. "Exposure to Opposing Views on Social Media Can Increase Political Polarization." *Proceedings of the National Academy of Sciences*, 115(37): 9216–9221.

Bayer, Patrick, Marcus Casey, W. Ben McCartney, John Orellana-Li, and Calvin Zhang. 2025. "Distinguishing Causes of Neighborhood Racial Change: A Nearest-Neighbor Design." *American Economic Review*, 115(11): 3999–4039.

Bazzi, Samuel, Martin Fiszbein, and Mesay Gebresilasse. 2020. "Frontier Culture: The Roots and Persistence of "Rugged Individualism" in the United States." *Econometrica*, 88(6): 2329–2368.

Becker, Sascha O., Jeanet Sinding Bentzen, and Chun Chee Kok. 2026. "Gender and Religion: A Survey." *Journal of Demographic Economics*, 1–42. First View; published online 28 January 2026.

Bernick, Bob. 2002. "Few back closed primary." Deseret News, June 21, 2002.

Bishop, Bill. 2008. *The Big Sort: Why the Clustering of Like-Minded America Is Tearing Us Apart.* Boston:Houghton Mifflin.

Bisin, Alberto, and Thierry Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." *Journal of Economic Theory*, 97(2): 298–319.

Bisin, Alberto, and Thierry Verdier. 2011. "The Economics of Cultural Transmission and Socialization." In *Handbook of Social Economics*. Vol. 1A, , ed. Jess Benhabib, Alberto Bisin and Matthew O. Jackson, 339–416. Amsterdam:North-Holland.

Blume, Lawrence E., William A. Brock, Steven N. Durlauf, and Rajshri Jayaraman. 2015. "Linear Social Interactions Models." *Journal of Political Economy*, 123(2): 444–496.

Brock, William A, and Steven N Durlauf. 2001. "Interactions-based Models." *Handbook of Econometrics*, 84.

Brown, Jacob R. 2025. "Partisan Conversion Through Neighborhood Influence: How Voters Adopt the Partisanship of Their Neighbors." *The Journal of Politics*, 87(4): 1482–1498.

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

Brown, Jacob R, Enrico Cantoni, Ryan Enos, Vincent Pons, and Emilie Sartre. 2025. "Sources and Extent of Rising Partisan Segregation in the U.S. – Evidence from 143 Million Voters." National Bureau of Economic Research Working Paper 33422.

Brown, Jacob R., Enrico Cantoni, Sahil Chinoy, Martin Koenen, and Vincent Pons. 2023. "The Effect of Childhood Environment on Political Behavior: Evidence from Young U.S. Movers, 1992–2021." National Bureau of Economic Research NBER Working Paper 31759.

Bursztyn, Leonardo, Alessandra L. González, and David Yanagizawa-Drott. 2020. "Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia." *American Economic Review*, 110(10): 2997–3029.

Butts, Kyle. 2023. "Geographic difference-in-discontinuities." *Applied Economics Letters*.

Caeyers, Bet, and Marcel Fafchamps. 2024. "Exclusion Bias and the Estimation of Peer Effects." *Journal of Human Resources*. Published online ahead of print; not yet assigned to a print issue.

Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter.* New York:Wiley.

Campbell, David E. 2004. "Acts of Faith: Churches and Political Engagement." *Political Behavior*, 26(2): 155–180.

Campos, Camila F. S., Shaun Hargreaves Heap, and Fernanda Leite Lopez De Leon. 2016. "The political influence of peer groups: experimental evidence in the classroom." *Oxford Economic Papers*, gpw065.

Cantoni, Enrico, and Vincent 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.

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

Carrell, Scott E., Bruce I. Sacerdote, and James E. West. 2013. "From Natural Variation to Optimal Policy? The Importance of Endogenous Peer Group Formation." *Econometrica*, 81(3): 855–882. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10168.

Carter, D., G. Herschlag, Z. Hunter, and J. Mattingly. 2019. "A merge-split proposal for reversible Monte Carlo Markov chain sampling of redistricting plans." Working Paper, Duke University.

Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology*, 113(3): 702–734.

Chinoy, Sahil, and Martin Koenen. 2024. "Political Sorting in the U.S. Labor Market: Evidence and Explanations." Working paper.

Chyn, Eric, and Kareem Haggag. 2023. "Moved to Vote: The Long-Run Effects of Neighborhoods on Political Participation." *Review of Economics and Statistics*, 105(6): 1596–1605.

Cnaan, Ram, Van Evans, and Daniel W. Curtis. 2012. "Called to Serve: The Prosocial Behavior of Active Latter-day Saints." Working paper.

Croson, Rachel, and Uri Gneezy. 2009. "Gender Differences in Preferences." *Journal of Economic Literature*, 47(2): 448–474.

Currarini, Sergio, Matthew O. Jackson, and Paolo Pin. 2009. "An Economic Model of Friendship: Homophily, Minorities, and Segregation." *Econometrica*, 77(4): 1003–1045.

DeFord, Daryl, Moon Duchin, and Justin Solomon. 2021. "Recombination: A Family of Markov Chains for Redistricting." *Harvard Data Science Review*.

DellaVigna, Stefano, and Ethan Kaplan. 2007. "The Fox News Effect: Media Bias and Voting." *The Quarterly Journal of Economics*, 122(3): 1187–1234.

DellaVigna, Stefano, and Matthew Gentzkow. 2010. "Persuasion: Empirical Evidence." *Annual Review of Economics*, 2: 643–669.

Dellavigna, Stefano, John A. List, Ulrike Malmendier, and Gautam Rao. 2017. "Voting to Tell Others." *The Review of Economic Studies*, 84(1): 143–181.

DiPasquale, Denise, and Edward L. Glaeser. 1999. "Incentives and Social Capital: Are Homeowners Better Citizens?" *Journal of Urban Economics*, 45(2): 354–384.

Djupe, Paul A., and Christopher P. Gilbert. 2006. "The Resourceful Believer: Generating Civic Skills in Church." *The Journal of Politics*, 68(1): 116–127.

Djupe, Paul A., and Christopher P. Gilbert. 2009. *The Political Influence of Churches.* New York:Cambridge University Press.

Ellingsen, Tore, Magnus Johannesson, Johanna Möllerström, and Sara Munkhammar. 2013. "Gender Differences in Social Framing Effects." *Economics Letters*, 118(3): 470–472.

Enikolopov, Ruben, Maria Petrova, and Ekaterina Zhuravskaya. 2011. "Media and Political Persuasion: Evidence from Russia." *American Economic Review*, 101(7): 3253–3285.

Enos, Ryan D. 2016. "What the Demolition of Public Housing Teaches Us about the Impact of Racial Threat on Political Behavior." *American Journal of Political Science*, 60(1): 123–142. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12156.

Feld, Scott L. 1981. "The Focused Organization of Social Ties." *American Journal of Sociology*, 86(5): 1015–1035.

Frake, Justin, Reuben Hurst, and Max Kagan. 2026. "Political Segregation in the US Workplace." *Nature Human Behaviour*, 10(8): 1455–1473.

Gelman, Andrew, and Donald B Rubin. 1992. "Inference from iterative simulation using multiple sequences." *Statistical science*, 7(4): 457–472.

Gerber, Alan S., Gregory A. Huber, and Ebonya Washington. 2010. "Party Affiliation, Partisanship, and Political Beliefs: A Field Experiment." *American Political Science Review*, 104(4): 720–744.

Gerber, Alan S., Jonathan Gruber, and Daniel M. Hungerman. 2016. "Does Church Attendance Cause People to Vote? Using Blue Laws' Repeal to Estimate the Effect of Religiosity on Voter Turnout." *British Journal of Political Science*, 46(3): 481–500.

Golub, Benjamin, and Matthew O. Jackson. 2010. "Naïve Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics*, 2(1): 112–149.

Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology*, 83(6): 1420–1443.

Green, Donald P., Bradley Palmquist, and Eric Schickler. 2002. *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters.* New Haven, CT:Yale University Press.

Gruber, Jonathan, and Daniel M. Hungerman. 2008. "The Church Versus the Mall: What Happens When Religion Faces Increased Secular Competition?" *The Quarterly Journal of Economics*, 123(2): 831–862.

Guryan, Jonathan, Kory Kroft, and Matthew J Notowidigdo. 2009. "Peer Effects in the Workplace: Evidence from Random Groupings in Professional Golf Tournaments." *American Economic Journal: Applied Economics*, 1(4): 34–68.

Hafner, Shay R., and Andre P. Audette. 2023. "The Politics of Church Shopping." *Politics and Religion*, 16(1): 73–89.

Hegselmann, Rainer, and Ulrich Krause. 2002. "Opinion dynamics and bounded confidence: models, analysis and simulation." *Journal of Artificial Societies and Social Simulation*, 5(3).

Hoxby, Caroline M., and Gretchen Weingarth. 2005. "Taking Race Out of the Equation: School Reassignment and the Structure of Peer Effects." Harvard University Unpublished manuscript.

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

Hurst, Reuben, Max Kagan, Matthew Lee, and Justin Frake. 2025. "Organizational Civic Culture: Workplaces as Engines of Democratic Participation." *SSRN Electronic Journal.*

Hutson, Sonja. 2020. "73,000 Utahns register as Republican ahead of June 30 primary." KUER 90.1, June 24, 2020.

Imai, Kosuke, and Kabir Khanna. 2016. "Improving Ecological Inference by Predicting Individual Ethnicity from Voter Registration Records." *Political Analysis*, 24(2): 263–272.

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

Kaplan, Ethan, Jörg L. Spenkuch, and Rebecca Sullivan. 2022. "Partisan Spatial Sorting in the United States: A Theoretical and Empirical Overview." *Journal of Public Economics*, 211: 104668.

Keele, Luke J., and Rocío Titiunik. 2015. "Geographic Boundaries as Regression Discontinuities." *Political Analysis*, 23(1): 127–155.

Keele, Luke, Rocío Titiunik, and José R. Zubizarreta. 2015. "Enhancing a Geographic Regression Discontinuity Design Through Matching to Estimate the Effect of Ballot Initiatives on Voter Turnout." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 178(1): 223–239.

Kenny, Christopher T., Cory McCartan, Ben Fifield, and Kosuke Imai. 2022. "redist: Simulation Methods for Legislative Redistricting." R package.

Ladam, Christina, Ian Shapiro, and Anand Sokhey. 2019. "Social Network Influence on Political Behavior in Religious Contexts." In *Oxford Research Encyclopedia of Politics.* Oxford University Press.

Lee, David S. 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *The Review of Economic Studies*, 76(3): 1071–1102.

Lim, Chaeyoon, and Robert D. Putnam. 2010. "Religion, Social Networks, and Life Satisfaction." *American Sociological Review*, 75(6): 914–933.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, 60(3): 531–542.

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

Marsden, Peter V. 1988. "Homogeneity in Confiding Relations." *Social Networks*, 10(1): 57–76.

Mason, Philip B., Michael B. Toney, and Youngtae Cho. 2011. "Religious Affiliation and Hispanic Health in Utah." *The Social Science Journal*, 48(1): 175–192.

McBride, Michael. 2007. "Club Mormon." *Rationality and Society.*

McCartan, Cory, and Kosuke Imai. 2023. "Sequential Monte Carlo for sampling balanced and compact redistricting plans." *The Annals of Applied Statistics*, 17(4): 3300 – 3323.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology*, 27: 415–444.

Miller, Alan S., and Rodney Stark. 2002. "Gender and Religiousness: Can Socialization Explanations Be Saved?" *American Journal of Sociology*, 107(6): 1399–1423.

Moffitt, Robert A. 2001. "Policy Interventions, Low-Level Equilibria, and Social Interactions." In *Social Dynamics.* , ed. Steven N. Durlauf and H. Peyton Young, 45–82. Cambridge, MA:MIT Press.

Moody, James. 2001. "Race, School Integration, and Friendship Segregation in America." *American Journal of Sociology*, 107(3): 679–716.

Mutz, Diana C. 2006. *Hearing the Other Side: Deliberative versus Participatory Democracy.* Cambridge:Cambridge University Press.

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

Nickerson, David W. 2008. "Is Voting Contagious? Evidence from Two Field Experiments." *American Political Science Review*, 102(1): 49–57.

Paluck, Elizabeth Levy. 2009. "Reducing Intergroup Prejudice and Conflict Using the Media: A Field Experiment in Rwanda." *Journal of Personality and Social Psychology*, 96(3): 574–587.

Perez-Truglia, Ricardo, and Guillermo Cruces. 2017. "Partisan Interactions: Evidence from a Field Experiment in the United States." *Journal of Political Economy*, 125(4): 1208–1243.

Pew Research Center. 2015. "America's Changing Religious Landscape." Based on the 2014 U.S. Religious Landscape Study.

Pope, Devin G. 2024. "Religious Worship Attendance in America: Evidence from Cellphone Data." *NBER Working Paper.*

Putnam, Robert D., and David E. Campbell. 2010. *American Grace: How Religion Divides and Unites Us.* New York:Simon & Schuster.

Roche, Lisa Riley. 2020. “New efforts to get voters to switch to GOP for primary underway — including an unlikely new Republican.” Deseret News, May 25, 2020.

Sacerdote, Bruce. 2001. “Peer Effects with Random Assignment: Results for Dartmouth Roommates.” *The Quarterly Journal of Economics*, 116(2): 681–704.

Sacerdote, Bruce. 2011. “Peer Effects in Education: How Might They Work, How Big Are They and How Much Do We Know Thus Far?” In *Handbook of the Economics of Education*. Vol. 3, 249–277. Elsevier.

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

Strother, Logan, Spencer Piston, Ezra Golberstein, Sarah E. Gollust, and Daniel Eisenberg. 2021. “College roommates have a modest but significant influence on each other’s political ideology.” *Proceedings of the National Academy of Sciences*, 118(2): e2015514117.

The Church of Jesus Christ of Latter-day Saints. 2023. “First Presidency Announces Uniform Worldwide Standards for Ward and Stake Boundaries.” Published: Newsroom, The Church of Jesus Christ of Latter-day Saints.

The Church of Jesus Christ of Latter-day Saints. 2026*a*. “General Handbook: Section 36, Creating, Changing, and Naming New Units.” General Handbook: Serving in The Church of Jesus Christ of Latter-day Saints; accessed April 13, 2026.

The Church of Jesus Christ of Latter-day Saints. 2026*b*. “Mormon Lay Ministry.” Accessed April 13, 2026.

U.S. Census Bureau. 2013. “Voting and Registration in the Election of November 2012, Detailed Tables (P20-568), Table 4a.” Current Population Survey, November 2012 Voting and Registration Supplement.

Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics.* Cambridge, MA:Harvard University Press.

Wald, Kenneth D., Dennis E. Owen, and Samuel S. Hill. 1988. “Churches as Political Communities.” *American Political Science Review*, 82(2): 531–548.

Appendix

Partisan Effects of Religious Peer Groups

Contents

A Setting: Wards and Utah Politics

This appendix provides supplementary evidence on the setting described in Section 2. Appendices A.1 and A.2 present survey evidence that wards are socially dense but politically reserved environments, supporting Section 2.1. Appendix A.3 documents party switching, the size and presidential vote choice of the Unaffiliated electorate, and the geography of partisanship, supporting Section 2.4, and Appendix A.4 characterizes how Utah Latter-day Saints compare to Latter-day Saints nationally and to the U.S. population.

A.1 LDS Wards as Peer Groups

The figures below draw on the Faith Matters Survey, run by Robert Putnam and David Campbell as the empirical backbone of *American Grace* (Putnam and Campbell, 2010). Alongside standard questions about religious affiliation and political attitudes, their survey included a uniquely rich battery on respondents' social networks, which we use to characterize the role wards play in the social lives of LDS Church members.

Figures A1a and A1b show that wards are an unusually important peer group for LDS members. Figure A1a shows that LDS members report having more close friends in their congregation than members of any other major religious tradition. Figure A1b restricts attention to LDS respondents and compares ward friends to work friends across five tie-strength measures; ward friends rank above work friends on every one.

One reason wards play such a central social role is that members spend a great deal of time together. Figure A1c compares weekly attendance at religious services by religious tradition in the Faith Matters Survey. At 88%, LDS respondents report the highest rate of weekly or more frequent attendance of any tradition, far above the next-highest, Evangelical Protestants (49%). Because self-reported attendance may overstate the true rate, it is notable that Pope (2024), using cellphone-mobility data rather than self-reports, similarly finds that LDS Church members attend at the highest rate among the religious groups he studies.

Beyond formal worship, LDS members are also more likely than members of any other major U.S. religious tradition to volunteer for a religious group or congregation. Figure A1d shows that LDS respondents in the Faith Matters Survey 2011 report having done so in the previous twelve months at the highest rate of any tradition. Most LDS religious volunteering takes place through lay-ministry "callings" that, by design, involve repeated collaboration with other members of the same ward, so a high rate of religious volunteering plausibly translates into a high rate of repeated within-ward interaction. Cnaan, Evans and Curtis (2012) estimate that the average LDS member spends about eight hours per week on church-related volunteering, much of it connected to these callings.

McBride (2007) provides an institutional rationale for these patterns. A central goal of the LDS Church, and one of the three missions stated in its handbook, is to "perfect the saints," meaning to help members progress toward salvation by becoming more godly people. McBride argues that the Church pursues this goal in large part by structuring ward life around dense, frequent social interaction: members are encouraged to serve one another, and ward programs deliberately bring members into contact with people they would not otherwise choose as companions. This emphasis on the ward as a cohesive social unit also serves a more practical purpose. By embedding new and low-commitment members in a thick web of within-ward interactions, the ward helps them accumulate religious capital and, over time, transition into higher-commitment participants.

Together, the patterns in Figure A1 suggest that an LDS member's ward is a socially consequential peer group, and one where peer effects, insofar as they exist, are likely to operate.

A.2 Political Activity in Wards

Although wards anchor the social lives of their members, they are comparatively reserved venues for explicit political activity. Two indicators from the Faith Matters Survey illustrate the contrast with other religious traditions. Only 2% of LDS respondents report that political issues are discussed from the pulpit at least monthly, compared with 41% of Jewish respondents (Figure A2a); and just 16% report that their congregation organizes voter registration or distributes voter guides, compared with 49% of Black Protestants (Figure A2b). On both indicators LDS attenders rank among the lowest of any tradition shown. These patterns suggest that the effects documented above operate through exposure to ward peers rather than through institution-level political messaging or mobilization.

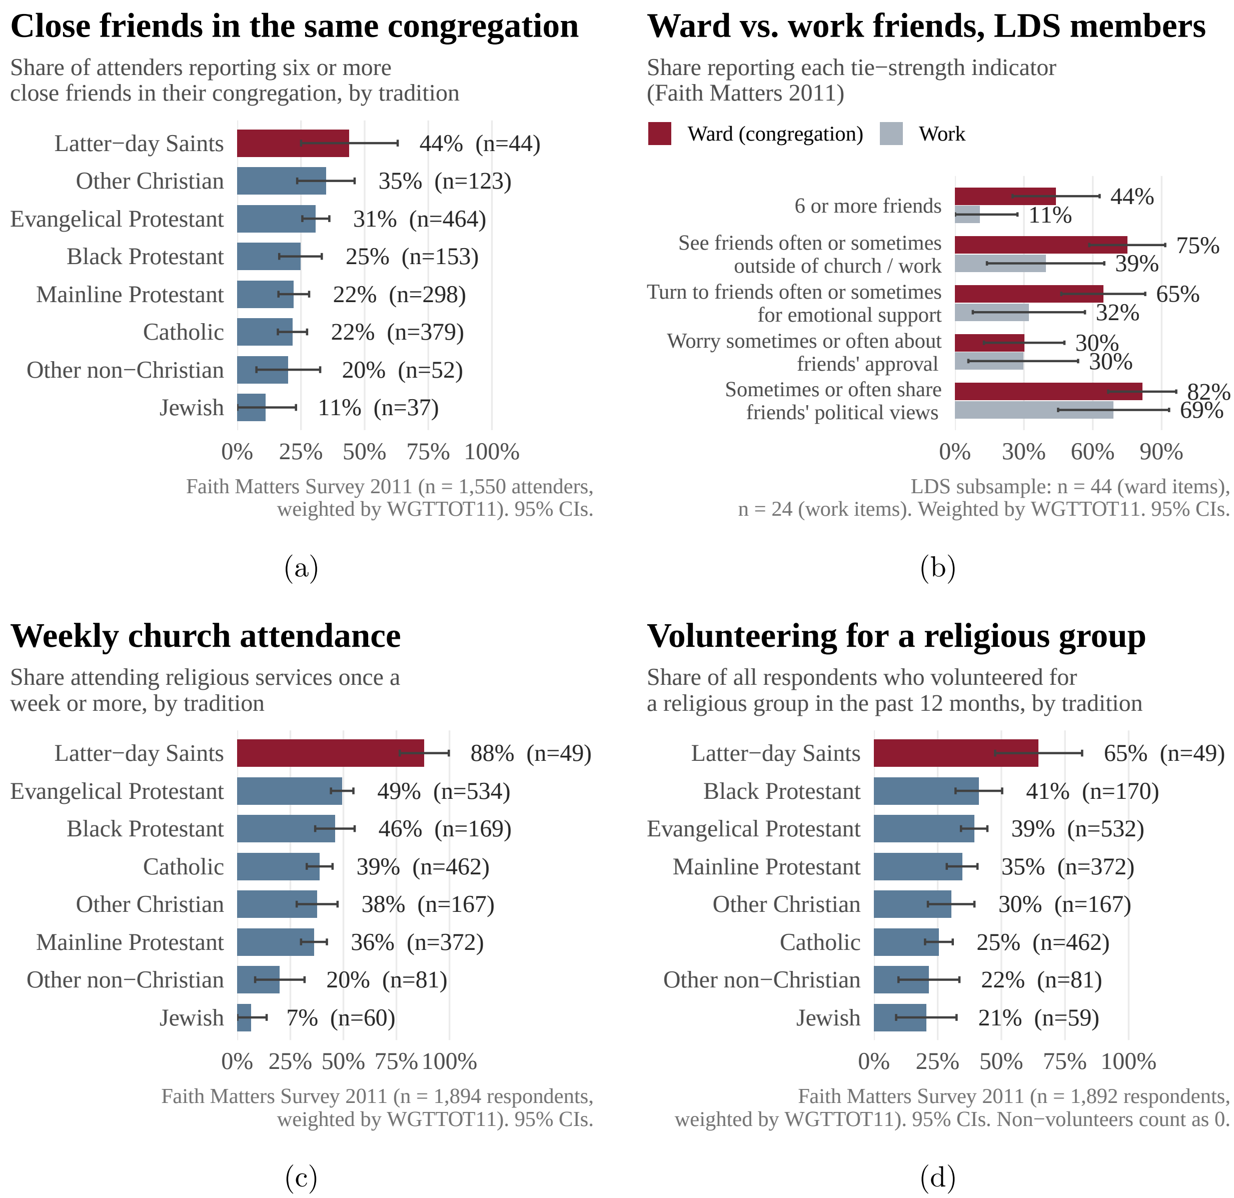

Figure A1: Wards as a social peer group, and the time members invest in them, by religious tradition. Panel (a): share of attenders reporting six or more close friends in their congregation. Panel (b): importance of ward vs. work friends for LDS members across five tie-strength measures (LDS subsample only). Panel (c): share reporting that they attend religious services once a week or more. Panel (d): share reporting they volunteered for a religious group in the past 12 months (non-volunteers counted as 0). Faith Matters Survey 2011 (weighted by WGTTOT11). Error bars are 95% confidence intervals computed using the Kish effective sample size, which accounts for the precision lost to survey weighting.

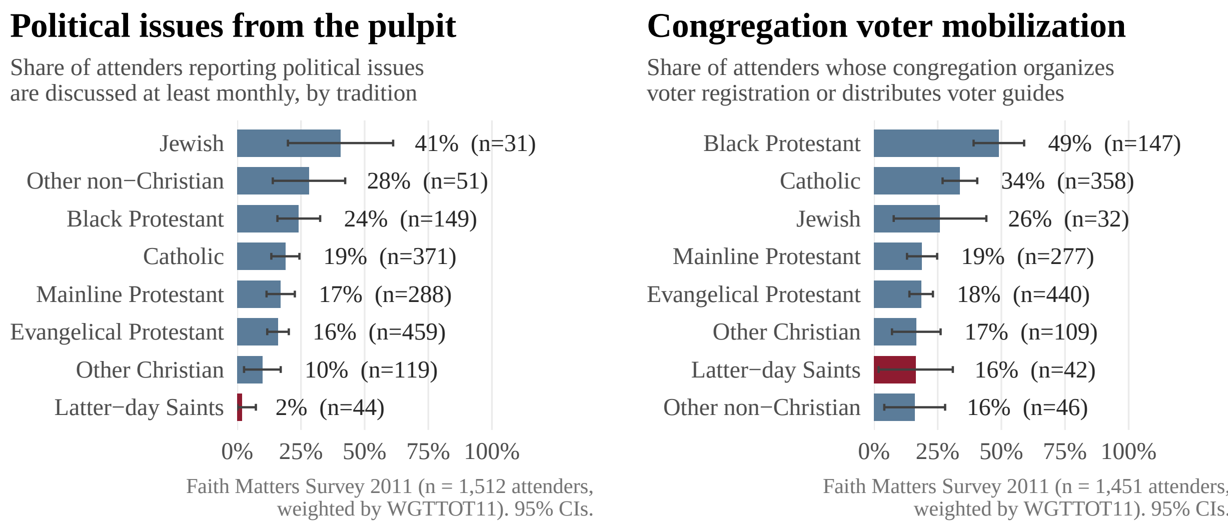

(a) (b)

Figure A2: Congregations as political institutions, by religious tradition. Panel (a): share of attenders reporting that political issues are discussed at least monthly from the pulpit. Panel (b): share reporting that their congregation organizes voter registration or distributes voter guides. Faith Matters Survey 2011 (weighted by WGTTOT11). Error bars are 95% confidence intervals computed using the Kish effective sample size, which accounts for the precision lost to survey weighting.

A.3 Additional Background on Politics

This subsection documents four features of Utah's partisan landscape that support Section 2.4: party switching was common over our study window; Unaffiliated voters make up a large share of the electorate; self-identified Independents, the closest survey proxy for Unaffiliated registrants, split their presidential votes between the major parties rather than leaning consistently toward either; and partisanship is sorted geographically at a broad scale while still varying considerably between neighboring wards.

Figure A3 traces registration flows between 2012 and 2020 among all registered Utah voters in our panel observed in both years, and shows the first two features. Party switching was common: 27.9% of 2012 Unaffiliated voters moved to the Republican Party and 12.7% to the Democratic Party, while 12.7% of 2012 Republicans and 28.3% of 2012 Democrats left their party. And Unaffiliated voters are a large share of the electorate, 46% of registrants in 2012 and 34% in 2020.

Figure A4 reports presidential vote shares among self-identified Utah Independents (including "leaners") in the Cooperative Election Study (CES). We use self-reported party identification here because the CES's Catalist-validated registration is not available before

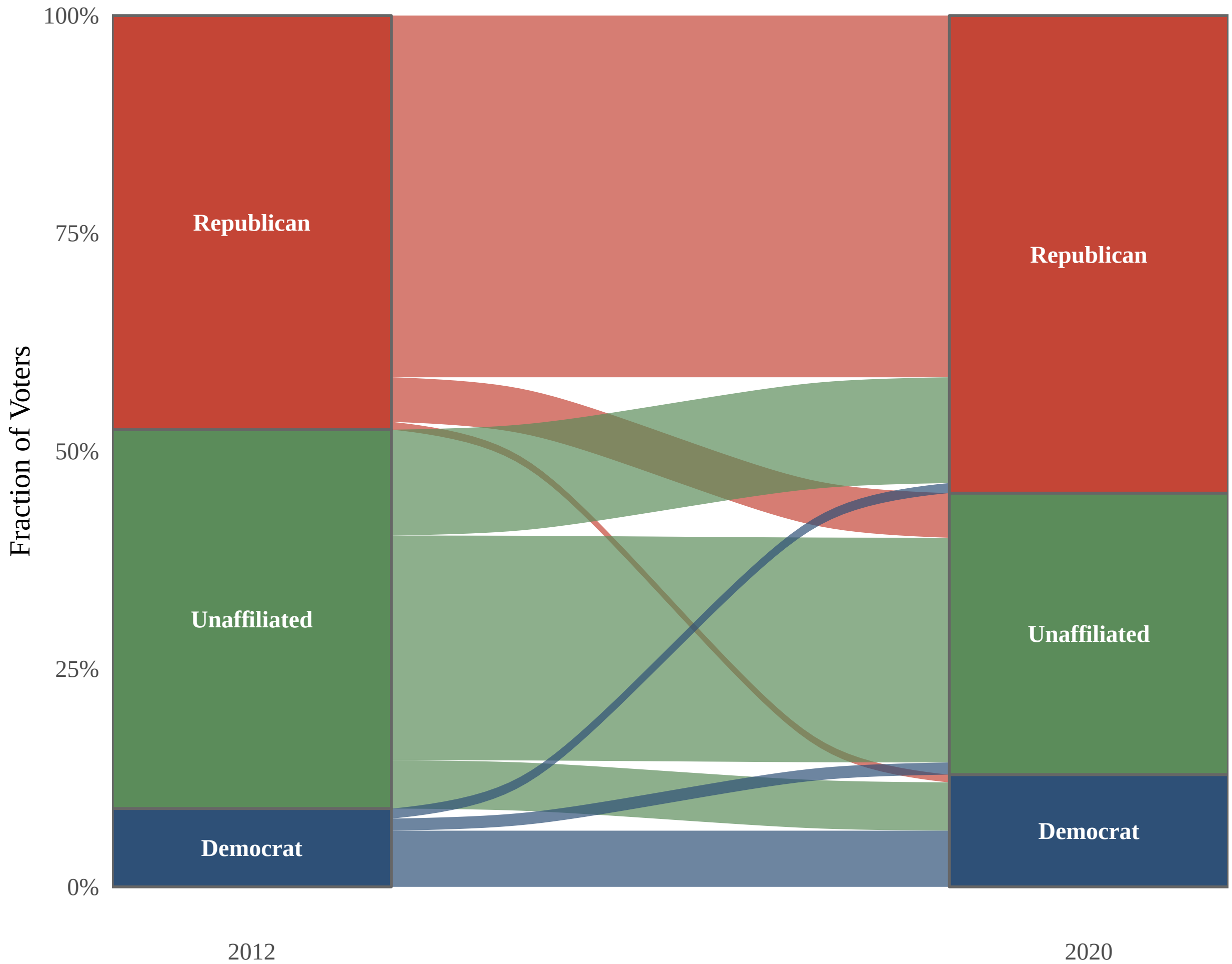

Figure A3: Party Registration Flows in the Utah Electorate, 2012–2020. Each band traces the movement of voters between party registrations from 2012 to 2020, among all registered Utah voters in our panel observed in both years.

2016; the two definitions agree closely where both are available. Independents split their votes between the major parties rather than leaning consistently toward either. They leaned left in 2008 and 2020, swung sharply right in 2012, when Mitt Romney's candidacy was unusually popular in Utah, and cast a substantial share of their 2016 votes for candidates outside the major parties, most notably Evan McMullin.[32]

Finally, Figure A5 maps ward-level Republican registration shares across the Utah Valley and Salt Lake areas, computed from the 2012 voter file over 2013 ward boundaries. As in much of the country, partisanship is sorted geographically at a broad scale: Republican

[32]Pooling the 2020 CES, Utah self-identified Independents cast 41% of their presidential ballots for the Republican, 52% for the Democrat, and 7% for an Other candidate. Utah Independents defined via the Catalist-validated registration variable cast 37%, 55%, and 7% for the same three categories.

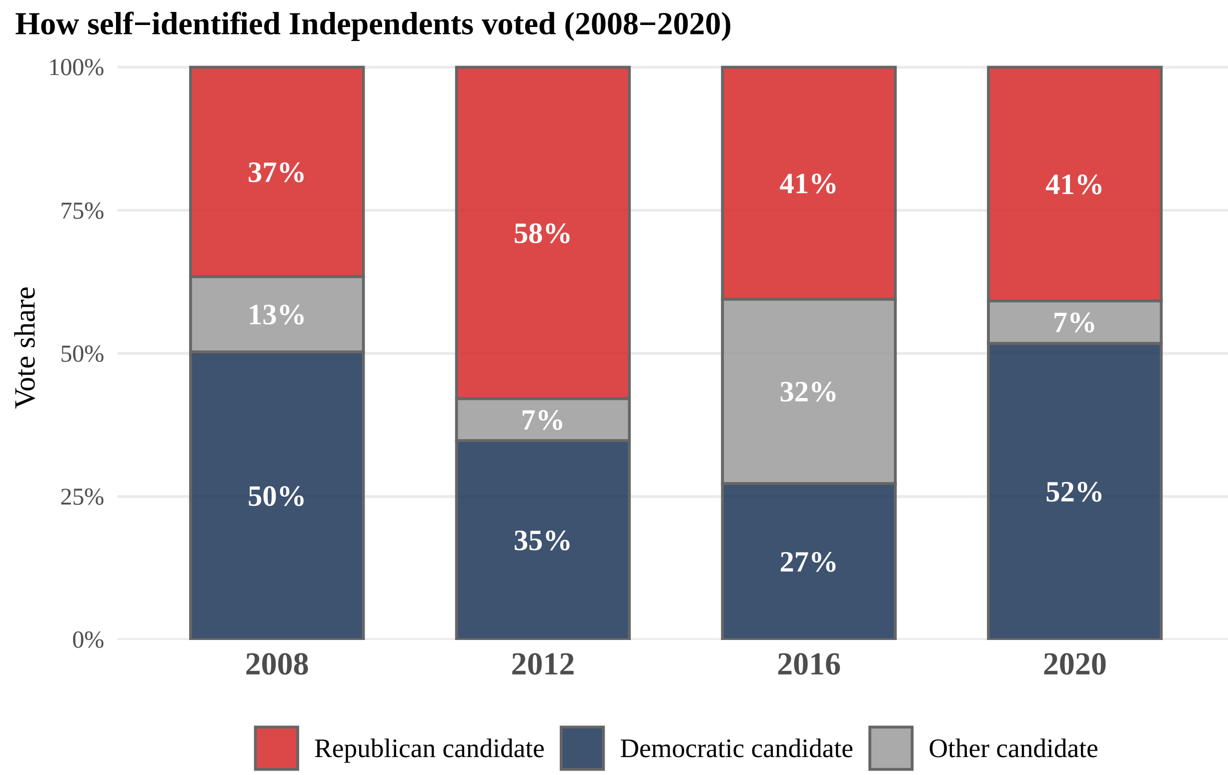

Figure A4: Presidential vote shares among Utah self-reported Independents (CES `pid7` in $\{3, 4, 5\}$, i.e., leaners and pure Independents) in the 2008, 2012, 2016, and 2020 elections. Bars stack to 100% among respondents who reported a presidential vote choice (non-voters and "not sure" excluded). "Other" aggregates all non-major-party candidates and includes the substantial Evan McMullin share in 2016. Bars are weighted using each wave's post-election common weight.

share is lowest in and around Salt Lake City and rises toward the suburbs, and Utah County is heavily Republican. Within these broad patterns, however, partisan composition varies considerably between neighboring wards, and it is this local variation that our design exploits (Section 4).

A.4 Representativeness of Utah Latter-day Saints

Our estimates come from Latter-day Saints in Utah, so a natural question is how well they generalize to other populations. Table A1 compares Utah LDS respondents with all Utah respondents and with the full U.S. adult population, using the CES common content pooled across the 2012–2020 even-year waves (the study window); a column for U.S. Latter-day Saints is included for reference.

We first compare Utah Latter-day Saints to Utah residents as a whole. On demographics

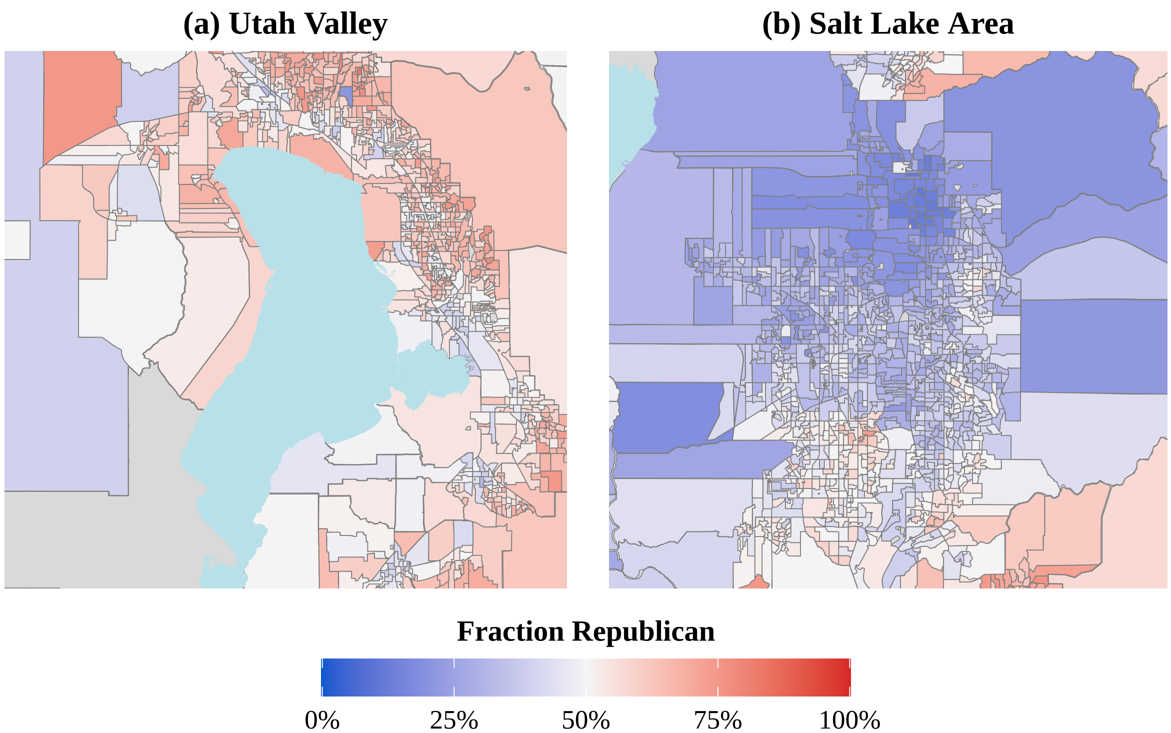

Figure A5: Ward-Level Republican Share, 2012. Each polygon is a single ward (2013 boundaries), shaded by the fraction of 2012 registered voters affiliated with the Republican Party. Gray areas have no mapped ward.

and socioeconomic status the two groups are close: they are within 4 percentage points of each other in every age bracket and every income bracket, in gender, and in education (34% vs. 31% with a bachelor's degree). Utah Latter-day Saints are somewhat more likely to be White (93% vs. 88%), married (70% vs. 61%), and to have children at home (41% vs. 36%). They differ sharply only on religion and politics: 76% attend church weekly or more against 46%, and 57% identify as Republican against 40%, with Democratic identification correspondingly lower (8% vs. 17%). This partisan gap matters for our design: because we cannot observe religion, every Utah voter enters the sample, and the fact that Latter-day Saints are disproportionately Republican is what drives the different patterns of attenuation across political party which we quantify in Appendix E.

Next, we compare Utah Latter-day Saints to all U.S. adults. Here the differences are larger, but much of the demographic gap is one Utah shares as a whole rather than one specific to Latter-day Saints. Utah Latter-day Saints are younger (30% are aged 18–29 vs. 21% nationally, and 13% are 65 or older vs. 20%), much more likely to be married (70% vs. 49%) and to have children at home (41% vs. 25%), and much more likely to be White (93%

vs. 72%); on each of these, all Utah respondents sit closer to Utah Latter-day Saints than to the national figure. On socioeconomic status the gaps are small: Utah Latter-day Saints are modestly more educated (34% vs. 28% with a bachelor's degree) and about as likely to have family income below $50,000 (47% vs. 50%). The largest differences are again religion and politics: weekly attendance is 76% against 28%, and Republican identification 57% against 27%.

For external validity, the relevant benchmark is the populations in which political peer effects have previously been identified. The cleanest existing evidence assigns peers directly, to college roommates (Strother et al., 2021) and classmates (Campos, Hargreaves Heap and Leite Lopez De Leon, 2016): populations concentrated in a single age bracket, selected on education, and still forming their political identities. Utah Latter-day Saints span the full adult age distribution and differ only modestly from the U.S. population on education and income, so on these dimensions they are considerably more representative than a student sample. The dimension on which they are least representative is religious observance. The mechanisms we document operate through a congregation with unusually high attendance and dense social ties (Appendix A.1). Populations with weaker group attachment should experience weaker doses of the same peer exposure, so we would expect effects of assigned peer composition to be attenuated, not absent, in less tightly knit institutional settings such as schools, workplaces, or less observant congregations. Conversely, the fact that our subjects are adults with presumably settled political identities argues against dismissing the estimates as a feature of an unusually persuadable population.

Table A1: Comparing Utah Latter-day Saints to Utah, U.S. Latter-day Saints, and the U.S. Population

	Utah LDS	Utah (all)	U.S. LDS	U.S. (all)
Demographics				
Age 18–29	29.7	28.7	27.1	21.4
Age 30–44	29.4	28.5	30.2	24.3
Age 45–64	27.5	28.1	29.0	34.7
Age 65+	13.4	14.6	13.8	19.6
Female	51.7	50.1	51.0	51.7
White	93.2	87.9	84.2	71.7
Family				
Married	69.6	60.5	63.7	49.3
Has child under 18	41.1	35.6	39.9	25.4
Socioeconomic status				
Bachelor's degree or higher	34.0	30.5	31.5	28.0
Family income <$50k	47.2	49.5	50.1	50.4
Family income $50–100k	38.4	36.4	34.8	32.9
Family income $100k+	14.4	14.1	15.1	16.8
Party identification				
Democrat	7.9	16.7	14.0	34.2
Republican	57.1	39.5	49.8	27.5
Independent / other	35.0	43.8	36.2	38.3
Religiosity				
Attends church weekly or more	76.2	46.5	69.1	27.5
Prays daily or more	79.0	56.2	75.2	45.9
Religion "very important"	77.2	49.6	71.6	40.4
Respondents (unweighted)	1,286	2,643	4,220	296,335

Notes: CES common content, pooled 2012–2020 even-year waves. Cells are weighted percentages (each wave's main survey weight, normalized within wave); shares computed among respondents answering each item. LDS = CES religious-affiliation category "Mormon." Income brackets harmonized across the pre-2018 and post-2018 CES family-income codings. Party identification is the three-category self-report; "Independent / other" includes other-party identifiers and "not sure."

B How and Why Ward Boundaries Change

This appendix documents additional evidence behind the description of ward boundary changes given in section 2.3. As noted in there, ward boundary changes are undertaken by regional lay leaders called stake presidents, who are responsible for a collection of 5-10 contiguous wards. Appendices B.1–B.3 report a survey of former stake presidents, describing how it was fielded, who responded, and a summary of their responses. Because our identification rests on the claim that the placement of new ward lines is unrelated to the political composition of the households on either side, the goal of the survey is to characterize the decision rule stake presidents actually use. Appendix B.4 tests one implication of that rule, that lines are not drawn around a ward's most consequential member. Appendices B.5 and B.6 then document the 2013–2017 boundary changes that define our treatment variation: how many wards were created and dissolved, how changes cluster, and where they occurred.

B.1 A Survey of Former Stake Presidents: Design and Response

Sample. We assembled a roster of stake presidencies from *Church News*, which publishes each new stake presidency as it is called, and restricted it to men who presided over a Utah stake at some point during our 2013–2017 study window. Contact information was matched to this roster by name and birth year from a commercial people-search database. This yielded 492 former stake presidents with at least one candidate email address.

Fielding. The survey was fielded by cold email; exhibit B5 reproduces the message. Recipients were invited to answer in writing, or to schedule a short phone call.

Response rate. We sent 521 messages to 492 individuals, retrying at a second address where one was available. Of the 521 addresses, 131 bounced, leaving 113 individuals unreachable at every address we had. Six replies came from a different person of the same name, whom the address in fact belonged to. This leaves 373 intended recipients presumptively reached, of whom 19 gave a substantive response, a response rate of 5.1%. One additional recipient declined to respond. Fourteen respondents replied in writing and five spoke with us by phone.

Coverage. The 19 respondents led 19 distinct stakes in seven Utah counties—Utah (7), Salt Lake (4), Davis (3), Weber (2), and one each in Cache, Iron, and Washington—spanning the fast-growing suburban fringe, the aging urban cores of Salt Lake City and Ogden, and small towns in the state's rural south and north. Eighteen of the nineteen had personally overseen at least one boundary change; several had overseen four or more, and one had presided over a stake that grew from roughly a dozen wards to thirty-five or forty.

B.2 What Prompts a Change and Where Lines Are Drawn

Respondents were unanimous that boundary changes respond to changes in local membership counts. Often the catalyst for change was when a ward had either too few members to staff its lay ministry, or so many members that not all could have opportunities to contribute. In regions where the church is declining, respondents described a common trajectory. Typically serving older neighborhoods, the homes in their wards tended to be both small and expensive, a combination that disproportionately drove out the LDS population who tend to have large families (Table A1): young families were priced out of the market and older more established families were often too large to fit in the homes. The consequence was a shrinking pool of church members from which to staff a fixed set of positions. At the extreme, some respondents reported importing volunteers from neighboring wards to temporarily fill positions. In contrast, respondents from regions where the church is growing described the opposite problem—too many members to fill a limited number of positions and a heavy burden on leaders to care for a large and growing population. Most of these high growth regions are located in the suburbs where growth is driven by the construction of new developments, and in some areas, new development is happening at such a rate that some wards were splitting multiple times per year.

The primary criterion is a set of ward-level counts. Respondents consistently described drawing lines to equalize, across the resulting wards, the total number of members, the number of active members, the number of men eligible to serve in leadership positions, and the number of children and youth by age group. These are precisely the statistics the Church's mapping tool reports for any proposed configuration (Appendix B.3), and respondents described iterating until the configuration satisfied all of them at once. Several also mentioned the Church's published guidelines on minimum unit size as a binding constraint on the set of feasible maps.

Balance across wards, not similarity within them. Question 3 of our instrument asked directly whether, in dividing an area containing both a wealthier and a poorer neighborhood, respondents would mix the two across the resulting wards or separate them. Every respondent who addressed the question chose mixing, usually for practical reasons: leadership capacity, welfare need, and ministering burden are correlated with neighborhood affluence, so a ward drawn entirely from one type of neighborhood inherits either all of the burden or all of the capacity. Respondents described splitting a single apartment complex across three or four wards, dividing a retirement community among the wards that surrounded it, and reshuffling apartment complexes purely to equalize welfare demands. One summarized the principle at the level it actually operates: "the socioeconomic conditions were more just at the neighborhood level, not individual. For example, we didn't want a ward of all government-subsidized apartment buildings . . . and no single member homes. We wanted a mix." Two qualifications recur. In many stakes the question is moot, because the stake is economically homogeneous; and mixing is subordinate to the counts, since respondents were usually not willing to draw an awkward or fragmented boundary to achieve a demographic mix.

Neighborhoods are kept intact where possible. Respondents used major roads, rivers, and other natural features as boundaries where they existed. Where a neighborhood had to be divided, the strongly preferred practice was to run the line along the backs of lots—down the middle of the block—rather than down the center of a street, so that facing neighbors stay in the same ward. Respondents were explicit that this is about preserving existing social ties, and several volunteered that a boundary change does in practice attenuate friendships across the new line, even between households a few dozen feet apart, because the shared weekly and mid-week activities that sustained the friendship no longer coincide. School boundaries received little weight because wards are much smaller than school catchment areas so it was seldom a concern.

Individual households enter only at the margin, and politics not at all. Respondents differed in how much they considered individual households, some reviewing the map household by household and others focusing mostly on assignment at the neighborhood level. Some respondents also mentioned considering where a sitting bishop lived to avoid disrupting a functioning ward's leadership, while others deliberately ignored it on the grounds that a bishop serves a few years while a boundary lasts much longer. All respondents denied accounting for member's partisanship, stating that it was never discussed and is not something they were in a position to observe. The same applied to household income with respondents reporting that they had no individual-level income data, and only considered socioeconomic conditions at the neighborhood level.

Finally, respondents consistently described the last step of the process as spiritual rather than analytical. Having assembled the data and narrowed the feasible maps, presidencies prayed about the result and, in several accounts, changed it. Whatever one makes of this substantively, it is worth noting for our purposes that it is a source of variation in the final line placement that is likely unrelated to the observable characteristics of the households on either side.

B.3 Mechanics, Approval, and Compliance

Respondents who served recently described a Church-provided online mapping tool that lets a stake president draw hypothetical ward boundaries over his stake and immediately see, for each resulting unit, key membership statistics: total members, men eligible for leadership positions, and children and youth by age. The tool is used iteratively—one respondent recalled going through seven or eight configurations before finding one that met his criteria simultaneously—and respondents who served earlier described improvising the same capability. Respondents reported little formal training on how to draw boundaries: the Church publishes guidelines on unit size and on the counts that must be met, and these are treated as constraints, but the choice among feasible maps is left to the stake presidency's judgment.

Once a configuration is settled, the stake president submits it to Church headquarters in Salt Lake City for review and approval, a process respondents described as taking weeks to months. Approval is not automatic: headquarters evaluates the proposal against unit-size requirements and can reject a configuration that would leave a ward, or the stake, below minimum size. Respondents also described headquarters as having its own view of where growth and decline are headed, informed by data the stake does not see. Changes that move wards between stakes involve additional coordination: an area authority convenes the affected stake presidents, gathers ward-level statistics from each stake's clerks, and mediates disagreements. Only after approval is the change announced to members, typically the Sunday before it takes effect. Ward members—including, in several accounts, the sitting bishops—learn of it at that point rather than during deliberations.

Compliance with the resulting assignments is high. Asked how often members of their stake attended a ward other than the one they were assigned to, the modal respondent could recall one or two instances in his entire tenure; one estimated non-compliance at under one percent of members, and another said he had been asked for an exception exactly once. The exceptions respondents described are individually unusual and require formal approval from Church headquarters, requested through the stake president, which respondents described headquarters as generally willing to grant conditional on their recommendation.[33]

B.4 Bishop Placement Relative to New Boundaries

This diagnostic addresses a specific gerrymandering concern: of all ward members, the sitting bishop is the individual whose ward assignment leaders are most likely to consider when drawing a line, since each ward requires a bishop and bishops must reside within their ward's boundaries. If boundary-drawers were carefully engineering which side of a new line particular members fell on, bishops—the most consequential such members—should end up disproportionately close to new boundary lines, where small placement choices determine their assignment. Bishop identities are observed only in the later ward boundary data covering 2023–2026, which is what we use for this analysis (Appendix G), rather than the 2013–2017 boundaries from our main analysis. Figure B1 compares the distribution of distance to the nearest new boundary line for bishops and for rank-and-file voters. The two densities are nearly identical: bishops are no closer to new boundaries than ordinary voters, consistent with boundary placement that does not finely engineer the assignment of even its most salient members.

[33]The exception to this pattern is the Church's language-specific units, which serve a defined language community across an area much larger than a typical ward and which respondents said do draw substantial attendance from outside their boundaries. These units are not part of our analysis, which is restricted to geographically defined English-language family wards.

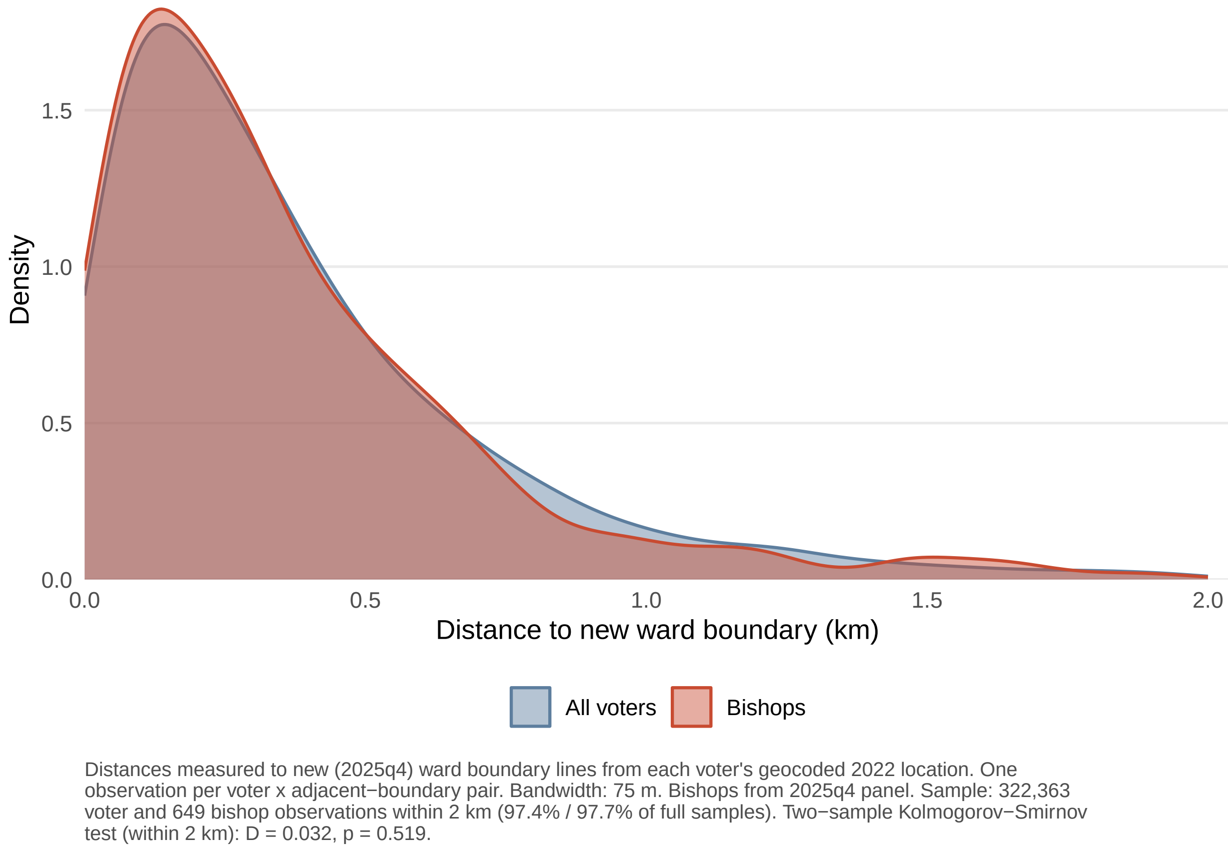

Figure B1: Kernel densities of distance from each voter's geocoded 2022 location to the nearest new (post-2023) ward boundary line, separately for all voters and for bishops identified in the 2025q4 bishop panel.

B.5 Boundary Changes, 2013–2017

This subsection expands on the boundary-change patterns summarized in Section 2.3, documenting in detail how LDS ward boundaries changed between 2013 and 2017, the years for which we have ward boundary data. Over this period, 239 new wards were created and 99 old wards were dissolved, raising the total number of wards in our dataset from 3,975 to 4,115; Table B1 gives the year-over-year changes, broken down into these two flows. The changes documented here define the treatment variation used throughout the paper: the analysis sample of Section 3 consists of voters living near these newly drawn boundaries, and the placebo analysis of Section 5 repeats the design on analogous post-2023 changes.

A boundary change creates a new ward, closes an old one, or simply shifts a line between neighbors. A new ward is carved out of pieces of one or more existing wards, and a closed ward's area is distributed among the continuing wards around it. Because every affected ward

Table B1: Annual ward accounting, 2013–2017

	Period			
	2013-14	2014-15	2015-16	2016-17
Wards (start of period)	3,975	4,022	4,060	4,077
+ *New wards formed*	59	55	64	61
− *Wards dissolved*	12	17	47	23
Wards (end of period)	4,022	4,060	4,077	4,115

must still meet the staffing requirements and membership counts described in Appendix B.2, one change often sets off a local chain reaction in which the boundaries of several neighboring wards are adjusted at once. We call each such group of contiguous wards whose boundaries changed simultaneously a boundary-change cluster.

Between 2013 and 2017, we observed 488 such clusters. 196 (40.2%) of these result in the creation of at least one new ward, 72 (14.8%) result in the closing of at least one old ward, and the other 220 (45.1%) involve simple realignments of boundaries. The vast majority of these clusters (83%) fall entirely within a single stake, while the other 17% cross stake boundaries. Changes involving multiple stakes are more logistically complicated, since they require coordination across stake presidencies (Appendix B.3). Table B2 reports these summary statistics year by year.

Most boundary-change clusters are small, involving only a few wards. Figure B2 shows the distribution of cluster sizes, characterized by the number of wards in that cluster at the start and end of the period. The modal cluster (27.3% of cases) involves adjusting a ward boundary between two neighboring wards without a change in the number of wards or larger adjustment spillovers. Other common boundary change events include splitting two wards into three (9.8%), splitting one ward into two (8.8%), and readjusting the boundaries of three wards (7.6%). While most clusters are small, at the extreme, around 1% of boundary change events involve 15 or more wards.

Table B2: Boundary-change clusters by year pair

	Period			
	2013-14	2014-15	2015-16	2016-17
Boundary change clusters	118	92	141	137
By type of cluster				
Ward count increased	51 (43.2%)	46 (50.0%)	49 (34.8%)	50 (36.5%)
Ward count decreased	11 (9.3%)	14 (15.2%)	27 (19.1%)	20 (14.6%)
Boundaries realigned only	56 (47.5%)	32 (34.8%)	65 (46.1%)	67 (48.9%)
By stake geography				
Contained within one stake	98 (83.1%)	77 (83.7%)	111 (78.7%)	119 (86.9%)
Crossing stake boundaries	20 (16.9%)	15 (16.3%)	30 (21.3%)	18 (13.1%)

Notes: A boundary-change cluster is a contiguous group of wards whose boundaries were redrawn simultaneously. Pooled across the four year pairs, the average cluster contained 3.44 wards at the start of its period. Percentages are shares of that period's clusters.

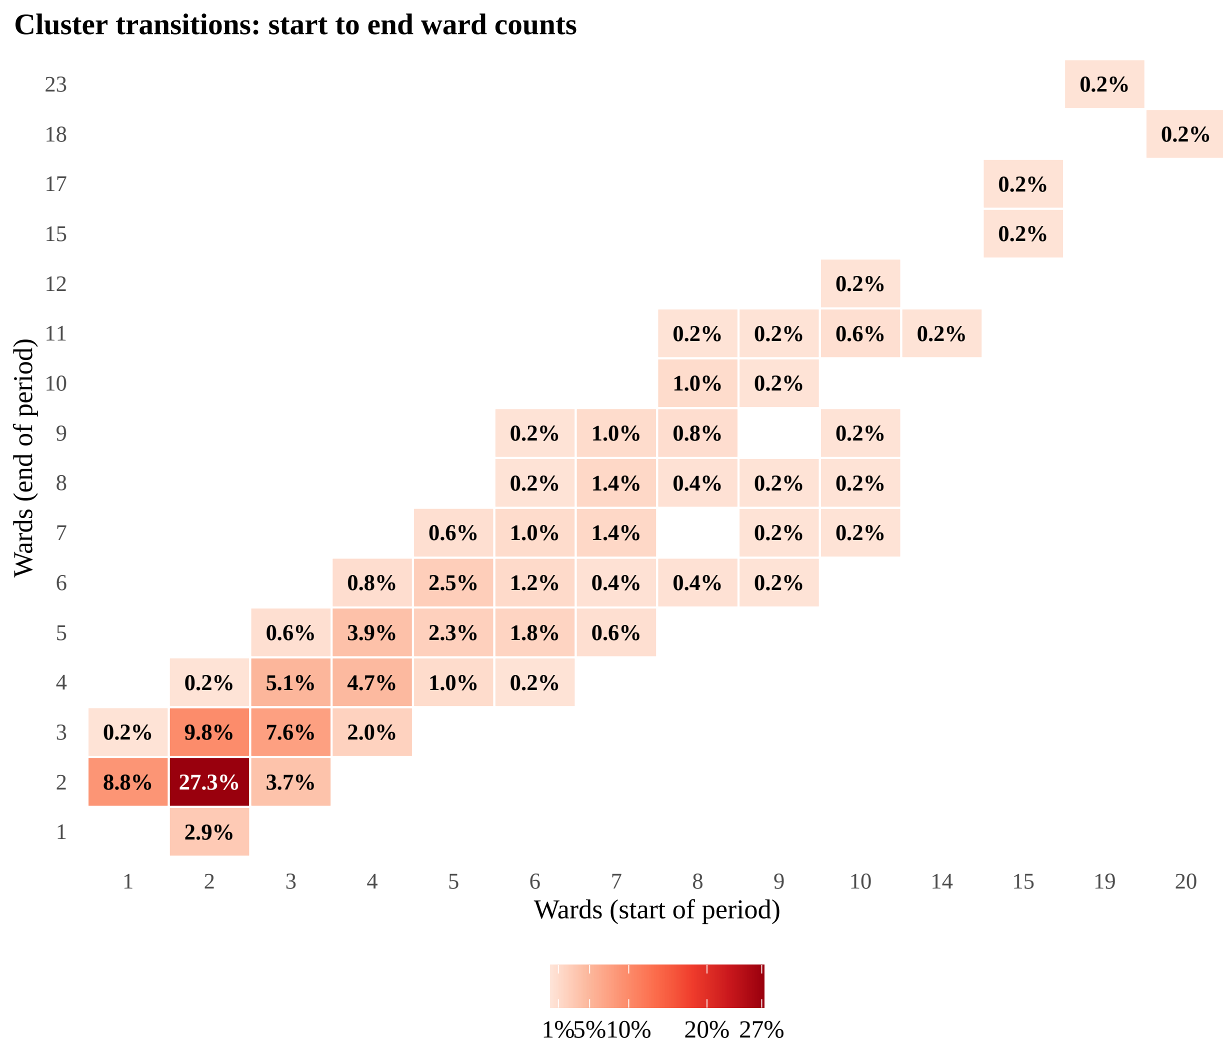

Figure B2: Joint distribution of cluster size characterized by number of wards at the start (x-axis) and end (y-axis), pooled across changes from four periods: 2013–14, 2014–15, 2015–16, and 2016–2017.

B.6 Where Boundaries Change

Figure B3 maps the spatial distribution of boundary change clusters in Utah's two most populous counties—Salt Lake and Utah Counties—colored by whether they resulted in the creation of new wards, the dissolving of old wards, or a simple boundary adjustment. In general, the maps show that new wards are more likely to be created in high LDS areas (the southwest suburbs of Salt Lake and most of Utah County), while low LDS areas (e.g., downtown Salt Lake City) are more likely to close a ward. These patterns suggest a dynamic of increasing geographic polarization along religious lines during this period–high LDS areas attracted more LDS move-ins or converted more of their non-LDS neighbors, while in low LDS areas, church members either moved away or stopped attending church.

Table B3 reports the growth statistics summarized in Section 2.3: over our study window, population grew far faster in clusters that added a ward (108.0%) than where boundaries were merely realigned (38.5%) or a ward closed (14.3%), with the same ordering for counts of registered voters (95.9% vs. 35.5% and 8.0%). Figure B4 shows the corresponding pattern in new single-family home construction: areas that added a ward experienced pronounced construction booms in the years leading up to their boundary change. Notably, the booms precede the change and then taper off as developers move on to new undeveloped areas.

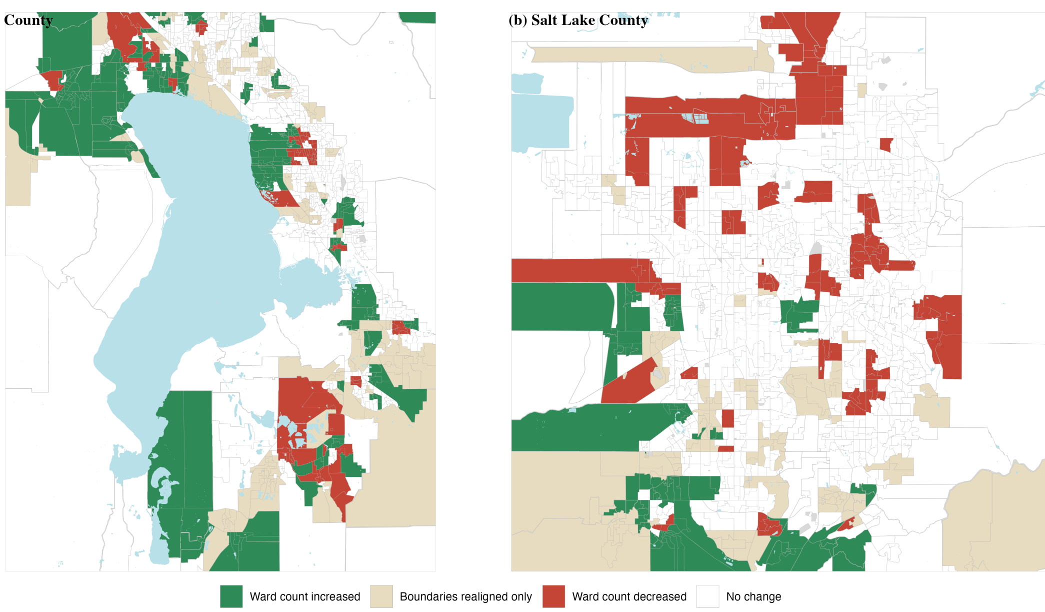

Figure B3: Wards in Utah County (left) and Salt Lake County (right), with boundary change clusters shaded by their 2013–2017 net change type (e.g., green indicates a cluster where the ward count increased between 2013 and 2017). Wards are plotted based on their 2017 boundaries.

Table B3: Ward-level summary statistics by cluster type

	Ward count increased	Boundaries realigned only	Ward count decreased
Wards × Periods (N)	691	640	349
Clusters (N)	196	220	72
Population and growth			
Avg population (2010 imputation)	549	557	792
Avg population (2020 imputation)	1,141	772	905
Population growth, 2010–2020	108.0%	38.5%	14.3%
Voter-file dynamics			
Avg voters (2012)	192	213	315
Avg voters (2020)	393	291	340
Voter-file growth, 2012–2020	95.9%	35.5%	8.0%
Stayed in ward, 2012–2020	37.3%	38.0%	35.8%

Notes: The unit of observation is a ward-period pair where the periods are 2013–14, 2014–15, 2015–16, and 2016–17. Wards are defined using their start-of-period boundary. Population endpoints come from intersecting the start-of-period ward polygons with the 2010 and 2020 Census blocks. Voter-file rows use the 2012 and 2020 voter-file cross sections joined to the same start-of-period ward definition.

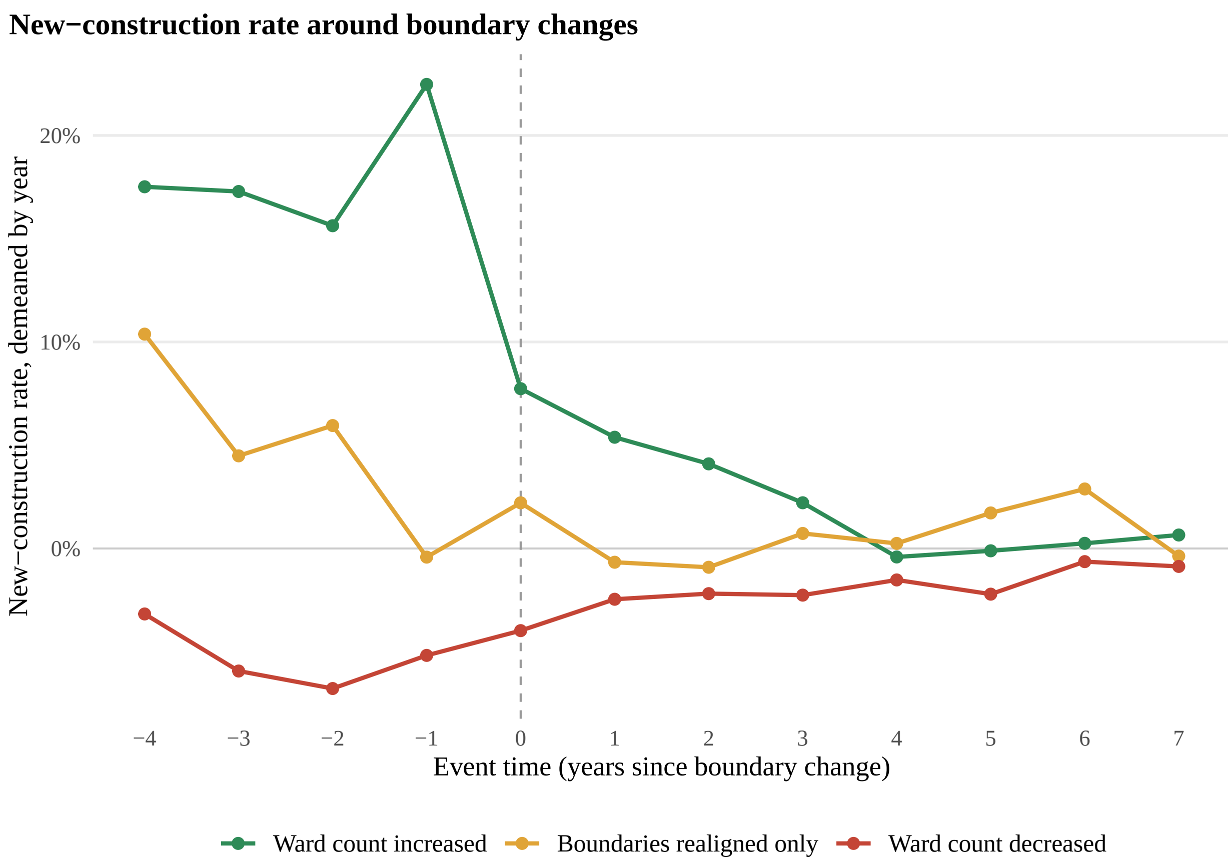

Figure B4: Year-over-year new-construction rate in event time, pooled across the 2013–2014 through 2016–2017 cohorts. Event time is calendar year minus the cohort's change year; the dashed line marks the change. Each ward-year rate is demeaned by the cross-ward mean for that calendar year.

Subject: Your experience with ward boundary decisions — UC San Diego dissertation research

Dear President *[Last Name]*,

My name is Kurtis Gilliat, and I am a PhD candidate in economics at the University of California, San Diego. My dissertation research studies how people's communities and peer groups influence their beliefs and behavior (for example, the decision to start a new business). One of my projects examines how ward communities in the Church of Jesus Christ of Latter-day Saints shape outcomes for members.

As part of that work, I am trying to better understand how ward boundary decisions are made in practice. I understand you previously served as president of the *[Stake Name]*, and I was hoping you would be willing to share a little about your experience.

I have listed four short questions below (also attached as a one-page document, in case that is easier to print or share). If you are able, simply replying to this email with your thoughts—on as many or as few of the questions as you like—would be tremendously helpful. A few sentences is plenty. If a conversation would be easier, I would also be glad to set up a brief 5–10 minute phone call at whatever time works for you.

If you did not personally oversee a boundary adjustment during your service, I would still value your perspective on how you would have approached such a decision.

1. What circumstances typically led to ward boundary adjustments in your stake? When changes were made, what goals were you generally trying to accomplish?
2. When deciding where to draw new boundary lines, how much weight was given to the circumstances of individual households or members, compared with broader factors such as neighborhoods, major roads, school boundaries, or other geographic features?
3. When creating or adjusting wards, did you tend to prioritize similarity within wards or balance across wards? For example, if you were dividing an area containing both a wealthier neighborhood and a lower-income neighborhood into two wards, would you be more likely to (a) draw boundaries so that each ward includes a mix of both neighborhoods, or (b) draw boundaries so that one ward comes mostly from one neighborhood and the other ward mostly from the other?
4. More generally, what principles or considerations most strongly influenced where boundaries were ultimately drawn?

Any responses you provide will be kept confidential and used only for academic research. If I were ever to quote your response in my work, I would do so anonymously and only with your permission.

Thank you very much for your time.

Best regards,
Kurtis Gilliat
PhD Candidate, Department of Economics
University of California, San Diego

Figure B5: Recruitment email sent to former Utah stake presidents, August 11–25, 2026. Bracketed italics denote mail-merge fields filled in for each recipient. The four questions in the body were also attached as a one-page document.

C Data Construction and Validation

This appendix documents how each data source used in Section 3 was built and validated. Appendices C.1 and C.2 describe the cleaning and linking of the TargetSmart voter panel and validate it against published state totals. Appendices C.3 and C.4 describe our process for digitizing ward boundaries and provide ward-level statistics and distributions. Appendices C.5 through C.8 cover the CoreLogic and school-boundary linkages and report summary statistics for the analysis sample and its comparison groups, including our placebo sample. Finally, Appendix C.9 describes the Meetinghouse Locator scrape.

C.1 Cleaning and Linking TargetSmart Records

The procedure described below is adapted from Appendix C.1 of Brown et al. (2023), which describes an analogous cleaning and within-state linking pipeline applied to the full 50-state TargetSmart voter-file snapshots. We apply the same steps to the Utah-only subset used in this paper.

Cleaning the raw TargetSmart state-by-year files

The raw TargetSmart data consist of annual snapshots, one for each year between 2012 and 2021. These snapshots contain each voter's name, residential address, gender, race, date of birth, whether or not they were registered and voted in prior elections, and their party affiliation.

TargetSmart also provides a "voterbase ID" (henceforth VBID), which uniquely identifies each row for a given state and year, and an "exact track ID" (henceforth ETID), which represents its efforts to link individuals across years and states. We use this information as well as voters' first name, middle name, last name, date of birth, and vote history to de-duplicate the TargetSmart data, so that for each state-by-year file, the observation used in the analysis is the most likely current record for each voter. We also use this information to expand on TargetSmart's linkage model to further link records that likely correspond to the same individual.

To begin, we take the following steps to clean and de-duplicate the raw TargetSmart state-by-year files:

1. We remove hyphens and spaces from first and last names. We capitalize all letters of first and last names.

2. We recode invalid ZIP codes and census block group IDs as missing.

3. We use TargetSmart's information on whether a voter is found in Social Security death records to drop deceased voters.[34]

4. We use TargetSmart's information from the United States Postal Service National Change of Address database to drop voters that no longer reside at their listed residence.

5. We de-duplicate records with the same ETID, first name, and last name, giving preference to the record whose registration status is "Registered" (vs. "Unregistered"), whose voter status is "Active" (vs. "Inactive," based on recent election participation), and with the most recent registration date. We do this because, for a given individual, outdated records may exist alongside the active, most recent record in the same state voter file snapshot. When de-duplicating, we harmonize the vote history so that if any of the dropped records indicate that the person voted in a particular election, we maintain this turnout history in the record that we preserve for that individual.

6. We repeat the previous step, de-duplicating records with the same first name, middle initial, last name, date of birth, and ZIP code.[35]

7. We drop any record where the voter's age is listed as under 18 and the individual is listed as "Registered."

Linking records across years

To link rows within the state corresponding to the same individual across multiple years—in other words, to assign a unique voter identifier—we take the following steps:

1. We split the date of birth field into year, month, and day. Even though TargetSmart uses commercial data to reconstruct the exact day of birth for voters, the TargetSmart data feature high frequencies of dates of birth ending in 01 (i.e., indicating a voter born on the first day of a given month) and 0101 (i.e., indicating a voter born on January 1). When the date of birth ends in 01, we set the day of birth to missing. When the date of birth ends in 0101, we set the month and day of birth to missing.[36]

[34]Specifically, we drop voters indicated as "deceased" and who did not vote in the most recent general election before the year of the voter file.

[35]We consider groups of rows in which the middle initial is missing for one or more and the rest of these variables are identical to be duplicates.

[36]We lose information by setting the month or day of birth to missing for some people who were actually born on the first of the month or on January 1. However, there is no reliable way to determine whether a date of birth ending in 0101 actually corresponds to a January 1 birthday, or whether it indicates that the month and day of birth are missing.

2. The VBID uniquely identifies a row in each year snapshot, but when the same individual appears in multiple years, they typically appear with the same VBID. Therefore, we assume records that share a VBID are the same person, and assign them the same identifier. However, if TargetSmart assigned the same VBID to multiple rows where the first names, last names, and dates of birth are all different, or where the maximum difference in birth years is more than five years and the months and days of birth are different too, we break this link.

3. We group observations by ETID.

    - If at least one of the first names, last names, and dates of birth are the same among all observations sharing the same ETID, there is only one record per year, and the maximum age difference is less than or equal to five years, then we assign the same identifier to all these observations.
    - If not everyone in that set of observations shares either a first name, last name, or date of birth, we group them by name and date of birth and—as long as there is only one record per year—assign rows within each group the same identifier.

4. We group by first four letters of first name, last name, year of birth, and address. We require that each record has non-missing information for all of the grouping variables and that each record is unique within a year by these variables. If so, we assign these rows the same identifier.

5. We repeat the previous step using the following sets of grouping variables:[37]

    - First four letters of first name, last name, year of birth, and address.[38]
    - First name, last name, and address.
    - First name, last name, and date of birth.
    - First name, last name, and year and month of birth.
    - First name, last name, and year of birth.
    - First name, middle name, last name, and date of birth.
    - First name, middle name, last name, and year and month of birth.
    - First name, middle name, last name, and year of birth.

[37] If month and day of birth are not part of the grouping variables, we require that when non-missing, they do not conflict. We also require that middle initials, when non-missing, do not conflict, except if the set of grouping variables includes the address and year of birth, or if it includes the exact date of birth.

[38] We only de-duplicate records if no more than 15 people ever lived at the same address, and if the difference between the maximum and minimum years of birth is no more than five years.

C.2 Year-by-Year Voter Counts and Validation

Table C1 reports the number of unique Utah registered voters in our cleaned TargetSmart cross sections for each year from 2012 through 2021. Table C2 compares our cleaned counts, party-registration shares, and general-election turnout to statistics published by the Utah Lieutenant Governor's Office. Our registered-voter counts are systematically lower than state totals because our cleaning pipeline retains only voters with a reliable street-level geocode at a Utah address (see Section 3); individuals who first appear at a Utah address in a later year are excluded from earlier years' counts. Nevertheless, the share of voters registered with each party in our data closely tracks the state-reported shares.

Table C1: Year-by-Year Counts of Cleaned Utah Registered Voters

Year	Registered Voters
2012	1,091,892
2013	1,100,930
2014	1,098,110
2015	1,049,645
2016	1,158,595
2017	1,250,535
2018	1,222,348
2019	1,172,936
2020	1,276,559
2021	991,032

Figure C1 plots year-over-year party-switching rates between adjacent voter-file snapshots. Switching is heavily concentrated in the presidential election years and appears in the election-year snapshot itself: the 2015–2016 and 2019–2020 transitions show pronounced spikes (6.3% and 10.1% of linked voters, respectively), while every other transition is far quieter (at most 3.7%, and only 0.7% between the 2020 and 2021 files). Switches into the Republican Party account for the majority of the 2019–2020 spike (6.2% of linked voters), consistent with the registration surge ahead of the June 2020 Republican primary described in Section 2.4. Two implications follow for our design. First, presidential election years are indeed when registration changes occur, supporting our use of the 2012 and 2020 elections as measurement endpoints. Second, the election-year snapshots capture election-season affiliation changes: the 2020 file already reflects the 2020 registration surge, and the 2021 file adds post-general updates for voters not observed in 2020.

Table C2: Validation: Cleaned Voter File vs. Utah State-Published Statistics

	Registered		% Republican		% Democratic		Turnout	
Year	TS	State	TS	State	TS	State	TS	State
2015	1,049,645	1,412,099	45.6%	45.3%	9.5%	9.5%	—	—
2016	1,158,595	1,550,583	47.7%	46.5%	11.3%	11.4%	77.5%	82.0%
2017	1,250,535	1,538,748	46.3%	46.6%	11.8%	11.5%	—	—
2018	1,222,348	1,633,837	47.5%	46.4%	12.3%	12.4%	71.2%	—
2019	1,172,936	1,693,286	47.4%	45.0%	12.5%	13.0%	—	—
2020	1,276,559	1,857,861	52.1%	49.9%	14.2%	14.7%	87.6%	90.1%
2021	991,032	1,851,945	53.4%	49.3%	13.7%	14.7%	—	—

Notes: "TS" columns are computed from our cleaned TargetSmart Utah cross sections. "State" columns are from voter registration reports published by the Utah Lieutenant Governor's Office (registration and party counts) and from published turnout statistics for general elections. TargetSmart registered-voter counts are lower than state totals because our cleaning pipeline retains only voters with a reliable street-level geocode at a Utah address; inactive voters, non-geocoded voters, and individuals who first appear at a Utah address in a later year are excluded. For the turnout column, "TS" turnout for year Y is computed from the year $Y + 1$ TargetSmart snapshot, because the year-Y snapshot is taken mid-year and its vote-history field does not yet record the November Y general election. Turnout is reported only for even-year general elections. Dashes indicate data unavailable from the source.

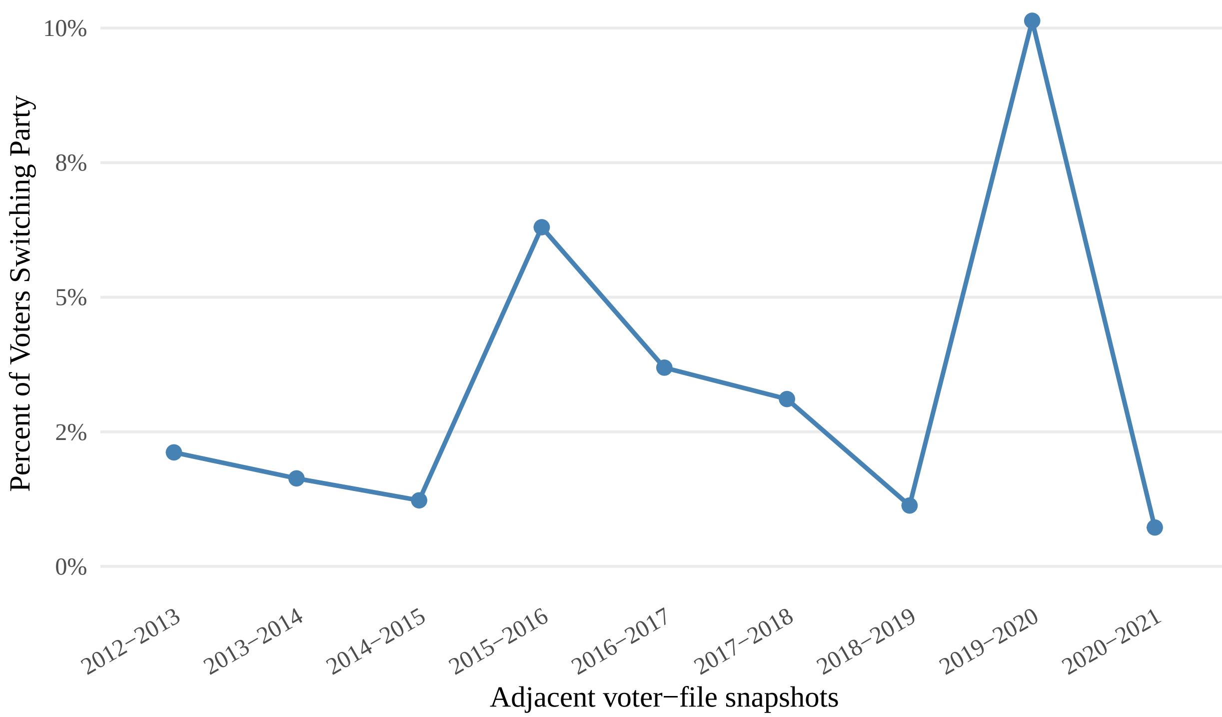

Figure C1: Year-over-year party-switching rates, 2012–2021. Each point is the share of voters observed in two adjacent voter-file snapshots whose three-way registration category (Republican, Democratic, Unaffiliated) differs between them; because Unaffiliated pools all minor parties, moves among minor parties do not count as switches. Between 0.9 and 1.1 million voters are linked in each adjacent pair.

C.3 Digitizing Ward Boundary Maps

We collected publicly available stake maps covering each of the approximately 4,000 Utah wards active between 2013 and 2017, and georeferenced, cleaned, and converted each map into a digital shapefile. To produce an initial set of polygons, we wrote a program that identifies boundary lines in the original map images by pixel color and converts the resulting non-boundary areas into polygon geometries. After this automated step, research assistants manually checked and corrected the shapefile output against the original PDF, adjusting boundary lines where the automated procedure had mis-traced a road, river, or fence line. As a final check, we manually reviewed each research assistant's work and approved the resulting shapefile before adding it to our ward panel.

C.4 Ward-Level Descriptive Statistics

Figure C2 reports ward-level distributions of land area, registered voter counts, and Republican and Democratic registration shares among 2013 wards. The weighted versions use each ward's number of registered voters as the histogram weight, so the weighted histograms approximate the distribution seen by a randomly chosen voter, while the unweighted histograms give equal weight to each ward. Note that the land-area distribution is heavily right-skewed: the median ward covers 0.5 km^2 and the interquartile range is [0.3, 1.8] km^2, but rural frontier wards span hundreds or thousands of square kilometers, so the histogram's x-axis is truncated at 25 km^2 for legibility.

Because registration shares sum to one, the cross-boundary gap in one party's share decomposes exactly into offsetting gaps in the other two. Figure C3 shows the structure of this trade-off. Panel (a) plots 2012 Democratic against Republican registration shares across the 3,869 Utah wards on the 2013 map: in a two-party electorate the points would lie on the dashed line (Dem $= 1-$ Rep), and the vertical distance below the line is the ward's Unaffiliated share. Panel (b) plots, for each of the 1,604 two-sided boundaries in our sample, the cross-boundary gap in Democratic share against the gap in Republican share, oriented so the Republican gap is positive, with two benchmarks: pure two-party substitution (slope -1) and pure substitution against the Unaffiliated share (slope 0). The voter-weighted fitted slope is -0.23: of a typical one percentage-point cross-boundary increase in Republican share, 0.23 points come at the expense of Democratic share and 0.77 points at the expense of Unaffiliated share. The reverse regression is asymmetric: a one-point increase in Democratic share comes with a 1.11-point decrease in Republican share—more than one-for-one, with Unaffiliated share slightly higher (by 0.11 points) on the more-Democratic side. Republican-share variation is thus largely variation against the Unaffiliated margin, while Democratic-

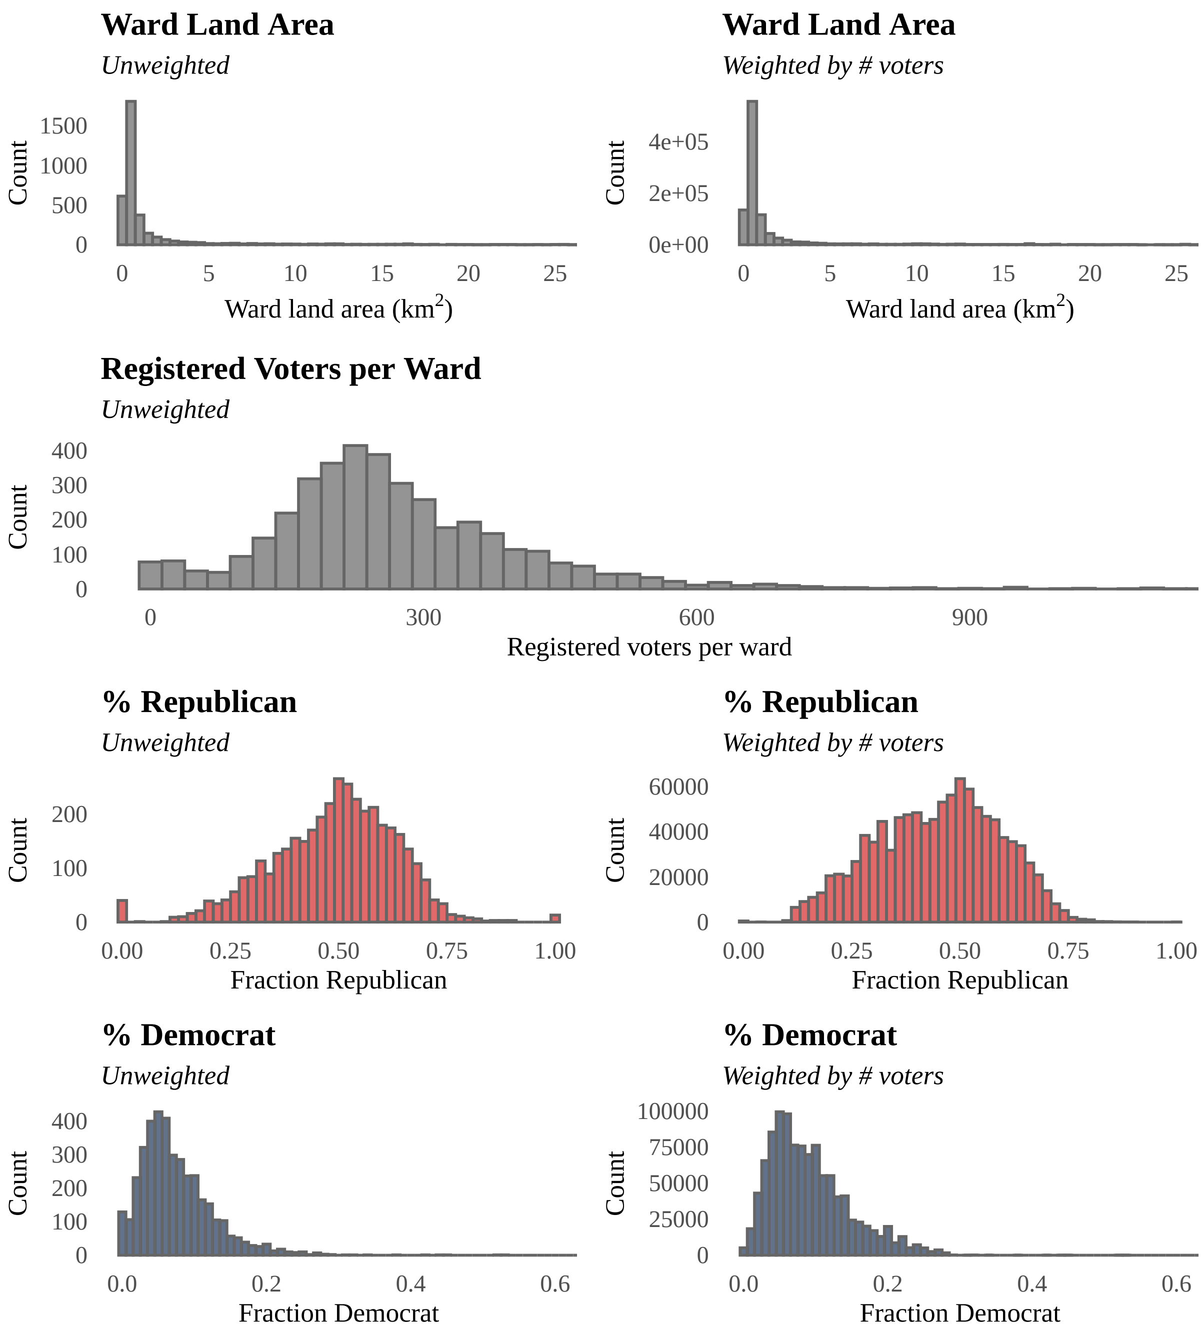

Figure C2: Ward-Level Distributions, 2013. Row 1 shows ward land area, truncated at 25 km^2. Row 2 shows the number of registered voters per ward (unweighted). Rows 3 and 4 show the fraction of each ward's registered voters affiliated with the Republican and Democratic parties, respectively.

share variation trades off with the Republican share.

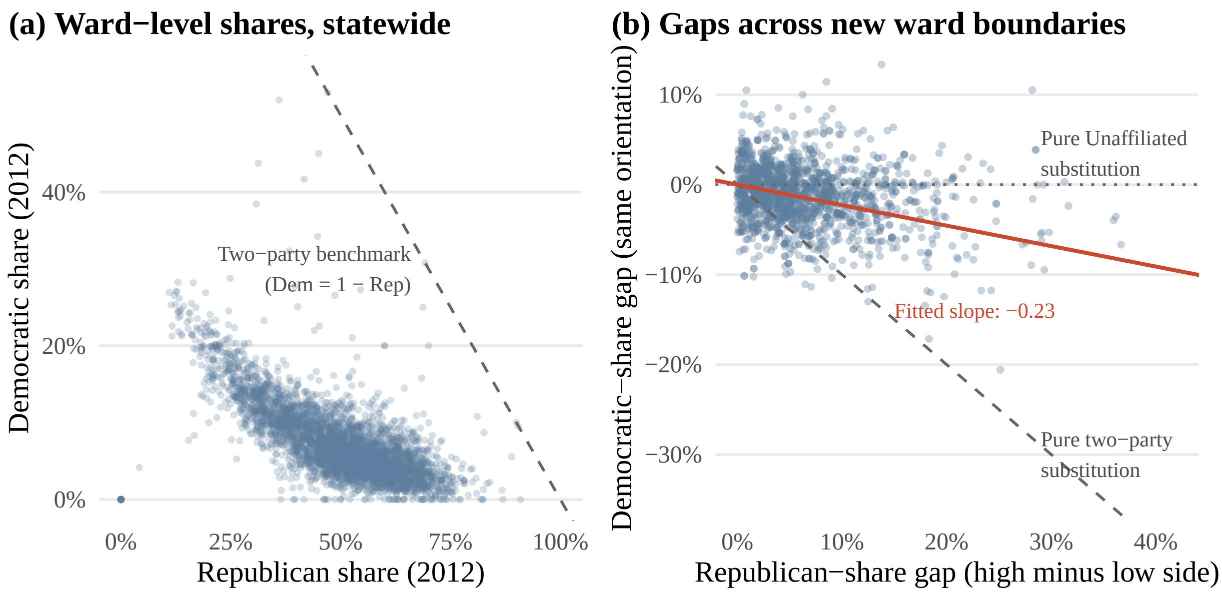

Figure C3: Relationship between ward-level Republican and Democratic shares and cross-boundary treatment gaps. Panel (a): ward-level 2012 Democratic vs. Republican registration shares for all Utah wards with at least ten registered voters (2013 boundaries); the dashed line is the two-party benchmark (Dem = 1− Rep), and vertical distance below it equals the ward's Unaffiliated share. Panel (b): cross-boundary gaps in 2012 shares for the 1,604 new boundaries with voters on both sides, oriented so the Republican gap is positive; the dashed and dotted lines mark pure two-party and pure-Unaffiliated substitution, and the solid line is the voter-weighted through-origin fit (slope −0.23; unweighted, −0.18). The reverse (Republican-on-Democratic) slopes are −1.11 weighted and −0.94 unweighted.

C.5 Linking CoreLogic Property Records to the Voter File

To attach property characteristics to voters, we link each CoreLogic owner record to at most one TargetSmart voter. We first standardize both sides: property situs addresses and voter residential addresses are reduced to a common normalized key (street address, city, and five-digit ZIP), and names are normalized with last name, first name, and middle initial stored separately. Voter addresses are pooled across the 2012–2021 cross sections, keeping the years over which each voter–address pair is observed, so a property can match a voter at any of the voter's address spells.[39]

Matching proceeds in two passes. The first pass joins owners to voters on the exact

[39] We allow links between voter files and property files of different years because of the fact that voters do not always update their registration address at the time of a move.

standardized address key and exact normalized last name. For owners still unmatched, the second pass joins on five-digit ZIP and exact last name and requires the property and voter geocodes to lie within 100 meters. In both passes, candidate pairs must also agree on first name: we score first-name pairs 100 for exact matches, 99 for known nickname pairs (e.g., "Bill"/"William"), 98 when one name contains the other (with at least three characters), and by Jaro–Winkler similarity otherwise, retaining candidates scoring at least 87. Among an owner's surviving candidates we keep a single best voter, preferring exact-address matches, then agreement in middle initial, then higher first-name scores, then smaller distances, then the address spell closest in time to the CoreLogic vintage. Finally, each voter may be linked to only one property; when a voter survives as the best match for owner records on multiple properties, we keep the highest-quality link by the same criteria. The large majority of matched owners link in the exact-address pass, with the remainder recovered by the 100-meter fallback.

C.6 Linking CoreLogic Property Records to HMDA Income

To recover applicant income for CoreLogic homeowners, we link CoreLogic mortgage origination records to Home Mortgage Disclosure Act (HMDA) loan-application records, which report applicant income but not property address (Bayer et al., 2025). We begin by spatially joining each CoreLogic mortgage record to its census block (and thus to its census tract) using parcel-level coordinates. We then aggregate HMDA loan records to the tract level using their reported census tract. Within each (tract, year) cell, we merge the two datasets on loan amount, rounded to the nearest \$1,000, and on lender name. We clean lender names of common abbreviations (e.g., "Bk" $\rightarrow$ "Bank", "Fcu" $\rightarrow$ "Federal Credit Union") and then fuzzy-match them using token-sort ratios, retaining only pairs with a full-name score of at least 67 and a first-word score of at least 70. To avoid spurious many-to-one matches, we further restrict to CoreLogic mortgages and HMDA loans that are uniquely identified within (tract, amount, lender). Because HMDA began rounding loan amounts to the nearest \$5,000 in 2018, we implement this linkage only for origination years 2012–2017. The resulting match rate is low—for example, we link roughly 27% of 2012 CoreLogic mortgage originations to HMDA income—so we omit income from our main analyses and use it only to describe the linked-property sample in Appendix C.8.

C.7 School District Data Availability

Table C3 lists all 41 Utah school districts, showing for each whether we obtained elementary-school boundary shapefiles from the district and the number of elementary schools reported either by the district directly or counted from the shapefile.

Table C3: Utah School Districts: Elementary Boundary Data Availability

School district	Data availability	Elementary schools
Davis	2012-2025	62
Granite	2012-2025	62
Alpine	2012-2025	53
Jordan	2013-2025	33
Nebo	2025	30
Canyons	2013-2025	29
Salt Lake	2012-2025	29
Weber	2015-2025	29
Washington	2020-2025	28
Cache	2018-2025	17
Tooele	2022	16
Ogden	2010-2025	15
Provo	2024	15
Box Elder	2024	12
Iron	2024	9
Duchesne	Not available	8
Murray	2013-2025	7
Uintah	2012-2025	7
Emery	Not available	6
Logan	2012-2025	6
San Juan	Not available	6
Carbon	2012-2025	5
Garfield	Not available	5
North Sanpete	Not available	5
Sevier	Not available	5
Wasatch	2025	5
Millard	Not available	4
Park City	Not available	4
Juab	Not available	3
South Sanpete	Not available	3
Beaver	Not available	2
Daggett	Not available	2
Kane	Not available	2
Morgan	Not available	2
Piute	Not available	2
Rich	Not available	2
South Summit	Not available	2
Tintic	Not available	2
Wayne	Not available	2
Grand County	2012-2025	1
North Summit	2012-2025	1

Notes: "Data availability" lists the range of years for which the district provided elementary-school boundary shapefiles; "Not available" indicates the district did not supply shapefiles. "Elementary schools" is the number of attendance zones in the 2012 shapefile (for districts that supplied one) or the count reported directly by the district (for those that did not).

C.8 Sample Comparison Summary Statistics

This subsection reports summary statistics for the analysis sample and balance-style comparisons of subsamples of our voter and property data. Table C4 reports summary statistics separately for each baseline-party subsample, the samples our regressions use, including each column's pre-attrition universe, boundary count, linkage rates, demographics, mobility, and 2020 outcomes. Table C5 compares the full 2012 TargetSmart cross section, the subset of voters observed in both 2012 and 2020 or 2021 (our panel sample), and the further subset used in our main regressions (panel voters living within 400m of a new ward boundary with voters on both sides). Tables C6 and C7 report the voter-side and property-side linkage-balance checks cross-referenced in Section 3. Table C8 compares the main analysis sample to the placebo sample used in the future-boundaries exercise of Section 5.1.

Boundary counts differ across exhibits because different analyses impose different requirements: the party-specific subsamples retain only boundaries with voters of that party on both sides (Table C4); the missing-religion verification (Appendix E) uses the boundaries with both flanking wards observed in the 2012 voter file; and the randomization inference (Appendix F.7) flips the roughly 1,440 boundaries with estimable exposure contrasts on both sides.

C.9 Scraping the Meetinghouse Locator

To track ward boundary changes occurring after our 2013–2017 digitized map window, we built a scraping pipeline that queries the LDS Church's Meetinghouse Locator tool, which returns the assigned ward for any Utah address submitted to it. We run the pipeline once per quarter, beginning in Spring 2023. For each address and quarter we record the returned ward name and identifier as well as the name of its bishop (lay pastor). We then detect boundary changes by comparing each address's ward assignment across consecutive quarters: addresses whose ward identifier changes between quarters mark a point on a new boundary line, and we reconstruct the full boundary by connecting such points within each stake. Boundary changes observed through Spring 2026 form the basis for our placebo sample in Section 3.5.

Table C4: Summary Statistics: Analysis Sample by Baseline Party

	Republicans	Democrats	Unaffiliated
2012 voters near an eligible boundary[a]			
Number of voters	82,307	12,200	75,654
... not observed in 2020/2021 (attrition)	26.1%	31.9%	34.3%
Analysis sample[b]			
Number of voters	60,795	8,314	49,684
New boundary lines[a]	1,439	876	1,450
Linked to CoreLogic single-family home	49.3%	40.5%	43.1%
2012 characteristics			
Age	46.3 (16.3)	44.9 (16.1)	42.7 (16.0)
Male	48.8%	44.2%	51.6%
White	97.6%	87.6%	93.6%
Rural	11.2%	5.6%	9.1%
Property value[c]	$274k [215k, 365k]	$247k [195k, 313k]	$253k [201k, 328k]
Year built[c]	1997 [1980, 2004]	1994 [1972, 2002]	1994 [1976, 2003]
Mobility, 2012–2020/21			
Moved[d]	52.2%	57.5%	57.4%
... out of Utah[d]	9.5%	15.3%	12.9%
Political outcomes			
Registered Republican in 2020	86.2%	12.6%	29.4%
Registered Democratic in 2020	1.8%	70.6%	10.6%
Registered Unaffiliated in 2020	12.0%	16.8%	59.9%
Voted in 2012 primary	40.1%	4.7%	2.7%
Voted in 2012 general	88.2%	78.8%	73.4%
Voted in 2020 primary[e]	64.2%	26.6%	25.1%
Voted in 2020 general[e]	87.8%	78.2%	78.2%

Notes: [a] An eligible boundary for each column is a newly drawn 2013–2017 line with voters of that baseline party within 400m on both sides, so both voter and boundary counts differ across columns and columns do not sum to the pooled analysis sample (Section 3.4). [b] All rows from this block on are computed over the analysis sample: voters near an eligible boundary who are observed in 2020 or 2021. [c] Among voters linked to a CoreLogic single-family home; property characteristics are fixed at their baseline (2012) values. Cells report the median [interquartile range]; Age reports the mean (standard deviation). [d] Share of analysis-sample voters whose 2020/2021 address is more than 100m from their 2012 address, and the subset observed at an address outside Utah. Movers remain in the sample; treatment is defined at 2012 residence. Voters who left and cannot be linked appear as attrition instead; see Section 6.4. [e] 2020 votes are recorded in the 2021 voter-file snapshot (see Section 3.1).

Table C5: Summary Statistics: Sample Comparison

	Full 2012 XS	Panel	Analysis	p
N	1,091,892	753,451	121,876	—
Political				
Fraction Republican	44.9%	47.4%	50.9%	***
Fraction Democrat	9.1%	9.0%	7.7%	***
Fraction Unaffiliated	46.0%	43.6%	41.4%	***
Voted 2012 general	74.4%	81.6%	81.4%	***
Voted 2012 primary	20.6%	22.4%	21.9%	***
Demographics				
Age	46.9 (18.3)	45.9 (16.7)	44.7 (16.3)	***
Fraction male	49.3%	49.4%	49.6%	*
Fraction white	95.0%	95.5%	95.2%	**
Fraction rural	10.7%	10.8%	9.9%	***
Linked to elementary school	94.9%	94.9%	97.3%	***
Property				
Fraction homeowner	40.4%	44.6%	46.1%	***
Home value (USD)	$262k [203k, 350k]	$267k [206k, 356k]	$263k [208k, 349k]	***
Building size (sq ft)	1,809 [1,373, 2,492]	1,836 [1,394, 2,528]	1,831 [1,408, 2,487]	***
Land size (sq ft)	10,019 [7,841, 13,852]	10,019 [7,841, 13,939]	10,019 [7,841, 13,939]	***
Year built	1986 [1970, 1999]	1987 [1971, 1999]	1996 [1978, 2003]	***
Bedrooms	3 [3, 4]	3 [3, 4]	3 [3, 4]	***
Bathrooms	2.0 [2.0, 3.0]	2.0 [2.0, 3.0]	2.0 [2.0, 3.0]	***
Fraction single-story	76.3%	75.3%	72.3%	***
HMDA income (USD)	$104k [74k, 146k]	$105k [75k, 148k]	$102k [74k, 143k]	*

Notes: "Full 2012 XS" is the cleaned 2012 TargetSmart cross section; "Panel" is the subset of voters also observed in 2020 or 2021; "Analysis" is the further subset used in our main DiD regressions — voters in the panel living within 400m of a new ward boundary between 2013 and 2017 on a boundary with voters on both sides. p column tests "Analysis" vs. "Full 2012 XS". At sample sizes above 10^5, nearly all mean differences are statistically significant; readers should interpret magnitude alongside significance. Significance: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Continuous variables reported as "mean (SD)"; skewed variables reported as "median [25th, 75th percentile]"; binary indicators reported as percentages. Property-row statistics are computed over CLIP-non-missing observations within each column.

Table C6: Summary Statistics: Voters Linked vs. Unlinked to CoreLogic (Analysis Sample)

	Linked to CoreLogic	Not linked	p
N	56,214	65,662	—
Political			
Fraction Republican	54.3%	47.9%	***
Fraction Democrat	6.9%	8.5%	***
Fraction Unaffiliated	38.8%	43.7%	***
Voted 2012 general	88.4%	75.0%	***
Voted 2012 primary	25.1%	19.0%	***
Demographics			
Age	48.1 (14.5)	41.9 (17.1)	***
Fraction male	50.1%	49.2%	**
Fraction white	95.8%	94.7%	***
Fraction rural	7.9%	11.6%	***
Linked to elementary school	98.9%	96.0%	***

Notes: Analysis sample only. Voters are split by whether their record was successfully linked to a CoreLogic property. Property variables are omitted because they are missing by construction for unlinked voters.

Table C7: Summary Statistics: Properties Linked vs. Unlinked to Voter File

	Linked to voter file	Not linked	p
N	357,240	129,820	—
Demographics			
Fraction rural	9.2%	—	
Property			
Home value (USD)	$257k [197k, 343k]	$215k [168k, 284k]	***
Building size (sq ft)	1,764 [1,340, 2,424]	1,534 [1,180, 2,018]	***
Land size (sq ft)	9,583 [7,405, 13,504]	8,276 [5,663, 11,326]	***
Year built	1987 [1969, 2000]	1985 [1964, 2001]	***
Bedrooms	3 [3, 4]	3 [3, 4]	***
Bathrooms	2.0 [2.0, 3.0]	2.0 [2.0, 3.0]	***
Fraction single-story	76.0%	75.9%	
HMDA income (USD)	$98k [68k, 141k]	$77k [53k, 113k]	***

Notes: Properties are 2012 CoreLogic owner-occupied single-family homes in Utah, deduplicated to one row per CLIP. Columns split properties by whether at least one owner was successfully linked to a TargetSmart voter record. Demographic and political variables are omitted because they come from the voter-file side of the join.

Table C8: Summary Statistics: Main vs. Placebo Analysis Sample

	Main analysis	Placebo analysis	p
N	121,876	103,335	—
Political			
Fraction Republican	50.9%	43.8%	***
Fraction Democrat	7.7%	10.3%	***
Fraction Unaffiliated	41.4%	46.0%	***
Voted 2012 general	81.4%	80.8%	***
Voted 2012 primary	21.9%	21.1%	***
Demographics			
Age	44.7 (16.3)	45.5 (16.7)	***
Fraction male	49.6%	49.3%	
Fraction white	95.2%	94.2%	***
Fraction rural	9.9%	5.4%	***
Linked to elementary school	97.3%	98.4%	***
Property			
Fraction homeowner	46.1%	44.6%	***
Home value (USD)	$263k [208k, 349k]	$244k [191k, 314k]	***
Building size (sq ft)	1,831 [1,408, 2,487]	1,716 [1,302, 2,330]	***
Land size (sq ft)	10,019 [7,841, 13,939]	9,148 [7,405, 11,761]	***
Year built	1996 [1978, 2003]	1983 [1966, 1998]	***
Bedrooms	3 [3, 4]	3 [3, 4]	***
Bathrooms	2.0 [2.0, 3.0]	2.0 [2.0, 3.0]	***
Fraction single-story	72.3%	77.7%	***
HMDA income (USD)	$102k [74k, 143k]	$96k [68k, 132k]	***

Notes: "Main analysis" is the sample used in `analysis_did_new` (model m1): voters in our panel living within 400m of a 2013–2017 new ward boundary with voters on both sides. "Placebo analysis" is the corresponding sample used in `placebo_did_new` (model m1): voters in the panel living within 400m of a 2023q2–2026q2 future ward boundary with voters on both sides.

D A Conceptual Model for Interpreting the Estimand

This appendix develops a friendship-formation model that gives the main-text specification, Equation (1), a structural interpretation (Appendices D.1 and D.2). The model assumes a linear response at the level of the friend circle, which we discuss substantively in Section 6.2, and takes up two questions given that assumption: what structural object the regression coefficient β recovers, and how that mapping depends on the number of friends a member has; and whether, when a boundary change severs existing friendships and forces new ones, the influence of retained and newly formed ties can be separately identified from ward rosters alone. The friend-level influence parameter ρ derived below is the object that Appendix E denotes β_{true}. We frame the derivation in terms of the Republican share; the argument is identical for any party.

Table D1 previews the result. Under random friending, β recovers the friend-circle conformity parameter ρ, the expected partisan movement from an all-out-party to an all-co-party friend circle, regardless of the number of friends. Homophily and a broadcast channel change the interpretation but not the specification: β becomes the sum of a homophily-attenuated dyadic component and a broadcast component. In every variant β remains the causal effect of ward peer composition on partisan change at the margin the design isolates, which is the object relevant to the reassignment counterfactual; the additional structure is needed only to recover deeper friend-level primitives. The disruption model of Appendix D.2 shows that the same rosters, linked across a boundary change, can in principle separate the influence of retained and newly formed ties through regression (18).

Table D1: Interpretation of the estimated coefficient under successive model variants. β is the coefficient of the main estimating equation (1), estimated within baseline-party subsamples.

Model variant	Interpretation of the estimated coefficient(s)
Random friending, m friends	$\beta = \rho$: expected movement from an all-out-party to an all-co-party friend circle, invariant to m; per-friend effect ρ/m.
Homophilous friending (weight λ)	$\beta = \rho\,(1-\lambda)$: ward-level effect, a lower bound on the friend-level ρ.
Broadcast channel (ρ_g) alongside friends (ρ_f)	$\beta = \rho_f(1-\lambda) + \rho_g$: sum of dyadic and broadcast influence; separated only by variation loading differently on the channels (Section 7.4).
Friendship disruption ($\rho_{\text{old}}, \rho_{\text{new}}$)	Regression (18) identifies the two rates where disruption intensity and newcomer distinctiveness vary; otherwise β recovers the disruption-weighted blend.

D.1 A Model of Friend Formation and Peer Influence

Setup. Consider a ward w with n_w members. Each member i has a baseline (2012) Republican indicator $R_i^0 \in \{0,1\}$. Define the **leave-one-out ward Republican mean**

$$\bar{Z}_w^{-i} \;\equiv\; \frac{1}{n_w - 1} \sum_{j \in w,\, j \neq i} R_j^0, \tag{4}$$

the baseline Republican share among i's ward-mates. Between baseline and endline, each member draws $m \geq 1$ friends uniformly at random, without replacement, from the other $n_w - 1$ members of the ward. Let $\bar{R}_i^f \equiv \frac{1}{m} \sum_{k=1}^{m} R^0_{f_k(i)}$ denote the baseline Republican share of i's friend circle. Peer influence operates on the friend circle's baseline composition, scaled by $\rho \in [0,1]$: the member's endline (2020) affiliation moves a fraction ρ of the way from their own baseline toward the friend-circle share, plus an idiosyncratic shock,

$$R_i^1 \;=\; R_i^0 \;+\; \rho\left(\bar{R}_i^f - R_i^0\right) \;+\; \nu_i, \qquad \mathbb{E}[\nu_i \mid R_i^0,\, \bar{R}_i^f,\, w] = 0, \tag{5}$$

so that $\rho = 0$ is no peer influence and $\rho = 1$ is full conformity to the friend circle. For $m = 1$ the linear form in (5) is without loss of generality—any function of a binary variable is affine—while for $m > 1$ it is the substantive linearity assumption discussed in Section 6.2. Differencing,

$$\Delta R_i \;\equiv\; R_i^1 - R_i^0 \;=\; \rho\left(\bar{R}_i^f - R_i^0\right) + \nu_i. \tag{6}$$

Integrating out the unobserved friends. The researcher never observes who is friends with whom. The only object on the right-hand side of (6) that depends on the friend draw is $\bar{R}_i^f$, and because the m friends are drawn uniformly from i's ward-mates, each ward-mate is equally likely to be in the friend circle, so

$$\mathbb{E}\Big[\bar{R}_i^f \mid w,\, i\Big] \;=\; \bar{Z}_w^{-i}. \tag{7}$$

Under random friending, then, the LOO ward mean is not merely a proxy for the unobserved peer measure: it is exactly the expectation of the friend circle's composition, for any m. Taking the expectation of (6) conditional on the ward and the member's own baseline party,

$$\mathbb{E}\big[\Delta R_i \mid w,\, R_i^0\big] \;=\; \rho\, \bar{Z}_w^{-i} \;-\; \rho\, R_i^0. \tag{8}$$

The regression and its interpretation. Estimating separately by baseline party fixes R_i^0 at a constant $r \in \{0,1\}$, so within each baseline-party subsample the conditional expectation of the partisan change is an exact linear function of the LOO ward mean,

$$\mathbb{E}\big[\Delta R_i \mid \bar{Z}_w^{-i}\big] \;=\; (-\rho\, r) \;+\; \rho \cdot \bar{Z}_w^{-i}, \qquad i : R_i^0 = r. \tag{9}$$

Because the regressor $\bar{Z}_w^{-i}$ is a function of predetermined baseline affiliations only, it is orthogonal to the endline shock ν_i, so the population OLS slope equals ρ with no further assumption. Importantly, the interpretation of ρ does not depend on the unobserved number of friends: ρ is the expected partisan movement induced by shifting a member's entire friend circle from containing no co-partisans to consisting entirely of co-partisans, whatever the circle's size. The marginal effect of converting a single friend is ρ/m, so the number of friends governs how influence is divided across ties rather than the total, and β recovers the total. This is the mapping we use when discussing magnitudes.

Why splitting by baseline party is necessary. Pooling across parties omits the $-\rho R_i^0$ term in (8), which is positively correlated with $\bar{Z}_w^{-i}$ because Republicans live in more-Republican wards, biasing the pooled slope toward zero. Splitting also makes the LOO mean identical for every member of a given ward-by-baseline-party cell—removing one Republican from the roster yields the same numerator and denominator for all Republicans in the ward—which is the property exploited in the main text (Section 4.1).

From the cross section to the boundary design. Equation (9) is a cross-sectional linear-in-means regression; the main-text specification embeds the same relationship in the boundary discontinuity. A redrawn boundary reassigns i from a shared pre-boundary ward to one of two new wards A or B, changing the pool from which friends are drawn. Applying (8) on each side and differencing across the boundary, the own-baseline term $-\rho\, R_i^0$ and any pre-boundary common component cancel, leaving

$$\mathbb{E}[\Delta R_i \mid \text{new}(i) = A] - \mathbb{E}[\Delta R_i \mid \text{new}(i) = B] \;=\; \rho\left(\bar{Z}_A^{-i} - \bar{Z}_B^{-i}\right) \tag{10}$$

within a baseline-party subsample. Since the treatment in Equation (1) is the cross-boundary difference in LOO peer composition, switched on for the high-exposure side ($D_{ib} \times I_b$, with the side indicator δD_{ib} entering separately), the slope coefficient there is exactly $\beta = \rho$; the boundary fixed effects μ_b absorb the per-boundary intercept.

The two departures from random friending that follow leave the estimating equation unchanged but alter what β measures, and each limits how far the estimate can be pushed structurally.

Homophilous friending. The mapping $\beta = \rho$ rests on friends being drawn at random *within* the ward. If instead members preferentially befriend co-partisans among their ward-mates, the expected friend-circle composition tilts toward own type, $\mathbb{E}[\bar{R}_i^f \mid w, R_i^0 = r] = \lambda\, r + (1-\lambda)\, \bar{Z}_w^{-i}$ for a homophily weight $\lambda \in [0,1]$, and the regression slope becomes $\beta = \rho\,(1-\lambda)$. The design still identifies the causal effect of *ward* composition—the object relevant to the reassignment counterfactual—but the friend-level ρ is then bounded below by β, and recovering it exactly would require knowing λ.

A broadcast channel. Some influence is plausibly broadcast rather than dyadic: a comment over the pulpit or in Sunday School reaches every member, not only one's friends. Letting ρ_f scale the friend channel and ρ_g a broadcast channel operating on the full ward composition, $R_i^1 = R_i^0 + \rho_f(\bar{R}_i^f - R_i^0) + \rho_g(\bar{Z}_w^{-i} - R_i^0) + \nu_i$, the two channels collapse onto the same regressor and the slope becomes

$$\beta \;=\; \rho_f\,(1-\lambda) \;+\; \rho_g. \tag{11}$$

From the ward-mean regression alone the channels are observationally equivalent: β recovers their sum. Distinguishing them requires variation that loads on them differently, which is what the same- versus opposite-demographic decomposition of Section 7.4 provides—a pulpit broadcast is demographically blind, whereas dyadic influence flows through homophilous friendships. The concentration of our effects among demographically similar peers is therefore evidence that the dyadic channel is active, though it does not pin down the two rates separately.

D.2 Boundary Changes as Peer-Network Disruption

The baseline model holds the friend circle fixed from baseline to endline. A boundary change does more than shift the composition of the pool from which friends are drawn: it can sever an existing friendship and force a new one, and a freshly formed tie need not transmit influence at the same rate as a long-standing one. This subsection lets the boundary change disrupt friendships directly, splitting the conformity parameter into a retained-tie rate ρ_{old} and a new-tie rate ρ_{new}, and derives the estimating equation that separately identifies the two from ward rosters linked across the redrawing. For clarity we return to a single friend ($m = 1$); the m-friend generalization proceeds tie by tie and delivers the same regression.

Timeline. Friending and a first round of influence occur in a pre-observation period exactly as in Appendix D.1, resulting in the baseline (2012) party R_i^0 we observe. A boundary change then **removes** a set $\mathcal{M}_w$ of M_w members from ward w, **retains** the remaining $n_w - M_w$ (the "stayers" $\mathcal{S}_w$, among whom is i), and **adds** a set $\mathcal{N}_w$ of N_w newcomers, leaving a post-change ward of $n_w - M_w + N_w$ members. The change disrupts i's friendship only if i's friend was among the removed. Since the pre-period friend was uniform over i's $n_w - 1$ ward-mates,

$$p_w \;\equiv\; \Pr(\text{friend removed}) \;=\; \frac{M_w}{n_w - 1}. \tag{12}$$

If the friend stays (probability $1 - p_w$), i keeps them. If the friend is removed, i draws a replacement uniformly from the post-change ward—the $n_w - M_w - 1$ remaining stayers other than i together with the N_w newcomers, since either is an equally available new acquaintance.

Old versus new ties. A retained friendship and a freshly drawn one transmit influence at possibly different rates:

$$\Delta R_i \;=\; \begin{cases} \rho_{\text{old}}\big(R^0_{f(i)} - R^0_i\big) + \nu_i, & \text{friend retained,} \\ \rho_{\text{new}}\big(R^0_{g(i)} - R^0_i\big) + \nu_i, & \text{friend redrawn,} \end{cases} \qquad \mathbb{E}[\nu_i \mid \cdot] = 0, \tag{13}$$

where $f(i)$ is the retained friend and $g(i)$ the replacement. The rate ρ_{new} attaches to the newness of the relationship, not to whether the friend is new to the ward: a long-standing ward-mate whom i befriends only after the change is a new tie too. The redraw pool therefore mixes remaining stayers and newcomers, and all of it carries ρ_{new}.

Integrating out the unobserved friend and redraw. As in the baseline, the researcher observes neither $f(i)$ nor $g(i)$, and need not: only the disruption probability and the composition of each pool enter the conditional expectation. Define the leave-one-out stayer mean, the newcomer mean, and the redraw-pool mean as their size-weighted mixture,

$$\bar{Z}^{S,-i}_w = \frac{1}{n_w - M_w - 1} \sum_{j \in \mathcal{S}_w,\, j \neq i} R^0_j, \qquad \bar{Z}^N_w = \frac{1}{N_w} \sum_{j \in \mathcal{N}_w} R^0_j, \tag{14}$$

$$\bar{Z}^{\text{pool},-i}_w \;=\; (1 - q_w)\, \bar{Z}^{S,-i}_w \;+\; q_w\, \bar{Z}^N_w, \qquad q_w \;\equiv\; \frac{N_w}{n_w - M_w - 1 + N_w}, \tag{15}$$

where q_w is the newcomers' share of the redraw pool. Conditional on a retained tie the friend is uniform over the stayers, so its expected baseline type is $\bar{Z}^{S,-i}_w$; conditional on a redraw the replacement is uniform over the pool, so its expected type is $\bar{Z}^{\text{pool},-i}_w$. The removed members' own composition never appears: it governs who left, but departed members influence no one.

Taking expectations over both the friend draw and the redraw,

$$\mathbb{E}\big[\Delta R_i \mid w,\, R_i^0 = r\big] \;=\; (1 - p_w)\,\rho_{\text{old}}\,\big(\bar{Z}_w^{S,-i} - r\big) \;+\; p_w\,\rho_{\text{new}}\,\big(\bar{Z}_w^{\text{pool},-i} - r\big). \tag{16}$$

The estimating equation. Equation (16) is linear in two covariates built entirely from observed counts and group party shares,

$$X_{ib}^{\text{old}} \;\equiv\; (1 - p_w)\big(\bar{Z}_w^{S,-i} - R_i^0\big), \qquad X_{ib}^{\text{new}} \;\equiv\; p_w\big(\bar{Z}_w^{\text{pool},-i} - R_i^0\big), \tag{17}$$

so that, embedded in the boundary design and estimated within baseline-party subsamples (fixing r, exactly as the main specification does), the model implies the regression

$$\Delta Y_{ib} \;=\; \rho_{\text{old}}\,X_{ib}^{\text{old}} \;+\; \rho_{\text{new}}\,X_{ib}^{\text{new}} \;+\; \mu_b \;+\; f(r_{ib}) \;+\; \varepsilon_{ib}, \tag{18}$$

the direct analog of Equation (1) with the single treatment replaced by the pair $(X^{\text{old}}, X^{\text{new}})$. Constructing the regressors requires only observable data. The voter file linked to the old and new boundary maps gives, for every ward, the counts M_w, N_w, n_w and the baseline partisan shares of the removed, retained, and added groups. No friendship data enters—the unobserved tie is integrated out through the disruption probability p_w, just as the friend circle was integrated out in the baseline.

E Adjusting for Missing Religion Data

Our main difference-in-discontinuities specification treats every Utah voter as if they were LDS, both as a subject of ward peer influence and as a contributor to the leave-one-out (LOO) ward peer-composition measure. In reality only a fraction of Utah voters are LDS, and that fraction varies systematically by party affiliation and by geography.

This appendix investigates how this measurement issue affects the interpretation of the published estimates, deriving re-scaling factors which can be applied to the reported coefficients to recover the underlying causal effect on LDS members. Appendix E.1 calibrates the party-conditional probability of being LDS from the Utah Colleges Exit Poll (UCEP), a large in-person survey of Utah voters conducted at polling places, pooling the 2002, 2006, 2012, 2014, and 2016 waves. We then derive closed-form expressions for the implied attenuation factor (Appendix E.2), validate them against Monte Carlo simulations across a sequence of stylized scenarios (Appendix E.3), and verify the implied corrections against the actual Utah ward-boundary configurations in our analysis sample (Appendix E.4). Appendix E.5 then constructs a ward-level measure of local LDS density, and Appendix E.6 uses it to test a prediction of the attenuation logic: effects should be stronger where LDS density is higher.

E.1 Variation in Probability LDS by Individual Characteristics

We use the Utah Colleges Exit Poll (UCEP), conducted by BYU's Center for the Study of Elections and Democracy at polling places across the state, to characterize how the probability of being LDS varies with individual characteristics. We use the UCEP because of its large sample size relative to other surveys like the CES.

Table E1 reports the share of UCEP respondents identifying as LDS in the 2012 wave, broken down by column party (Republican, Democrat, Independent, or all respondents) and by row characteristic (party affiliation, sex, age bucket, race). The party block shows the marginal probability of being LDS within each party-registration group: roughly 84% of Utah Republicans, 24% of Utah Democrats, and 52% of Utah Independents identify as LDS. Within parties, the probability is broadly stable across sex and age. Women are slightly more likely than men to be LDS among Republicans (0.86 vs. 0.81) and Independents (0.54 vs. 0.51), and equally likely among Democrats. Age shows no common gradient: among Democrats the LDS share rises from 0.19 in the 18–29 group to 0.34 among those 66 and older, while among Republicans it is essentially flat and slightly U-shaped (0.87, 0.84, 0.80, 0.88 across the four age groups) and among Independents it varies within a five-point band.

These party-conditional shares are the inputs to the missing-religion adjustment of Appendix E, which uses the rounded values $p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, and $p_{\text{LDS}}(D) = 0.25$.

Two further features of the UCEP data support treating party and geography as additive shifters, which is the structure Appendix E imposes. First, within each row of Table E1 the column-to-column differences in LDS share are similar across rows, consistent with no large interactions between party and demographics. Second, across Utah counties each party-conditional LDS share tracks the county's unconditional LDS share with a slope close to one (Pearson r of 0.82 for Republicans, 0.86 for Democrats, and 0.94 for Independents), consistent with county acting as an additive shifter on top of party.

Table E1: Estimated probability LDS by individual characteristics, Utah Colleges Exit Poll 2012.

	Republicans	Democrats	Independents	Overall
Party affiliation				
Republican	0.84	—	—	0.84
Democrat	—	0.24	—	0.24
Independent	—	—	0.52	0.52
Sex				
Male	0.81	0.24	0.51	0.63
Female	0.86	0.24	0.54	0.64
Age				
18–29	0.87	0.19	0.55	0.64
30–45	0.84	0.26	0.55	0.65
46–65	0.80	0.24	0.49	0.60
66+	0.88	0.34	0.53	0.71
Race				
White	0.85	0.27	0.55	0.67
Non-White	0.69	0.14	0.39	0.38

Notes: Cells report the share of respondents identifying as LDS within the row × column subgroup, UCEP 2012 wave only (unweighted). Sample sizes per party column: Republicans $n = 7,485$; Democrats $n = 3,143$; Independents $n = 1,486$; Overall $n = 22,797$. "—" indicates a cell that is empty by construction (the row and column conditions are contradictory). LDS classification follows the four-level CSED `relig_index` (any non-"Not LDS" response).

E.2 Conceptual Model and Closed-Form Solutions

For simplicity we begin by considering a single ward boundary line drawn to divide one previous ward into two equal halves. Once we have established this intuition, we consider the more realistic case of multiple ward boundaries across Utah.

Setup. A single ward boundary line divides N "boundary" individuals 50/50 between two newly drawn wards A and B; before the boundary change, all N shared a common pre-boundary ward, which we denote pre. Let $\text{new}(i) \in \{A, B\}$ denote voter i's post-boundary ward assignment. Each new ward also contains a large mass of "other" voters drawn from a ward-specific party mix $\pi_w = (\pi_{w,R}, \pi_{w,U}, \pi_{w,D})$, which is what gives each ward a distinct partisan composition. Each voter i has a baseline party $p_i \in \{R, U, D\}$ and an LDS indicator $\mathbf{1}\{\text{LDS}_i\}$ drawn Bernoulli with party-conditional probability $p_{\text{LDS}}(p_i)$.

For each ward $w \in \{A, B, \text{pre}\}$ and party $X \in \{R, D\}$ we define two peer-composition measures: $Z^{\text{obs}}_{X,w}$, the leave-one-out (LOO) fraction of party-X voters in ward w among all voters; and $Z^{\text{LDS}}_{X,w}$, the LOO fraction of party-X voters in ward w among LDS members only. Only LDS voters respond to their ward's peer composition; non-LDS voters contribute only idiosyncratic noise. Specifically, for each party X the post-period change in voter i's indicator for affiliation with X is

$$\Delta y_i^X \;=\; \mathbf{1}\{\text{LDS}_i\} \cdot \beta_{\text{true}} \cdot \left(Z^{\text{LDS}}_{X,\,\text{new}(i)} - Z^{\text{LDS}}_{X,\,\text{pre}}\right) + \varepsilon_i^X,$$

where ε_i^X is mean-zero idiosyncratic noise. The researcher, who does not observe LDS status, runs the regression

$$\Delta y_i^X \;=\; \beta \cdot (Z^{\text{obs}}_{X,\,\text{new}(i)} - Z^{\text{obs}}_{X,\,\text{pre}}) + e_i, \qquad i : p_i = p, \tag{19}$$

separately for each baseline party p, using the observed (all-voter) LOO peer share $Z^{\text{obs}}_{X,w}$.[40] This produces six specifications: three baseline parties, each with the two exposures $X \in \{\text{WardRep}, \text{WardDem}\}$. Our object of interest is the attenuation factor $k_{p,X} = \mathbb{E}[\hat{\beta}_{p,X}]/\beta_{\text{true}}$, with implied adjustment multiplier $1/k_{p,X}$.

Single-boundary attenuation. Fix a baseline party p and consider the subsample of voters with $p_i = p$. Taking the expectation of the DGP for voter i conditional on their

[40]Note that this specification is equivalent to: $\Delta y_i^X \;=\; \alpha + \beta \cdot Z^{\text{obs}}_{X,\,\text{new}(i)} + e_i$ by letting $\alpha = -\beta \cdot Z^{\text{obs}}_{X,\,\text{pre}}$.

post-boundary ward $\text{new}(i) = w$:

$$\begin{aligned}\mathbb{E}\big[\Delta y_i^X \mid \text{new}(i) = w\big] &= \mathbb{E}[\mathbf{1}\{\text{LDS}_i\}] \cdot \beta_{\text{true}} \cdot \big(Z_{X,w}^{\text{LDS}} - Z_{X,\text{pre}}^{\text{LDS}}\big) + \mathbb{E}\big[\varepsilon_i^X\big] \\ &= p_{\text{LDS}}(p) \cdot \beta_{\text{true}} \cdot \big(Z_{X,w}^{\text{LDS}} - Z_{X,\text{pre}}^{\text{LDS}}\big),\end{aligned}$$

where the first equality uses linearity of expectation and the fact that $Z_{X,w}^{\text{LDS}}$ and $Z_{X,\text{pre}}^{\text{LDS}}$ are properties of the wards (not of voter i), and the second equality uses $\mathbb{E}[\mathbf{1}\{\text{LDS}_i\}] = p_{\text{LDS}}(p)$ within the party-p subsample and $\mathbb{E}[\varepsilon_i^X] = 0$.

Within this subsample, the researcher's treatment $Z_{X,\text{new}(i)}^{\text{obs}}$ takes only two values: $Z_{X,A}^{\text{obs}}$ for voters with $\text{new}(i) = A$ and $Z_{X,B}^{\text{obs}}$ for voters with $\text{new}(i) = B$. For an OLS regression with a two-valued covariate, the slope coefficient is the difference of the two conditional means of the outcome divided by the difference of the two covariate values:

$$\begin{aligned}\mathbb{E}\Big[\hat{\beta}_{p,X}\Big] &= \frac{\mathbb{E}\big[\Delta y_i^X \mid \text{new}(i) = A\big] - \mathbb{E}\big[\Delta y_i^X \mid \text{new}(i) = B\big]}{Z_{X,A}^{\text{obs}} - Z_{X,B}^{\text{obs}}} \\ &= p_{\text{LDS}}(p) \cdot \beta_{\text{true}} \cdot \frac{Z_{X,A}^{\text{LDS}} - Z_{X,B}^{\text{LDS}}}{Z_{X,A}^{\text{obs}} - Z_{X,B}^{\text{obs}}},\end{aligned}$$

where the second equality plugs in the two conditional means from above and the $Z_{X,\text{pre}}^{\text{LDS}}$ terms cancel between them. Dividing by β_{true} gives the single-boundary closed-form attenuation factor:

$$k_{p,X} \;=\; \frac{\mathbb{E}\Big[\hat{\beta}_{p,X}\Big]}{\beta_{\text{true}}} \;=\; p_{\text{LDS}}(p) \cdot \frac{Z_{X,A}^{\text{LDS}} - Z_{X,B}^{\text{LDS}}}{Z_{X,A}^{\text{obs}} - Z_{X,B}^{\text{obs}}}. \tag{20}$$

Equation (20) decomposes the attenuation into a responsive-share factor $p_{\text{LDS}}(p)$ (only the LDS fraction of party-p voters responds) and a peer-mismeasurement ratio (the gap in LDS-only peer composition between the two wards, relative to the observed all-voter gap).

To evaluate Equation (20) for any assumed $p_{\text{LDS}}(\cdot)$, we need the LDS-only ward composition. For ward w and party $p \in \{R, U, D\}$, let $\pi_{w,p}^{\text{LDS}}$ denote the share of voters in ward w who have baseline party p among LDS members only, in contrast to $\pi_{w,p}$ which is the analogous share among all voters. Because the share of voters in ward w who are both party p and LDS equals $\pi_{w,p} \cdot p_{\text{LDS}}(p)$, dividing by the total share who are LDS gives

$$\pi_{w,p}^{\text{LDS}} \;=\; \frac{\pi_{w,p}\, p_{\text{LDS}}(p)}{\sum_{p' \in \{R,U,D\}} \pi_{w,p'}\, p_{\text{LDS}}(p')}. \tag{21}$$

In the large-sample limit, the LOO peer fraction $Z_{X,w}^{\text{LDS}}$ converges to $\pi_{w,X}^{\text{LDS}}$, so we can plug Equation (21) into Equation (20) to compute the analytical attenuation under any party-conditional LDS-probability assumption.

Multiple boundaries with fixed effects. Replicating the single-boundary setup at B boundaries indexed $b = 1, \ldots, B$ and pooling, the researcher runs

$$\Delta y_{i,b}^{X} \;=\; \alpha + \gamma_b + \beta \cdot Z_{X,\mathrm{new}(i,b)}^{\mathrm{obs}} + e_{i,b}, \qquad i : p_i = p, \tag{22}$$

with boundary fixed effects γ_b. By the Frisch-Waugh-Lovell theorem, $\hat{\beta}$ equals the precision-weighted average of within-boundary slopes. Substituting the per-boundary version of (20) and using that the per-boundary weight is proportional to $(Z_{X,A,b}^{\mathrm{obs}} - Z_{X,B,b}^{\mathrm{obs}})^2$ under equal sample sizes per side, one factor of the observed gap cancels, leaving

$$k_{p,X}^{\mathrm{pooled}} \;=\; \frac{\sum_b p_{\mathrm{LDS}}(p)_b \left(Z_{X,A,b}^{\mathrm{obs}} - Z_{X,B,b}^{\mathrm{obs}}\right)\left(Z_{X,A,b}^{\mathrm{LDS}} - Z_{X,B,b}^{\mathrm{LDS}}\right)}{\sum_b \left(Z_{X,A,b}^{\mathrm{obs}} - Z_{X,B,b}^{\mathrm{obs}}\right)^2}. \tag{23}$$

Equation (23) is the master formula for the multi-boundary case. Note that we allow $p_{\mathrm{LDS}}(p)_b$ to vary across party and boundary to capture the survey evidence in Appendices E.1 and E.5.

Regimes for $k_{p,X}$. Combining Equations (20) and (21) gives a useful taxonomy of the possible regimes. The responsive-share factor $p_{\mathrm{LDS}}(p)$ is bounded above by 1, so whether $k_{p,X}$ itself exceeds 1 is governed entirely by the peer-mismeasurement ratio. Three regimes arise:

- **Peer ratio $= 1$:** when $p_{\mathrm{LDS}}(\cdot)$ is constant across parties, Equation (21) collapses to $\pi_{w,p}^{\mathrm{LDS}} = \pi_{w,p}$ and the LDS-only gap equals the observed gap. The full attenuation reduces to $k_{p,X} = p_{\mathrm{LDS}}$, which is always at most 1.

- **Peer ratio < 1 (compression):** when the high-p_{LDS} party is also the majority party in both wards, Equation (21) shifts each ward's LDS-only composition toward the high-p_{LDS} party by similar absolute amounts, so the LDS-only gap is smaller than the observed gap and $k_{p,X} < p_{\mathrm{LDS}}(p) \leq 1$.

- **Peer ratio > 1 (amplification):** when the high-p_{LDS} party is in the minority in both wards, each ward's LDS subpopulation is dominated by a small fraction of high-p_{LDS} voters. Small differences in observed shares of that party then translate into large differences in LDS-only shares, and the peer ratio can exceed 1. This can happen even when $p_{\mathrm{LDS}}(\cdot)$ is identical in both wards: it is the cross-ward variation in the marginal LDS rate $\sum_{p'} \pi_{w,p'}\, p_{\mathrm{LDS}}(p')$ that drives the non-linearity in Equation (21), and that marginal rate varies whenever the ward compositions π_A and π_B differ.

For each of our six specifications, the threshold above which $k_{p,X}$ exceeds 1 is $1/p_{\text{LDS}}(p)$ on the peer ratio. With the UCEP-implied values $p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, $p_{\text{LDS}}(D) = 0.25$, the thresholds are approximately 1.18 (Republican baseline), 2.00 (Unaffiliated), and 4.00 (Democrat). Utah is dominated by the compression regime—Republicans are both the high-p_{LDS} party and the majority party in most wards—so we expect $k_{p,X} \leq 1$ and adjustment multipliers above 1 in our setting. Appendix E.4 verifies this empirically against the actual ward-boundary configurations in our analysis sample.

E.3 Simulation Results

We validate the closed forms with Monte Carlo across four scenarios that progressively add features observed in the Utah data. To make the scenarios comparable, all four are calibrated so that the population-marginal probability of being LDS equals 0.65. That figure is implied by the party-conditional probabilities of Table E1, rounded to $p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, and $p_{\text{LDS}}(D) = 0.25$, together with an assumed population party mix of $(\pi_R, \pi_U, \pi_D) = (0.50, 0.40, 0.10)$. The four scenarios differ in which sources of heterogeneity they activate:

1. **Scenario 1.** p_{LDS} uniform at 0.65, single boundary. No within- or cross-boundary heterogeneity.

2. **Scenario 2.** Party-conditional $p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, $p_{\text{LDS}}(D) = 0.25$ (the rounded values from Table E1, with population-weighted average 0.65), single boundary with $\pi_A = (0.65, 0.25, 0.10)$ and $\pi_B = (0.45, 0.30, 0.25)$. Activates the peer-mismeasurement channel.

3. **Scenario 3.** $B = 40$ boundaries with $p_{\text{LDS},b}$ stratified on $[0.30, 1.00]$ (mean 0.65), party-uniform within each boundary, ward composition constant. Activates cross-boundary heterogeneity in p_{LDS} but holds precision weights uniform.

4. **Scenario 4.** $B = 40$ boundaries with party-conditional $p_{\text{LDS}}(p)_b$ and ward composition jointly driven by an intensity index $\mu_b \in [0, 1]$. Across-boundary averages match Scenario 2; high-μ_b boundaries are uniformly Republican (small within-boundary partisan gap and therefore small precision weight). Activates all three channels of bias.

Figure E1 summarizes the results: for each scenario it plots the implied adjustment multiplier $1/k$ for each of the six specifications, comparing the analytical attenuation from Equations (20) and (23) (red crosses) with the Monte Carlo mean across 500 (single-boundary) or 200 (multi-boundary) simulation draws (navy dots). Analytical predictions lie within

~ 0.01 of the Monte Carlo means in every cell, confirming that the closed forms match the simulation.

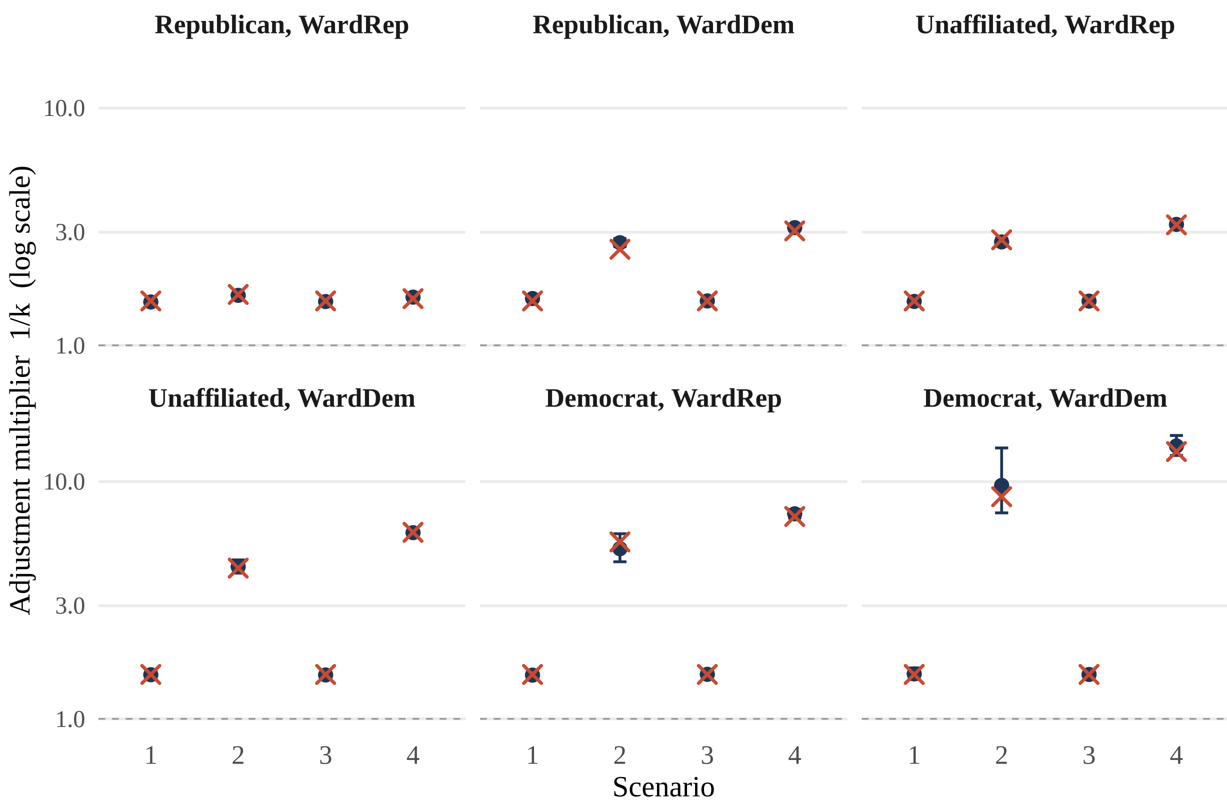

Figure E1: Implied adjustment multiplier $1/k$ for each of the six specifications across the four scenarios. Navy dots are the Monte Carlo mean with ± 1.96 standard-error bars on $1/k$; red crosses are the analytical prediction from Equations (20) or (23). The y-axis is on a log scale; the dashed horizontal line marks $1/k = 1$ (no adjustment).

The exercise yields three takeaways. First, adjusting for the fact that not every Utah voter is LDS scales the published coefficients up by a factor that depends on the assumed joint distribution of LDS status, party, and geography. Under uniform $p_{\text{LDS}} = 0.65$ (Scenarios 1 and 3), all six specifications adjust by the same modest factor $1/0.65 \approx 1.54$. Second, allowing p_{LDS} to vary by party (Scenario 2) introduces substantial cross-specification variation: the adjustment for baseline-Republican voters' response to a more Republican ward (WardRep) stays close to $1.64\times$, while the adjustment for baseline-Democrats' response to a more Democratic ward swells to $\sim 8.6\times$. Third, allowing p_{LDS} to vary across boundaries and to correlate with ward composition (Scenario 4) widens the spread further: the same Democrat-on-WardDem specification now requires a $\sim 13\times$ adjustment, while the Republican-on-WardRep specification barely moves. In all four scenarios, the implied adjustment multiplier exceeds 1 for every specification, consistent with the analytical compression regime above: under empirically grounded assumptions about p_{LDS}, our main estimates understate the underlying causal effect on LDS members. Whether this conclusion survives once the stylized ward compositions are replaced by Utah's actual boundary distribution is the subject of the next subsection.

E.4 Empirical Verification with Utah Boundary Pairs

The four scenarios above use stylized ward compositions. We now evaluate the closed-form attenuation on every newly-drawn ward-boundary pair from our 2013–2017 sample—the same boundary set that defines treatment in our main analysis. For each of the 1,588 boundaries with both flanking wards observed in the 2012 voter file, we read the observed party shares on either side, plug them into Equation (21) using the UCEP-implied values $p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, $p_{\text{LDS}}(D) = 0.25$, and form the implied $k_{p,X}$ for each of the six specifications. Table E2 reports both the precision-weighted pooled k from Equation (23), applied to the empirical boundary distribution, and per-boundary summary statistics.

Table E2: Empirical Utah attenuation factors. For each of the 6 specs we compute the implied k across the 1588 paired 2013–2017 ward-boundary pairs, using the standard party-conditional values $p_{\text{LDS}}(R) = 0.85$, $p_{\text{LDS}}(U) = 0.50$, $p_{\text{LDS}}(D) = 0.25$ from the UCEP table. "Pooled k" is the precision-weighted estimator from Equation (23) applied to actual boundary data; "Adj. mult." is its reciprocal. "Median k" and "Max k" summarize the distribution across boundary pairs; "$\#\{k > 1\}$" counts boundaries whose individual k exceeds 1 (corresponding adjustment multiplier < 1).

Baseline party	Exposure	$p_{\text{LDS}}(p)$	Pooled k	Adj. mult.	Median k	Max k	$\#\{k > 1\}$
Republican	WardRep	0.85	0.784	1.28	0.753	79.210	192
Republican	WardDem	0.85	0.392	2.55	0.332	6.708	20
Unaffiliated	WardRep	0.50	0.461	2.17	0.443	46.594	50
Unaffiliated	WardDem	0.50	0.230	4.34	0.195	3.946	10
Democrat	WardRep	0.25	0.230	4.34	0.221	23.297	18
Democrat	WardDem	0.25	0.115	8.68	0.098	1.973	5

The pooled k is below 1 for every specification, with implied adjustment multipliers ranging from 1.28× (Republican baseline, WardRep exposure) to 8.68× (Democrat baseline, WardDem exposure). These are the corrections that would actually apply to our published estimates under the UCEP-implied p_{LDS} values—qualitatively in line with Scenario 4, and concretely larger than the uniform 1.54× implied by Scenario 1.

Per-boundary k values, by contrast, are not all bounded below 1. Among the 1,588 boundary pairs, 192 (12%) have an implied $k_{R,\text{WardRep}} > 1$, with smaller fractions for the other five specifications (Table E2). These are boundaries that sit in the amplification regime characterized in Appendix E.2: both flanking wards have low Republican shares, and the cross-ward Republican gap is small, so that even a modest LDS-only gap divided by a tiny observed denominator produces a large peer ratio. Crucially, the same small observed gap that drives the large per-boundary ratio also yields a small precision weight $(Z^{\text{obs}}_{X,A,b} - Z^{\text{obs}}_{X,B,b})^2$ in Equation (23), so those boundaries contribute negligibly to the pooled estimator. Figure E2 plots the full distribution of peer ratios across the 1,588 paired boundaries; for the WardDem exposure (lower panel) the entire distribution sits well below even the lowest threshold ($1/p_{\text{LDS}}(R) = 1.18$), and even for the WardRep exposure (upper panel) the bulk of the mass lies below the Unaffiliated threshold (2.00) and far below the Democrat threshold (4.00).

Taken together, the empirical exercise confirms the conservative interpretation suggested by Scenarios 1–4: under the UCEP-implied p_{LDS} values applied to Utah's actual ward-boundary distribution, every published estimate should be scaled up to recover the underlying causal effect on LDS members, with adjustment multipliers in the 1.3×–8.7× range. The largest correction applies to baseline-Democrat estimates and the smallest to baseline-Republican estimates, mirroring the cross-specification spread already visible in the simulation results.

E.5 Estimating Local LDS Density

The most granular spatial variation available in public data is the county. The 2010 U.S. Religion Census tabulates the members on each congregation's roll and reports adherents by county; nationally, adherents are concentrated in the Intermountain West, with the highest shares running through Utah and the surrounding "Mormon corridor," and most Utah counties are over 70% LDS.

To better understand how LDS members are distributed within counties, we construct our own ward-level estimate from two ingredients: an estimate of the population living within each ward's boundary (from intersecting ward polygons with U.S. Census blocks), and an estimate of average weekly attendance per ward (from crowd-sourced data on `returnandreport.org`). Figure E3 plots the distribution of mean weekly attendance across the 112 Utah wards in the `returnandreport.org` sample; the cross-ward mean is 155 attendees per week. We then estimate each ward's attendance share by dividing an assumed weekly attendance of 150 (the cross-ward mean of 155, rounded) by the ward's residential population.[41]

This ward-level share reveals substantial within-county variation. Salt Lake County,

[41] Ward-level population estimates are constructed by assigning 2010 and 2020 census block populations to wards in proportion to the area of overlap between each block and each ward, then linearly interpolating between the two endpoints to get 2013–2017 estimates.

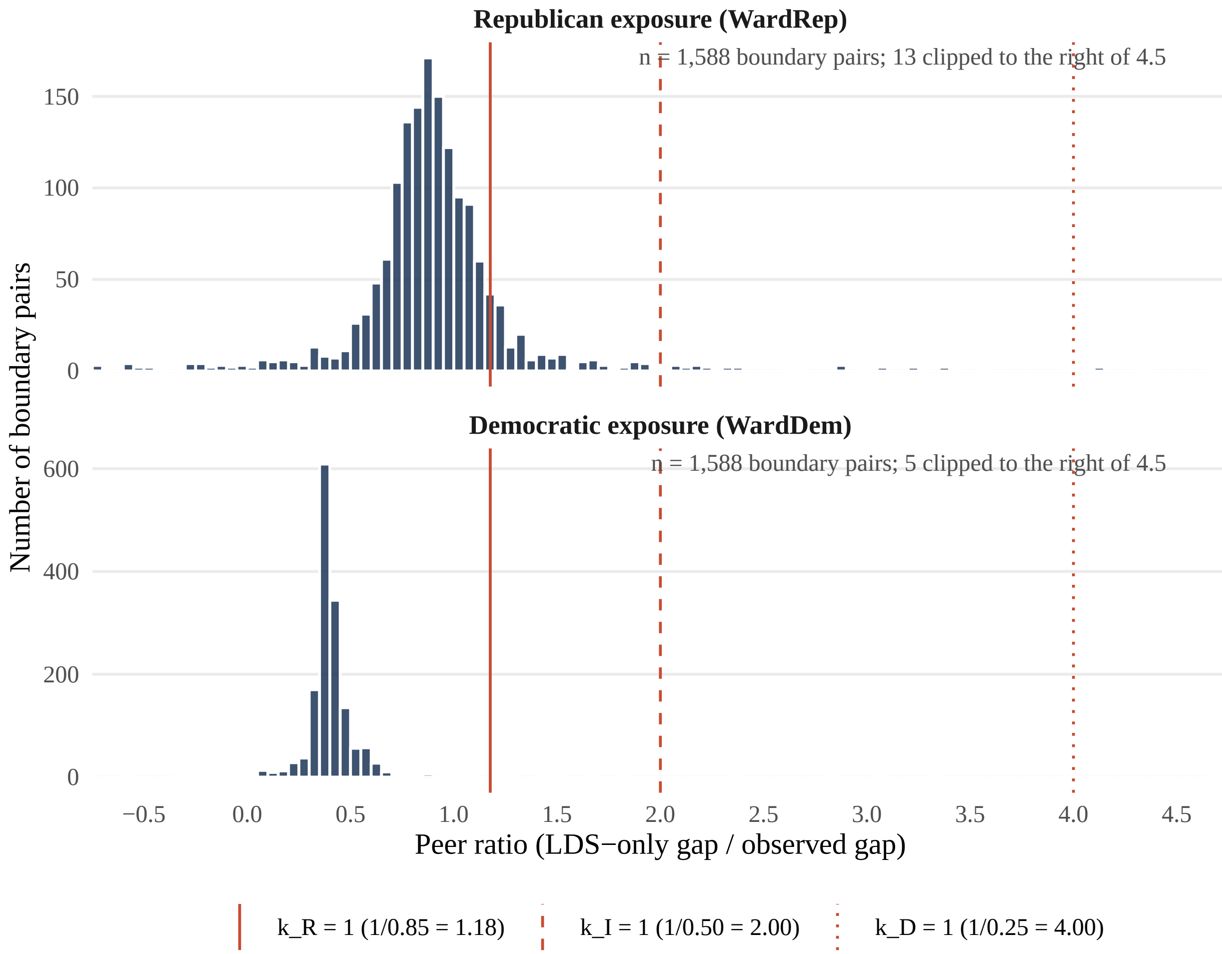

Figure E2: Distribution of the implied peer-mismeasurement ratio across 1,588 Utah ward-boundary pairs in our 2013–2017 sample, by exposure. Vertical lines mark the peer-ratio thresholds at which the implied $k_{p,X}$ would cross 1 for each baseline party (solid: k_R at 1/0.85; dashed: k_U at 1/0.50; dotted: k_D at 1/0.25). Boundaries above the relevant threshold are those for which the conservative "adjustment multiplier > 1" interpretation does not hold at the per-boundary level. The right tail beyond 4.5 (a handful of boundaries with very small observed gaps) is clipped from the figure for legibility; those boundaries contribute negligibly to the pooled estimator (Table E2) because their small observed gaps yield small precision weights in Equation (23).

whose county-wide LDS share is relatively low in the U.S. Religion Census, contains pockets of very high LDS density in the suburbs south of Salt Lake City. Utah County, by contrast, has a very high overall LDS share but contains pockets of low LDS density around Utah Lake. Most of the area north of Salt Lake County in Weber and Davis Counties is heavily LDS, except for the city of Ogden, which historically drew non-LDS migration as a railroad hub. Rural counties tend to have high LDS shares overall, with historically non-Mormon mining towns (e.g., Price and Park City) and gateway communities to the national parks (e.g., Moab) as notable exceptions. Appendix E.6 uses this ward-level measure to test whether peer effects are stronger where LDS density is higher.

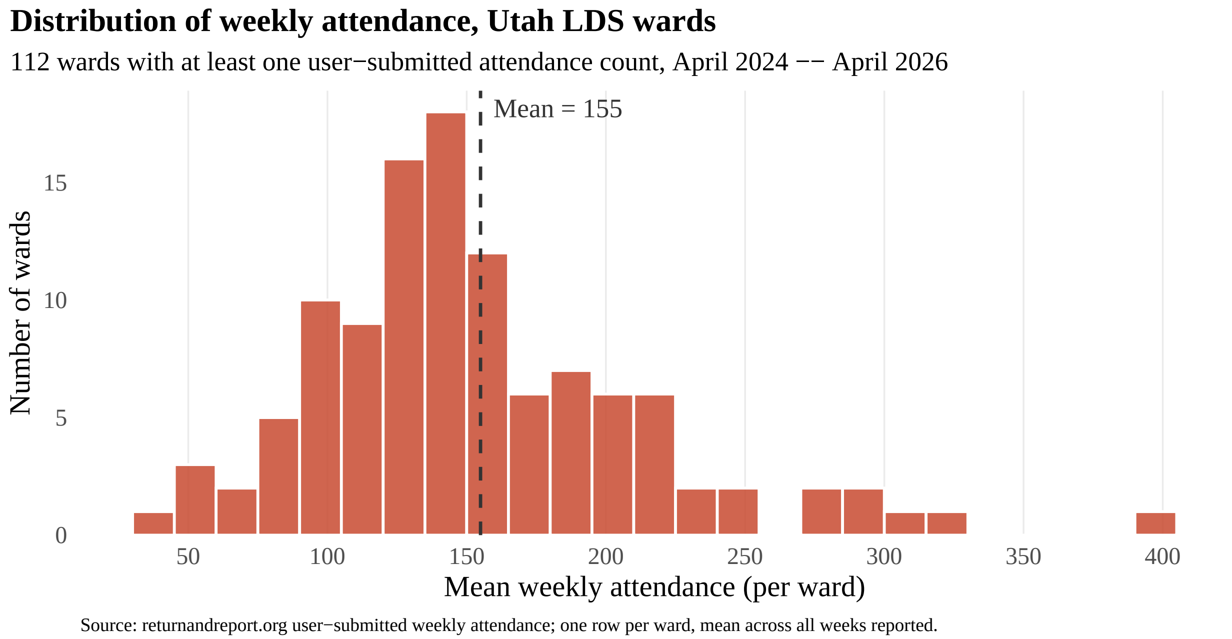

Figure E3: Distribution of mean weekly LDS church attendance across Utah wards. One observation per ward; each ward's value is the mean of all user-submitted weekly attendance counts on `returnandreport.org` between April 2024 and April 2026. The dashed line marks the cross-ward mean.

E.6 Heterogeneity by Local LDS Density

Appendices E.2–E.4 show that our inability to observe religious affiliation attenuates our estimates: every voter near a boundary enters the sample, but only LDS members plausibly respond to their assigned ward's composition, so the coefficients blend a member response with a non-member null. That logic carries a testable implication across regions. Where a larger share of a ward's residents are Latter-day Saints, a larger share of the voters we treat as exposed actually are exposed, less of the estimate is diluted, and the measured effect should—all else equal—be stronger.

We measure local LDS density with the ward-level estimate constructed in Appendix E.5, the assumed weekly attendance count divided by the ward's residential population, evaluated on each voter's pre-change ward (median 25%). We interact the treatment with indicators for the quartile of that measure across the 1,430 boundaries in the sample. Table E3 and Figure E4 report the estimates.

Contrary to the prediction, effects are if anything weaker where LDS density is higher. Among 2012 Republicans facing Democratic exposure, the coefficient on Republican retention is -1.002 in the least dense quartile and 0.181 in the most dense ($p = 0.000$ for the difference); switching to Democrat runs 0.381 against -0.006 ($p = 0.003$). Among 2012 Unaffiliated voters facing Republican exposure the profile has the same shape, 0.411 against 0.128, though the difference is not significant ($p = 0.222$). A median split gives the same ordering on all three margins.

We note two possible explanations for this finding. First, many variables correlate with LDS share across space, so it's not the case that we vary LDS density while holding all else is equal. For example, our density measure is positively correlated with the ward's Republican registration share (0.44) and more heavily Republican wards offer less variation in out-party exposure to begin with. Second, true ward peer effects may simply be larger where LDS density is lower. The model of Appendix D treats the ward as the pool from which a member's friends are drawn, but in practice a social circle is assembled from the ward and from the surrounding neighborhood. If Latter-day Saints disproportionately befriend other Latter-day Saints, the share of that circle the ward supplies depends on what the neighborhood offers. In a heavily LDS area a member can find co-religionists next door regardless of which ward the line assigns them to, so reassignment shifts little of their friend circle; in a sparsely LDS area the ward is close to the only source, so reassignment shifts much of it. Under that account a redrawn boundary is a larger intervention on the friend circle precisely where local density is lowest. We cannot distinguish the two readings with these data.

Table E3: Peer Effects by Local LDS Density

	2012 Rep. → Rep. (Dem. exposure)	2012 Rep. → Dem. (Dem. exposure)	2012 Unaff. → Rep. (Rep. exposure)
Quartiles of pre-change ward LDS density			
Q1 (least LDS dense)	-1.002^{***}	0.381^{***}	0.411^{***}
	(0.267)	(0.118)	(0.107)
Q2	-0.312	0.131	0.033
	(0.241)	(0.116)	(0.174)
Q3	-0.271	0.174^{**}	0.008
	(0.232)	(0.073)	(0.220)
Q4 (most LDS dense)	0.181	-0.006	0.128
	(0.191)	(0.053)	(0.206)
Q4 − Q1	1.183^{***}	-0.387^{***}	-0.283
	(0.329)	(0.130)	(0.232)
Median split			
Below median	-0.708^{***}	0.272^{***}	0.298^{***}
	(0.180)	(0.083)	(0.094)
Above median	0.028	0.060	0.057
	(0.157)	(0.045)	(0.148)
Above − below	0.736^{***}	-0.212^{**}	-0.242
	(0.240)	(0.095)	(0.175)
Linear interaction with ward LDS density			
Effect at mean density	-0.363^{***}	0.173^{***}	0.178^{**}
	(0.119)	(0.048)	(0.090)
Slope per unit density	2.893^{***}	-0.885^{**}	-0.925
	(0.958)	(0.388)	(0.679)
N voters	72,254	72,254	59,069
N boundaries	1,410	1,410	1,430

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each column is one (baseline party, exposure, outcome) cell of the preferred specification (m13, 400-meter bandwidth), with the treatment interacted with bin indicators for the pre-change ward's estimated LDS density (Appendix E.5). Q1 is the least dense quartile. All bins enter a single regression, so the difference rows are Wald tests on that fit's clustered variance matrix rather than comparisons of separate confidence intervals. The bottom panel replaces the bins with a linear interaction between the treatment and the density measure centered at its mean over pre-change wards, reporting the effect at the mean and the slope per unit (0 to 1) of density. Standard errors, in parentheses, are clustered two-way by ward-boundary side and voter.

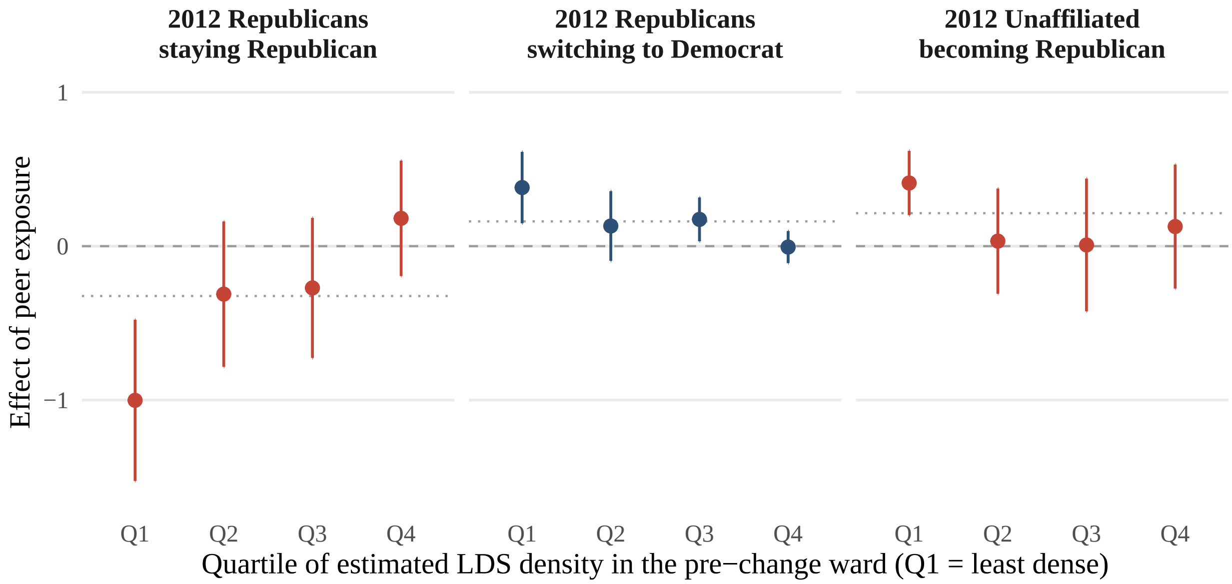

Figure E4: Peer effects by quartile of pre-change ward LDS density. Points are per-quartile treatment effects with 95% confidence intervals; Q1 is the least LDS-dense quartile. The dashed line marks zero and the dotted line the pooled preferred-specification estimate for that cell.

F Robustness and Inference

This appendix collects the design checks, robustness exercises, and alternative modes of inference referenced in Sections 4–6. Appendix F.1 reports the stand-alone boundary-side coefficients δ. Appendix F.2 reports convergence diagnostics for the simulated-boundary exercise of Section 5.2. Appendix F.3 supplies further evidence on the linear-in-means approximation of Section 6.2. Appendices F.4 through F.6 add progressively richer control sets, vary the distance bandwidth from 50 to 600 meters, and report alternative specifications and standard errors. Appendix F.7 reports design-based randomization inference for the main estimates, and Appendix F.8 re-estimates the main specification separately for movers and non-movers, bearing on the attrition discussion of Section 6.4.

F.1 Boundary-Side Level Differences

Table F1 reports the estimates of δ, the coefficient on the stand-alone high-exposure-side indicator D_{ib}, from the preferred specification. As discussed in Section 4.2, the conceptual framework implies that as I_b approaches zero there should be no difference between the two sides of a new boundary. The estimates are uniformly small—within three percentage points of zero—and mostly insignificant. The few marginally significant cells are concentrated in the Unaffiliated subsample and carry the opposite sign of the corresponding treatment effect, the pattern implied by the negative sampling correlation between an extrapolated intercept and its slope, rather than the aligned signs a level confound would produce.

Table F1: Estimates of δ (Stand-Alone Boundary-Side Term), Preferred Specification

	Republican exposure	Democratic exposure
2012 Republicans		
Republican registration	0.005	0.011
	(0.008)	(0.008)
Democratic registration	-0.002	-0.003
	(0.003)	(0.003)
Unaffiliated registration	-0.003	-0.007
	(0.007)	(0.007)
2012 Unaffiliated voters		
Republican registration	-0.025^{**}	0.027^{**}
	(0.011)	(0.011)
Democratic registration	0.004	0.001
	(0.008)	(0.008)
Unaffiliated registration	0.021^{*}	-0.028^{**}
	(0.012)	(0.012)
2012 Democrats		
Republican registration	0.027	-0.026
	(0.022)	(0.023)
Democratic registration	-0.014	0.011
	(0.027)	(0.027)
Unaffiliated registration	-0.014	0.016
	(0.021)	(0.022)

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each cell reports $\hat{\delta}$ from Equation (1) under the preferred specification (400-meter bandwidth, boundary fixed effects, quadratic distance controls, no individual controls), estimated separately by baseline-party subsample and exposure; outcomes are first-differenced registration indicators. Standard errors, in parentheses, are clustered two-way by ward-boundary side and voter.

F.2 Simulated-Boundary Convergence Diagnostics

Table F2 reports quantiles of the split-$\hat{R}$ convergence diagnostic across the stake-year simulation runs underlying the simulated-boundary exercise of Section 5.2.

Table F2: Convergence diagnostics for simulated ward plans

Variable	Stake-years	p05	p10	p25	Median	p75	p90	p95
Black-white dissimilarity	1,851	1.000	1.001	1.002	1.005	1.017	1.288	1.936
Hispanic/non-Hispanic dissimilarity	1,854	1.000	1.001	1.002	1.005	1.014	1.198	2.026
Maximum Black population share	1,838	1.000	1.001	1.002	1.004	1.013	1.151	1.854
Maximum Hispanic population share	1,839	1.000	1.001	1.002	1.005	1.012	1.159	1.835
Maximum White population share	1,846	1.000	1.001	1.002	1.005	1.011	1.075	1.635
Mean Black population share	1,851	1.000	1.001	1.002	1.005	1.013	1.167	1.900
Mean Hispanic population share	1,869	1.000	1.001	1.002	1.005	1.015	1.163	1.873
Mean Polsby-Popper compactness	1,876	1.000	1.001	1.002	1.006	1.015	1.189	1.677
Mean White population share	1,869	1.000	1.001	1.002	1.006	1.015	1.159	1.984
Minimum Black population share	1,120	1.000	1.000	1.001	1.003	1.005	1.017	1.041
Minimum Hispanic population share	1,850	1.000	1.001	1.002	1.005	1.011	1.063	1.911
Minimum Polsby-Popper compactness	1,867	1.000	1.000	1.001	1.003	1.007	1.082	1.595
Minimum White population share	1,842	1.000	1.001	1.002	1.004	1.013	1.120	1.624
Range in Black population	1,781	1.000	1.000	1.002	1.005	1.013	1.093	1.637
Range in Hispanic population	1,861	1.000	1.001	1.002	1.006	1.017	1.304	2.004
Range in White population	1,868	1.000	1.001	1.002	1.004	1.012	1.222	2.040
Range in total population	1,868	1.000	1.000	1.001	1.003	1.012	1.200	1.914
Variance in Black population	1,811	1.000	1.001	1.002	1.006	1.015	1.220	1.861
Variance in Black population share	1,851	1.000	1.001	1.002	1.005	1.014	1.179	1.798
Variance in Hispanic population	1,863	1.000	1.001	1.002	1.006	1.015	1.283	2.152
Variance in Hispanic population share	1,869	1.000	1.001	1.002	1.005	1.013	1.166	2.136
Variance in White population	1,869	1.000	1.000	1.001	1.004	1.012	1.185	2.062
Variance in White population share	1,869	1.000	1.001	1.002	1.005	1.012	1.139	1.836
Variance in total population	1,869	1.000	1.000	1.001	1.003	1.012	1.173	2.077

Notes: Entries summarize finite split R-hat statistics from the final per-stake, per-year simulation diagnostics. The table includes stake-years with final status passed, flagged, or not assessable when a finite R-hat statistic is available.

F.3 Additional Evidence on Linearity

Two figures complement the binned dose-response test of Section 6.2. Figure F1 reports the quartile-bin contrasts for the two exposure-sample combinations not shown in the main text (Republican exposure among 2012 Republicans; Democratic exposure among 2012 Unaffiliated voters); as in the main-text panels, the binned estimates track the linear fit. Figure F2 interacts the treatment with the baseline (pre-split) ward's partisan composition: per-quartile treatment effects are statistically indistinguishable across quartiles of the origin ward's Republican (or Democratic) share, with no rising response at high out-party shares, no sign reversals, and no threshold pattern, which is what tipping-point or backlash models would produce.

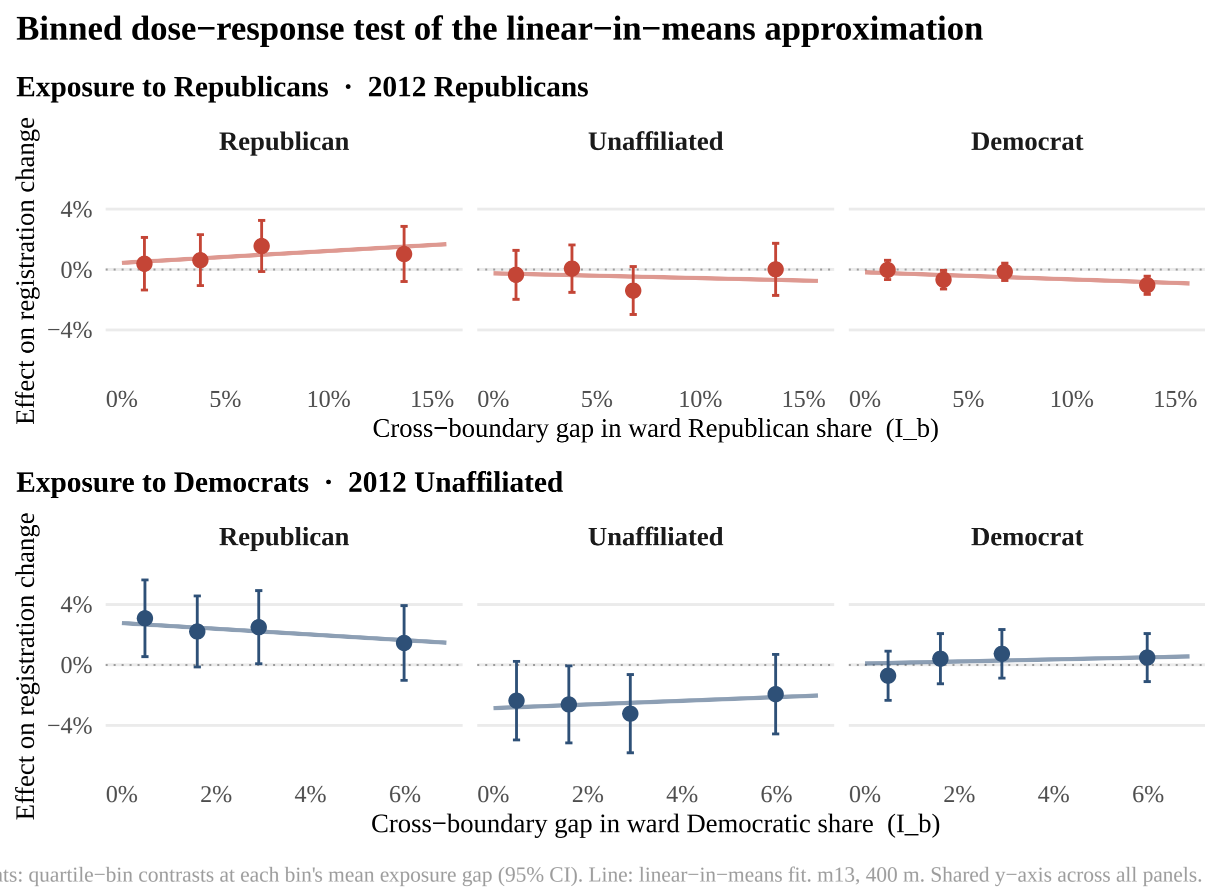

Figure F1: Binned Dose-Response Test: Remaining Exposure-Sample Combinations. See Figure 5 notes.

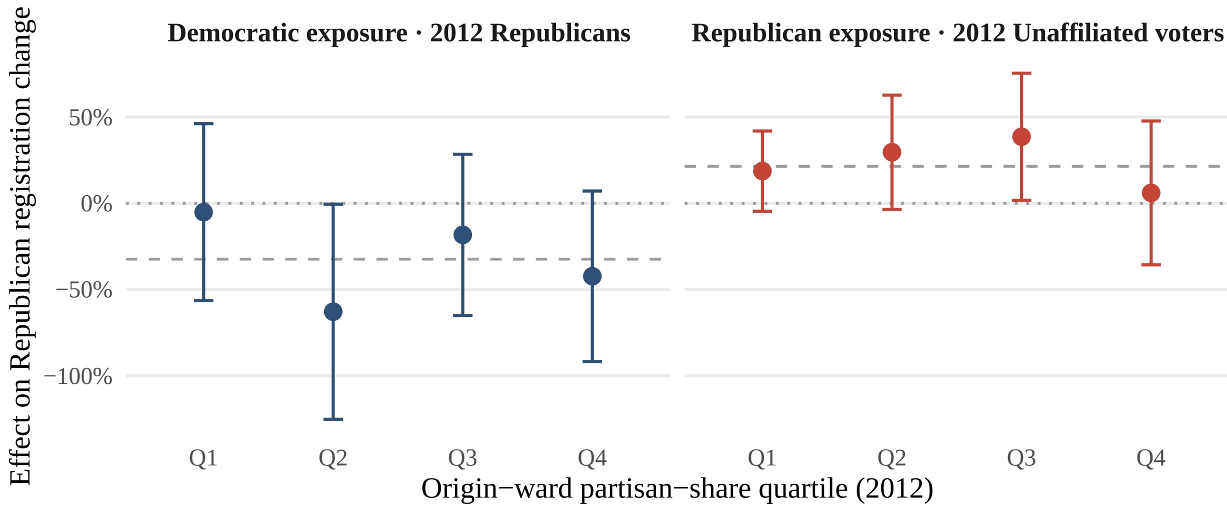

Figure F2: Treatment Effects by Baseline Ward Composition. Per-quartile treatment effects from interacting the treatment with quartiles of the origin ward's 2012 partisan share, for the two main exposure–sample pairs: Democratic exposure among 2012 Republicans and Republican exposure among 2012 Unaffiliated voters. Dashed line: pooled preferred-specification estimate at the 400m bandwidth. Whiskers: 95% confidence intervals.

F.4 Robustness to Controls

Our preferred specification includes boundary fixed effects and quadratic distance controls but no individual-level controls, on the reasoning that the level characteristics stake presidents balance on are absorbed by first-differencing the outcome (Section 4). The concern that remains is that those characteristics proxy for differential trends. Figure F3 therefore re-estimates each cell of Table 3 under progressively richer control sets: the preferred specification, then adding demographic trends (age, sex, race), then property trends (assessed value, square footage, year built), then school-catchment fixed effects. Point estimates are stable across the four control sets, and the significant cells retain their sign and significance throughout.

F.5 Robustness to Bandwidth

The 400-meter bandwidth of our preferred specification trades precision against the comparability of voters on either side of a new line (Section 4.1). Figure F4 re-estimates the six main estimates at bandwidths from 50 to 600 meters, holding the rest of the specification fixed. Estimates at narrow bandwidths are noisy, as expected given the small samples they retain, but the point estimates are stable in sign and magnitude across the range, and none of our conclusions depends on the choice of 400 meters.

F.6 Alternative Specifications and Inference

Figure F5 re-estimates the six main estimates of Figure F4 under five variants of the preferred specification: Conley (5 km) spatial standard errors in place of two-way clustering; a 5-meter "donut" that drops voters immediately adjacent to the boundary line; linear rather than quadratic distance controls; dropping, for each treatment, the boundary with the largest cross-boundary treatment gap; and restricting to boundaries whose adjacent wards changed boundaries exactly once during 2013–2017. Point estimates are stable across variants, and the Conley confidence intervals are only modestly wider than the clustered ones. Table F3 reports, for the six main estimates, the p-value under two-way clustering, clustering by boundary pair, Conley standard errors at 1, 2, 5, and 10 kilometers, and randomization inference. The changed-once restriction discards the most data and is accordingly the noisiest, but its estimates remain consistent with the baseline.

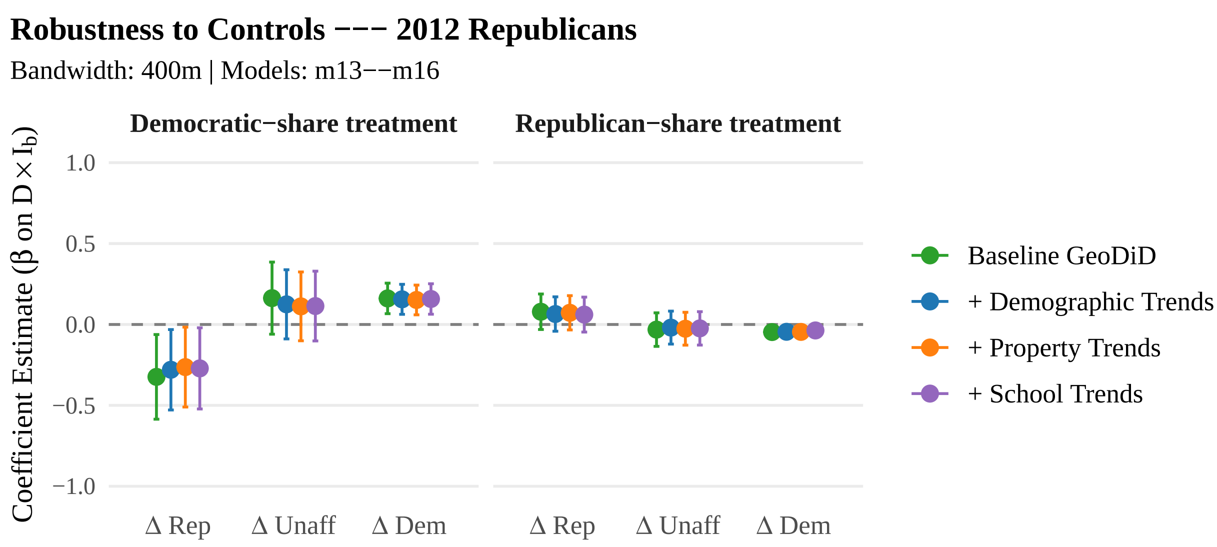

(a) 2012 Republicans

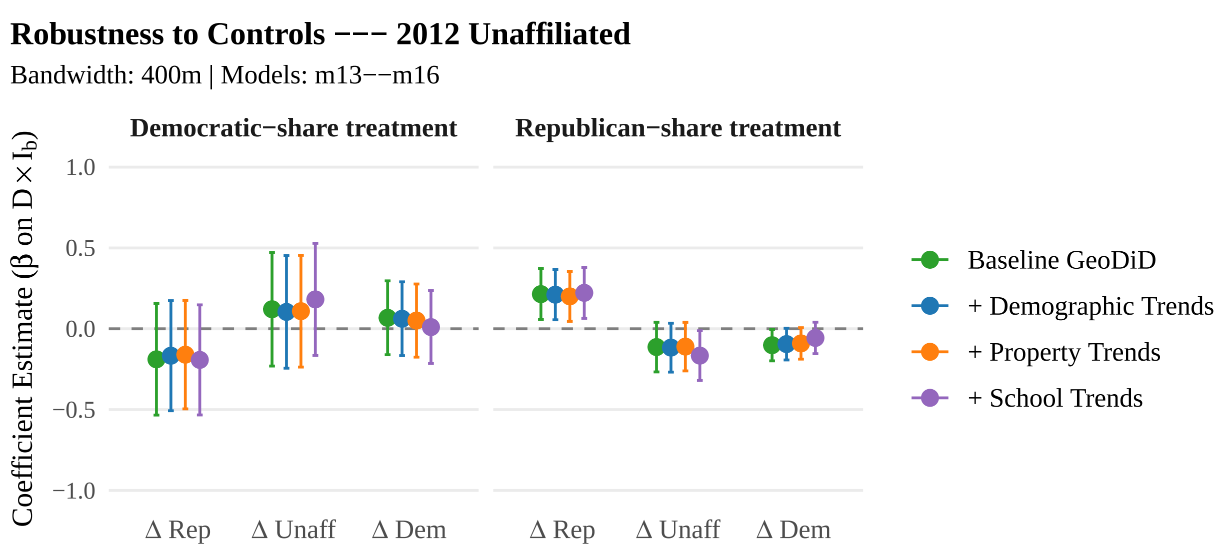

(b) 2012 Unaffiliated voters

Figure F3: Robustness of main effects to progressively richer control sets, by 2012 baseline party. Each panel within a subfigure plots the coefficient $\hat{\beta}$ on $D_{ib} \times I_b$ for one (treatment, outcome) cell of the corresponding bybase results table (Table 3). Points are coefficient estimates and whiskers are 95% confidence intervals, colored by control-set specification: the preferred baseline specification, then progressively adding demographic trends, property trends, and school trends.

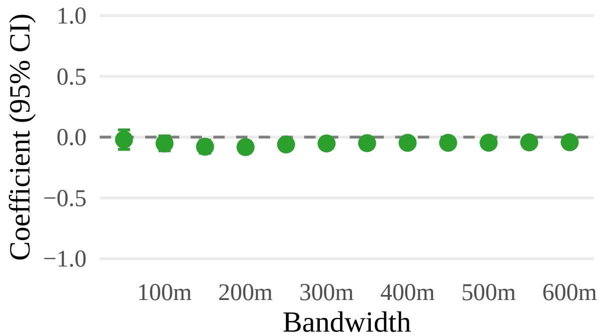

(a) 2012 Rep. → 2020 Dem. (Republican exposure)

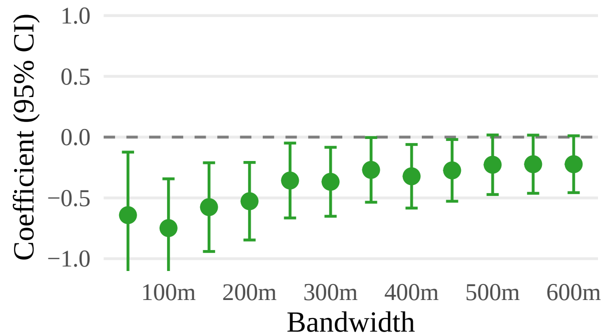

(b) 2012 Rep. → 2020 Rep. (Democratic exposure)

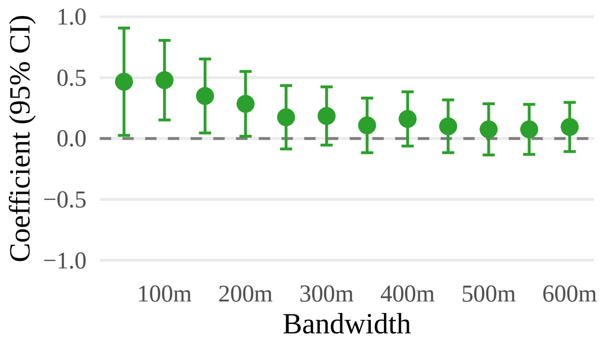

(c) 2012 Rep. → 2020 Ind. (Democratic exposure)

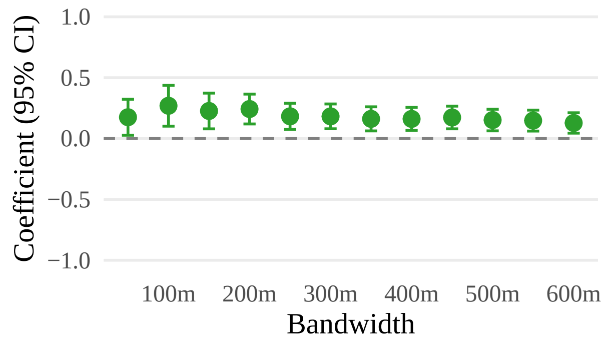

(d) 2012 Rep. → 2020 Dem. (Democratic exposure)

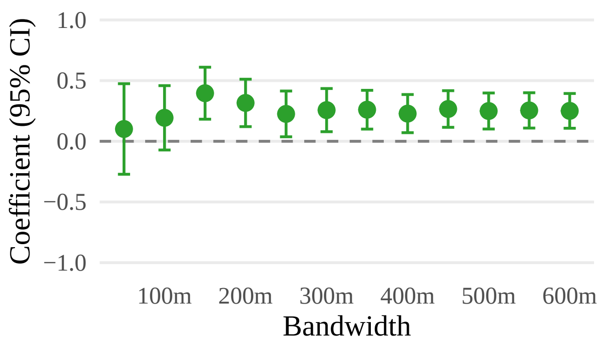

(e) 2012 Unaff. → 2020 Rep. (Republican exposure)

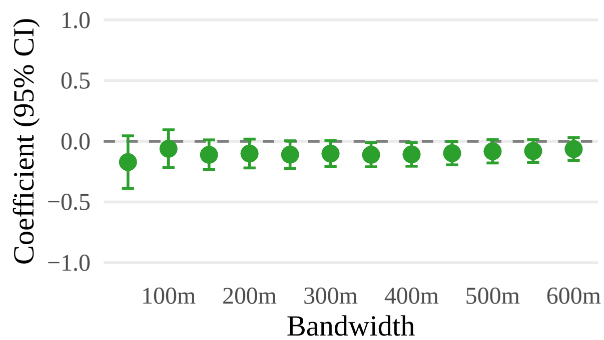

(f) 2012 Unaff. → 2020 Dem. (Republican exposure)

Figure F4: Bandwidth robustness for the six main estimates (the five significant effects in Table 4 plus Republicans' switching to Unaffiliated). Each panel plots $\hat{\beta}$ on $D_{ib} \times I_b$ with a 95% confidence interval at distance-to-boundary bandwidths from 50m to 600m, holding the preferred specification fixed.

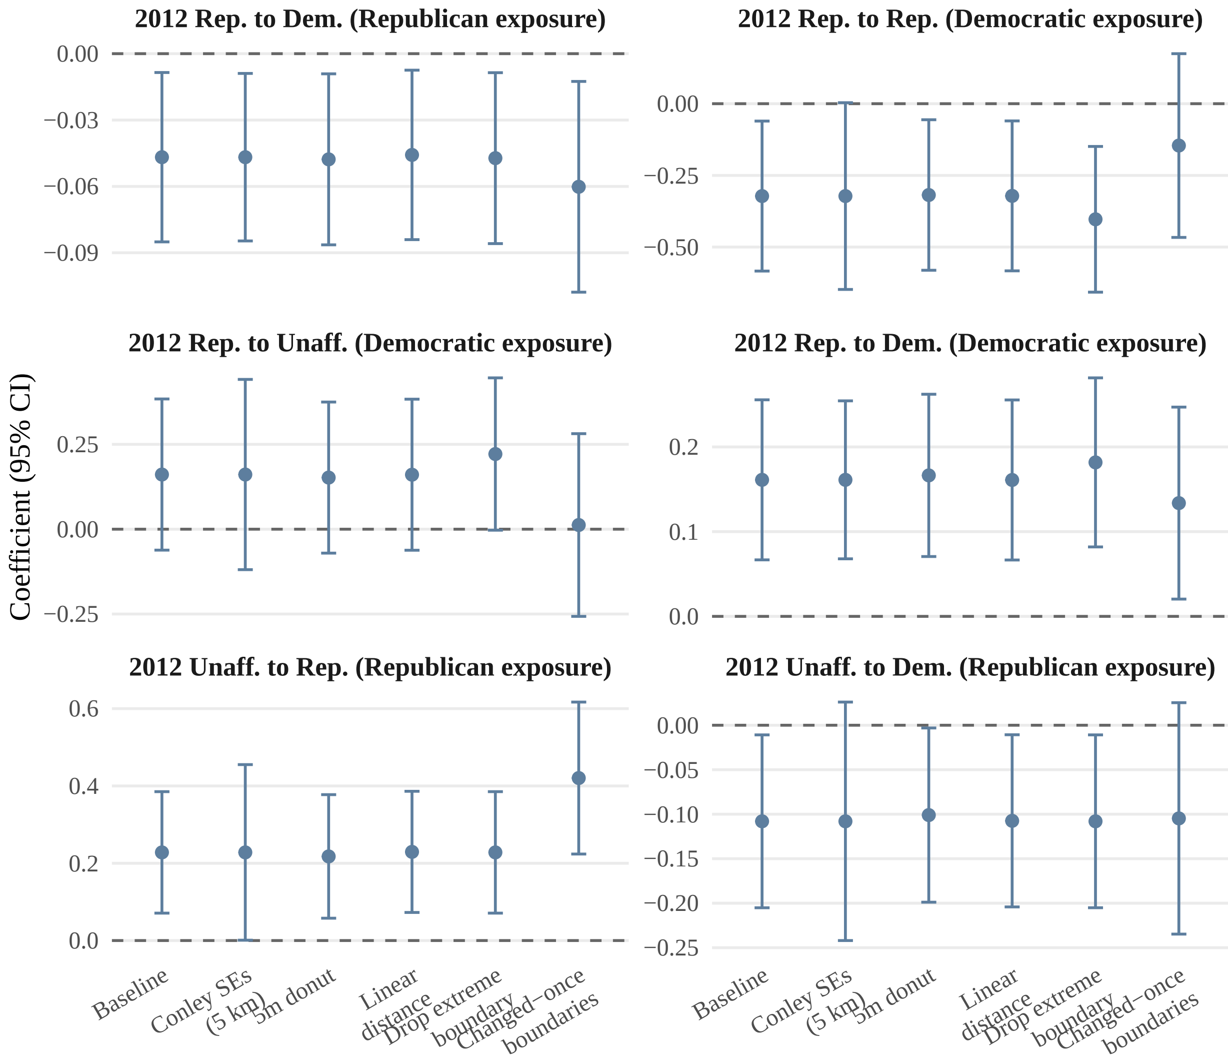

Figure F5: Robustness of the six main estimates to alternative specifications and inference. Each panel plots the coefficient $\hat{\beta}$ on $D_{ib} \times I_b$ with a 95% confidence interval under the preferred specification (Baseline) and five variants: Conley (5 km) spatial standard errors in place of two-way clustering; dropping voters within 5 meters of the boundary line (donut); linear instead of quadratic distance controls; dropping the boundary with the largest cross-boundary treatment gap; and restricting to boundaries whose adjacent wards changed exactly once during 2013–2017.

F.7 Randomization Inference

Randomization inference uses the design itself, rather than a large-sample approximation, to conduct inference. The identifying claim of the design is that, conditional on being near a given boundary and on distance to it, which side received the higher peer exposure is as-good-as-random. Under the sharp null hypothesis that peer composition does not affect

Table F3: Inference Robustness for the Six Main Estimates

		p-value under alternative inference						
	$\hat{\beta}$ (SE)	Two-way cluster	Cluster by boundary	Conley 1 km	Conley 2 km	Conley 5 km	Conley 10 km	Randomization
2012 Unaff. → 2020 Rep. (Rep. exposure)	0.228*** (0.080)	0.004	0.023	0.034	0.044	0.049	0.087	0.030
2012 Unaff. → 2020 Dem. (Rep. exposure)	−0.108** (0.050)	0.029	0.078	0.131	0.121	0.114	0.168	0.097
2012 Rep. → 2020 Rep. (Dem. exposure)	−0.322** (0.133)	0.016	0.045	0.055	0.074	0.053	0.028	0.059
2012 Rep. → 2020 Dem. (Dem. exposure)	0.161*** (0.048)	0.001	0.006	0.003	0.002	0.001	0.001	0.008
2012 Rep. → 2020 Unaff. (Dem. exposure)	0.161 (0.114)	0.157	0.239	0.274	0.310	0.260	0.183	0.278
2012 Rep. → 2020 Dem. (Rep. exposure)	−0.047** (0.020)	0.017	0.056	0.041	0.044	0.015	0.009	0.064

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$ on the coefficient, from the preferred specification (m13, 400-meter bandwidth) with standard errors two-way clustered by ward-boundary side and voter. Each remaining column reports the p-value for the same coefficient under one alternative: one-way clustering by boundary pair, which treats the two sides of a line as one cluster; Conley spatial standard errors allowing correlation among all voters within the stated distance regardless of the line; and randomization inference, which relabels which side of each boundary is the high-exposure side in each of 1,000 draws and re-estimates the specification (Appendix F.7). All rows use the same sample as Table 3.

outcomes, we re-randomize side assignment—flipping the high-exposure indicator, the sign of the distance running variable, and the dose interaction together, while leaving the boundary's geography, its cross-boundary exposure contrast, and all outcomes untouched—and re-estimate the main specification. Because the test conditions on the realized magnitude of each boundary's exposure contrast, it is a test of side assignment, not of boundary placement (the placebo and simulation exercises of Section 5 address placement).

Figure F6 shows the permutation distribution of the t-statistic for each of the six main estimates, based on 1,000 draws per estimate, with the observed t-statistic marked. The permutation distributions are wider than the standard normal (standard deviations near 1.31), making these more conservative tests. Republicans' switching to Democrat under Democratic exposure remains clearly significant ($p = 0.008$), as does Unaffiliated voters' Republican registration under Republican exposure ($p = 0.030$); Republicans' retention under Democratic exposure sits just above the 5% threshold ($p = 0.059$); and the remaining three estimates have p-values from 0.064 to 0.278. Appendix Table F3 displays these alongside the clustered and Conley p-values.

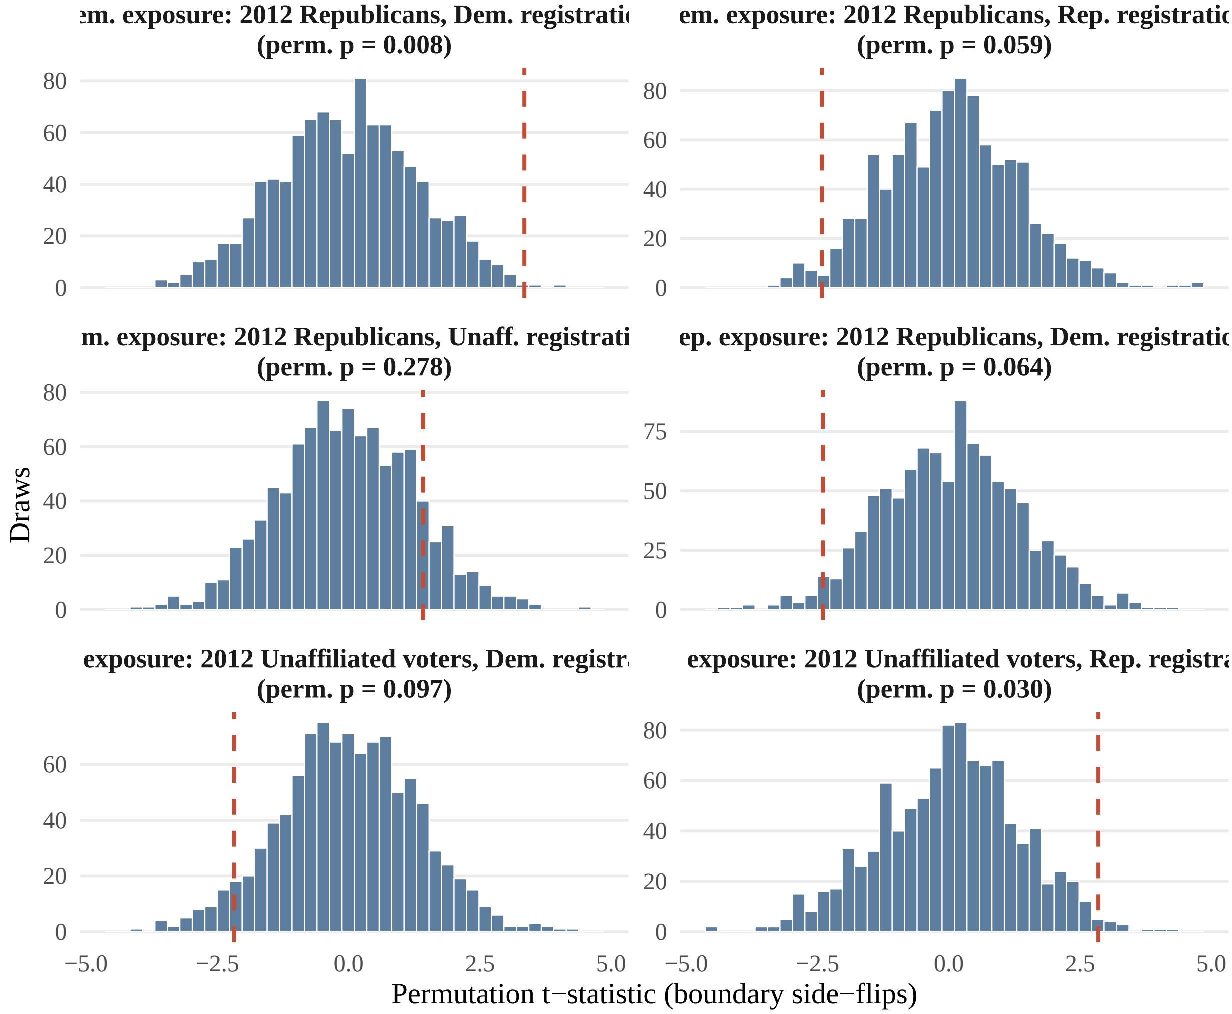

Figure F6: Randomization inference: permutation distributions of the t-statistic. Each panel shows the distribution of the main-specification t-statistic across 1,000 boundary side-flip permutations for one of the six main estimates; the dashed line marks the observed t-statistic. Permutation p-values (share of draws with $|t|$ at least as large as observed, with the observed draw included) are reported in each panel label.

F.8 Partisan Peer Exposure by Residential Mobility

Tables F4 and F5 re-estimate the main specification separately for movers and non-movers among 2012 Republicans and Unaffiliated voters, respectively. Unaffiliated includes minor-party registrants. We classify a voter as a mover if the address in either available 2020 or 2021 snapshot is more than 100 meters from the 2012 address. Non-movers are those whose address is observed in both 2012 and either 2020 or 2021 and where all observed addresses are within 100 meters. This does not rule out an intervening move and return. All voters in the two parent samples can be classified. Exposure is still measured using the 2012 address and baseline ward composition, not the destination ward.

Panels A and B report the subgroup estimates; Panel C directly tests their differences using fully interacted pooled models with shared cluster identifiers. None of the individual coefficient-equality tests or the six joint equality tests rejects at the 5% level. Some individual differences are marginal at the 10% level: the closest to the 5% threshold concerns the Republican-exposure coefficient for remaining Unaffiliated in the joint specification ($p = 0.0506$). Failure to reject equality does not establish that the effects are equivalent.

Because residential mobility is measured after baseline, it may respond to exposure or share causes with partisan change. These conditional comparisons therefore do not identify causal effect differences between fixed mover types, nor do they isolate persistence of peer influence from residential selection.

Table F4: Peer Partisan Exposure by Residential Mobility: 2012 Republicans

	Democratic-share treatment			Republican-share treatment			Both treatments		
	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem
Panel A: Non-movers									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	−0.206*	0.100	0.106*				−0.167	0.062	0.105*
	(0.115)	(0.100)	(0.056)				(0.117)	(0.101)	(0.058)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.080	−0.040	−0.041*	0.060	−0.033	−0.027
				(0.052)	(0.047)	(0.021)	(0.053)	(0.048)	(0.022)
Joint test p-value		0.099			0.111			0.217	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.423	0.806	0.211
Dep. Var. Mean	-0.051	0.042	0.010	-0.051	0.042	0.010	-0.051	0.042	0.010
# Voters	28,975	28,975	28,975	28,975	28,975	28,975	28,975	28,975	28,975
# Boundaries	1,373	1,373	1,373	1,373	1,373	1,373	1,373	1,373	1,373
Panel B: Movers									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	−0.243	0.053	0.190**				−0.200	0.029	0.171**
	(0.189)	(0.177)	(0.078)				(0.203)	(0.187)	(0.082)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.110	−0.058	−0.052*	0.076	−0.040	−0.037
				(0.085)	(0.085)	(0.031)	(0.089)	(0.088)	(0.031)
Joint test p-value		0.045			0.156			0.128	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.613	0.964	0.134
Dep. Var. Mean	-0.219	0.194	0.025	-0.219	0.194	0.025	-0.219	0.194	0.025
# Voters	31,646	31,646	31,646	31,646	31,646	31,646	31,646	31,646	31,646
# Boundaries	1,416	1,416	1,416	1,416	1,416	1,416	1,416	1,416	1,416
Panel C: Movers minus non-movers									
$\beta_{\text{Dem}}^{\text{M}} - \beta_{\text{Dem}}^{\text{N}}$	−0.037	−0.046	0.084				−0.034	−0.033	0.066
	(0.211)	(0.205)	(0.096)				(0.227)	(0.216)	(0.099)
p-value: Dem. equality	0.860	0.821	0.382				0.882	0.880	0.503
$\beta_{\text{Rep}}^{\text{M}} - \beta_{\text{Rep}}^{\text{N}}$				0.030	−0.019	−0.011	0.016	−0.007	−0.009
				(0.099)	(0.097)	(0.039)	(0.104)	(0.100)	(0.038)
p-value: Rep. equality				0.763	0.847	0.769	0.879	0.948	0.810
Joint equality p-value		0.680			0.932			0.971	
Boundary FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Distance Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Specification, outcomes, and exposure units follow Table 3. Panels A and B split its harmonized sample by residential mobility, keeping baseline exposures fixed. A mover's observed 2020 or 2021 address is more than 100 meters from the 2012 address; non-movers have at least one valid displacement and none above 100 meters. Counts exclude singleton observations removed by the estimator. Standard errors in parentheses are two-way clustered by ward-boundary side and voter. Panel C reports mover-minus-non-mover differences from pooled models with all slopes and boundary fixed effects interacted with mobility; cluster identifiers remain shared. Equality p-values use clustered t or F reference degrees of freedom. Within-panel joint tests ask whether all exposure coefficients are zero; Panel C's joint test asks whether the corresponding coefficients are equal across mobility groups. Both stack the Republican and Democratic outcome equations; the Unaffiliated equation is redundant. Within-panel joint tests retain the main table's legacy denominator degrees of freedom: stacked observations minus estimated parameters. All p-values are unadjusted for multiple testing. Mobility is measured after baseline, so these are conditional subgroup comparisons, not identified causal differences between fixed mover types.

Table F5: Peer Partisan Exposure by Residential Mobility: 2012 Unaffiliated Voters

	Democratic-share treatment			Republican-share treatment			Both treatments		
	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem
Panel A: Non-movers									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	-0.255	0.069	0.186				-0.110	-0.111	0.221
	(0.302)	(0.313)	(0.181)				(0.305)	(0.327)	(0.183)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.396^{***}	-0.331^{**}	-0.065	0.347^{**}	-0.353^{**}	0.006
				(0.137)	(0.135)	(0.088)	(0.141)	(0.143)	(0.088)
Joint test p-value		0.489			0.014			0.079	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.511	0.231	0.284
Dep. Var. Mean	0.293	-0.389	0.096	0.293	-0.389	0.096	0.293	-0.389	0.096
# Voters	21,096	21,096	21,096	21,096	21,096	21,096	21,096	21,096	21,096
# Boundaries	1,335	1,335	1,335	1,335	1,335	1,335	1,335	1,335	1,335
Panel B: Movers									
$D_{ib}^{\text{Dem}} \times I_b^{\text{Dem}}$	-0.334	0.242	0.092				-0.257	0.329	-0.072
	(0.228)	(0.233)	(0.166)				(0.245)	(0.257)	(0.173)
$D_{ib}^{\text{Rep}} \times I_b^{\text{Rep}}$				0.176^{*}	-0.043	-0.134^{**}	0.135	0.004	-0.139^{*}
				(0.105)	(0.104)	(0.067)	(0.112)	(0.115)	(0.071)
Joint test p-value		0.340			0.076			0.164	
p-value: $\beta_{\text{Dem}} = -\beta_{\text{Rep}}$							0.687	0.297	0.304
Dep. Var. Mean	0.297	-0.410	0.113	0.297	-0.410	0.113	0.297	-0.410	0.113
# Voters	28,455	28,455	28,455	28,455	28,455	28,455	28,455	28,455	28,455
# Boundaries	1,409	1,409	1,409	1,409	1,409	1,409	1,409	1,409	1,409
Panel C: Movers minus non-movers									
$\beta_{\text{Dem}}^{\text{M}} - \beta_{\text{Dem}}^{\text{N}}$	-0.079	0.173	-0.094				-0.147	0.439	-0.292
	(0.371)	(0.389)	(0.249)				(0.384)	(0.413)	(0.250)
p-value: Dem. equality	0.831	0.657	0.706				0.701	0.287	0.242
$\beta_{\text{Rep}}^{\text{M}} - \beta_{\text{Rep}}^{\text{N}}$				-0.220	0.289^{*}	-0.069	-0.212	0.358^{*}	-0.146
				(0.166)	(0.168)	(0.114)	(0.176)	(0.183)	(0.115)
p-value: Rep. equality				0.186	0.087	0.546	0.228	0.051	0.205
Joint equality p-value		0.884			0.230			0.283	
Boundary FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Distance Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Specification, outcomes, and exposure units follow Table 3. Panels A and B split its harmonized sample by residential mobility, keeping baseline exposures fixed. A mover's observed 2020 or 2021 address is more than 100 meters from the 2012 address; non-movers have at least one valid displacement and none above 100 meters. Counts exclude singleton observations removed by the estimator. Standard errors in parentheses are two-way clustered by ward-boundary side and voter. Panel C reports mover-minus-non-mover differences from pooled models with all slopes and boundary fixed effects interacted with mobility; cluster identifiers remain shared. Equality p-values use clustered t or F reference degrees of freedom. Within-panel joint tests ask whether all exposure coefficients are zero; Panel C's joint test asks whether the corresponding coefficients are equal across mobility groups. Both stack the Republican and Democratic outcome equations; the Unaffiliated equation is redundant. Within-panel joint tests retain the main table's legacy denominator degrees of freedom: stacked observations minus estimated parameters. All p-values are unadjusted for multiple testing. Mobility is measured after baseline, so these are conditional subgroup comparisons, not identified causal differences between fixed mover types.

G Bishops: Data, Matching, and Design Validation

This appendix details the auxiliary bishop analysis summarized in Section 7.3: the data, the procedure matching bishops to their voter records, the difference-in-differences design, the treatment variation, and the validation tests behind the null results in Table 6.

G.1 Data

Bishop identities are only observable in our data from 2023 onward. Because our TargetSmart voter files end in 2021, we instead use the L2 voter file, available to us from 2015–2025, to study the effects of bishop partisanship on voters in their ward. Like TargetSmart, L2 is an annual panel of every registered Utah voter with a stable person identifier, party registration, vote history, and geocoded residence. Each voter is assigned to a ward by matching their geocoded residence to the nearest point[42] in a ward-boundary directory covering all Utah wards quarterly from 2023q2 through 2026q2 (97–98% of voters match within 100 meters); ward assignment is fixed at the voter's 2020 residence throughout.

For each (ward, quarter) pair we match the bishop's name to a voter record within his own ward. The matcher parses the directory's name string (handling leading initials and middle names), then searches the ward's voter roster for an exact first-and-last-name match, disambiguating multiple candidates by middle name or initial, and falling back to adjacent voter-file years for bishops not found in the contemporaneous snapshot. Across 11 quarters and roughly 4,450 wards per quarter, the procedure matches about 79% of ward-quarters to a unique voter record, identifying 5,785 distinct bishops; unmatched cases are mostly indeterminate multiple-candidate names, which we exclude. Matched bishops are excluded from all member-level regression samples.

Each bishop's characteristics—party registration and turnout—are measured from his own 2020-election record (the 2021 snapshot). Fixing bishop characteristics at their pre-period values means the ward-level bishop treatment can change only when the person serving as bishop changes, never through an incumbent bishop's own registration switching.

G.2 Treatment Variation

Bishops are lay leaders selected from among a ward's members to preside over their congregation. They typically serve for a period of about five years, after which they are replaced with another ward member. Turnover is therefore routine and, critically for our design, its

[42] While our 2013–2017 ward boundary data provided us with ward polygons, our 2023–present ward boundary data comes in the form of address points attached to ward assignments.

timing is driven by the rotation norm rather than by congregational politics.

Figure G1 summarizes the treatment variation. Panel (a): at baseline (2023q2), among wards whose bishop is matched to a voter record with observable partisanship, 79.9% have a Republican bishop, 18.7% an Unaffiliated bishop (registered with no party or a third party), and 1.4% a Democratic bishop; the distribution among wards whose bishop subsequently turned over is nearly identical, indicating turnover is not selected on bishop partisanship. In the five quarters preceding the 2024 election (2023q2–2024q3), 24.9% of the 3,028 party-identified wards changed bishops. Panel (b) shows the incoming bishop's party conditional on the outgoing bishop's party among changed wards, pooling Democrats with Unaffiliated bishops as non-Republican because only six baseline bishops are Democrats. The new bishop's party is nearly independent of the old bishop's: 79% of wards that lose a Republican bishop and 76% of wards that lose a non-Republican bishop draw a Republican replacement. Appendix G.3 examines this pattern in more detail. Panel (c) plots the probability a ward has a Republican bishop against the ward's member Republican share: bishops are systematically more Republican than their congregations everywhere (the profile sits far above the 45-degree line), and the gradient in ward composition is modest.

G.3 Who Becomes Bishop

The design in Section 7.3 uses within-ward changes in the bishop's party due to bishop turnover. The key assumption is that the choice of a new bishop, specifically whether or not he's a Republican, is unrelated to underlying trends in the ward's partisanship. This section explores how bishops are chosen with a goal of assessing the validity of this assumption.

We take the roster of bishops and stake presidents serving in 2023, measure each leader's characteristics in the 2021 voter file, and compare him to the adult men registered in his own unit. A conditional logit treats each calling as one choice from that pool, the man called plus 100 randomly drawn eligible men from the same unit, so the unit's composition is conditioned on by construction and the coefficient on an attribute is the odds premium it carries within the unit. Age enters linearly and squared, and race enters as a covariate, so age-based eligibility is absorbed rather than imposed. Home values come from CoreLogic and are expressed as the leader's percentile within his own unit.

Bishops are drawn disproportionately from the Republican and wealthier men of their wards. Within the same ward, and holding age and race fixed, a Republican man has about 2.8 times the odds of being called bishop that an Unaffiliated man has, and a Democrat about three-tenths the odds; moving from the bottom to the top of the ward's home-value distribution multiplies the odds of being called by about 3.4. Selection is stronger on every

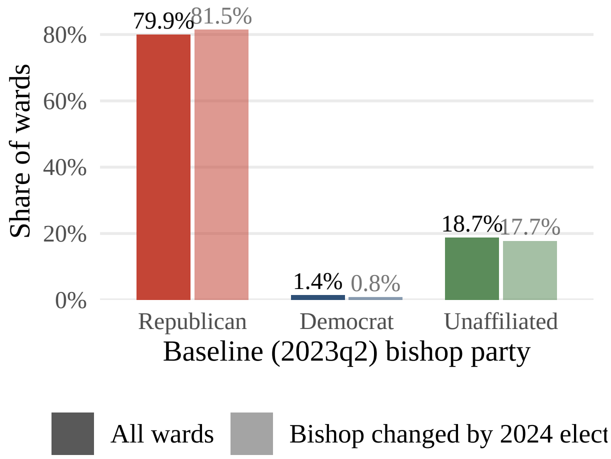

(a) Baseline bishop party: all vs. changed wards

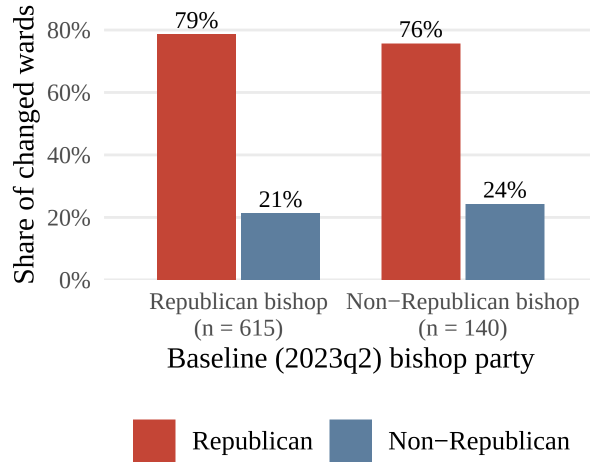

(b) New bishop's party by outgoing bishop's party

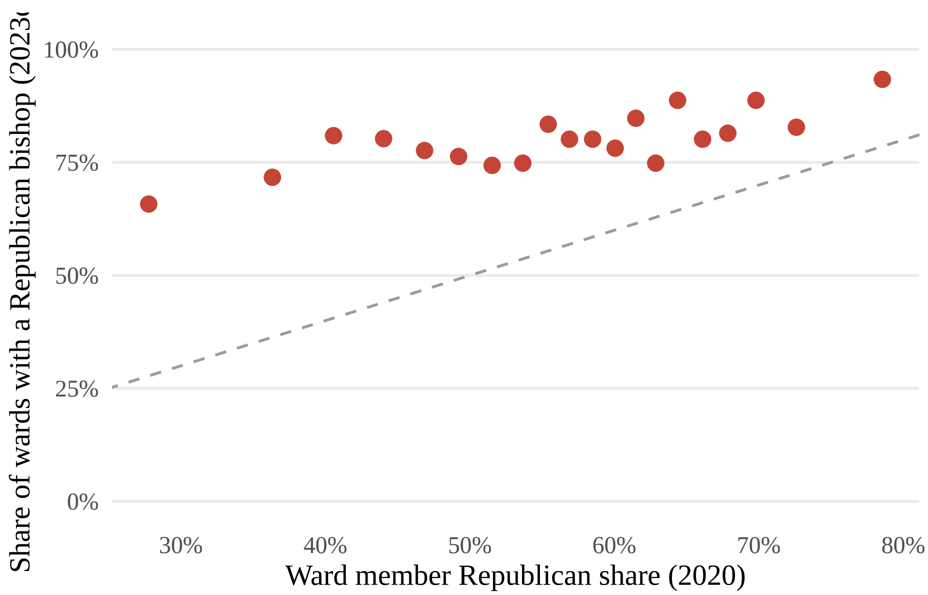

(c) Bishop party vs. ward member composition

Figure G1: Bishop Treatment Variation. Panel (a): distribution of the baseline (2023q2) bishop's party across all matched wards and across wards whose bishop changed before the 2024 election. Panel (b): the incoming bishop's party conditional on the outgoing bishop's party, among changed wards, with Democrats and Unaffiliated bishops pooled as non-Republican. Panel (c): binned scatter of the share of wards with a Republican bishop against the ward's 2020 member Republican share; dashed line is the 45-degree line.

dimension for stake presidents, for whom the Republican premium is 4.3 and the home-value premium roughly 15. Stated as a gap, bishops are about 22 percentage points more Republican than the registered voters of their own wards, with the shortfall falling on the Unaffiliated rather than on Democrats. One caveat applies to every party estimate here. The choice sets are registered voters, because the voter file does not identify religion, whereas callings go to active Latter-day Saints, who are likely more Republican than their non-member neighbors. The Republican premium therefore absorbs that difference on top of any preference for Republicans among those extending callings, and should be read as an upper bound on such a preference. The two exercises that follow compare bishops within the same ward and are not affected by unobserved religiosity.

The stake president selects the bishops of his stake.[43] We test whether a stake presidents' own characteristics predict the characteristics of the bishops he selects. Only nine matched stake presidents are Democrats, so we contrast Republican with non-Republican presidents, the latter overwhelmingly Unaffiliated. Table G1 shows a precise null. The Republican premium among candidates is 3.10 under Republican presidents and 3.08 under non-Republican ones; the interaction has an odds ratio of 1.00 ($p = 0.98$), and 0.95 once the ward's Republican share is allowed to shift the premium, which matters because non-Republican presidents lead less Republican wards. Wealthier presidents show a somewhat larger home-value premium (4.06 against 3.09), but the interaction is not significant.

Next we compare a ward's incoming bishop to its outgoing one. The ward's composition is held fixed and no reference population is needed, so this is the cleanest comparison available. Whoever selected the outgoing bishop revealed a willingness to choose a man with a given attribute; if the next bishop of the same ward shares it, that points to a preference. Both bishops are measured in the same 2021 snapshot so that the comparison is not confounded with measurement year. Of 2,788 within-ward transitions, 886 have both bishops' parties observed. Table G2 reports the results. The incoming bishop is Republican with probability 0.80 when the outgoing bishop was Republican and 0.81 when he was not; the persistence coefficient is -0.015 (s.e. 0.034), so any persistence above about five percentage points is ruled out, and it does not change when the ward's Republican share is controlled or when the sample is restricted to transitions where the same president made both calls. This null is conservative since two men drawn from the same eligible pool are mechanically correlated even absent any preference, which biases the statistic toward persistence. Home value does persist (a slope of 0.15, well below one), which could reflect persistence or the shared-pool bias.

[43]Technically stake presidents recommend bishops to church headquarters, who ultimately hold responsibility for approving them.

Together, the two exercises imply that bishops' Republican tilt comes from the composition of the eligible pool rather than from callers selecting on party. Which party a ward's next bishop belongs to is not predicted by the caller's own party or by the previous bishop's, which is what the difference-in-differences of Section 7.3 requires of its within-ward variation.

Table G1: Bishop Selection, by the Calling Stake President's Own Characteristics

	Odds ratio	Strata
Republican candidate premium, by president's party		
Republican president	3.10 [2.77, 3.47]	2,191
Non-Republican president	3.08 [2.42, 3.92]	447
Interaction	1.00 ($p = 0.98$)	2,638
Interaction, ward share controlled	0.95 ($p = 0.67$)	2,638
Home-value premium, by president's own percentile		
Less wealthy president	3.09 [2.44, 3.91]	978
Wealthier president	4.06 [3.18, 5.18]	979
Interaction	1.67 ($p = 0.15$)	1,957

Notes: Conditional logit within the ward's pool of eligible men, specified as in Appendix G.3, with 95% confidence intervals in brackets. The first two rows of each block estimate the premium separately for bishops called by each type of president; the interaction row puts both groups in one model and reports the odds ratio on (candidate attribute) × (president type). President attributes are constant within a stratum, so only interactions are identified. The "ward share controlled" row adds (candidate is Republican) × (ward's Republican share). The wealth split is at the median president home-value percentile.

Table G2: Does the Outgoing Bishop Predict the Incoming One?

	Estimate	Std. error	n
Party			
P(new Republican | old Republican)	0.799	—	730
P(new Republican | old not Republican)	0.814	—	156
Persistence	−0.015	(0.034)	886
Persistence, ward Republican share controlled	−0.026	(0.034)	877
Persistence, same president made both calls	−0.038	(0.061)	269
Home-value percentile within the ward			
Mean, outgoing bishop	0.650	—	779
Mean, incoming bishop	0.632	—	779
Persistence	0.150	(0.038)	779
Persistence, same president made both calls	0.121	(0.069)	221

Notes: Within-ward bishop transitions, 2023–2026. "Persistence" is the coefficient from regressing the incoming bishop's attribute on the outgoing bishop's. Standard errors cluster on stake. Both bishops are measured in the 2021 voter snapshot and the 2023 CoreLogic vintage, and the home-value percentile is computed against the same ward distribution for both men.

G.4 Design

We estimate a two-period difference-in-differences at the member level, comparing registration outcomes in the 2020 (pre) and 2024 (post) elections:

$$Y_{iwt} \;=\; \beta\, B_{w(i),t} \;+\; \alpha_i \;+\; \tau_t \;+\; \varepsilon_{iwt}, \tag{24}$$

where Y_{iwt} is a registration indicator (Republican, Democrat, or Unaffiliated), $B_{w(i),t}$ is the serving bishop's fixed 2020 characteristic (the 2023q2 bishop in the pre period, the bishop serving as of 2024q3 in the post period), and α_i and τ_t are individual and election-year fixed effects. Standard errors are clustered by ward. The sample comprises the roughly one million members observed in both periods, split by 2020 party registration; bishops themselves are excluded. Because B is fixed at the bishop's own pre-period record, β is identified purely off wards whose bishop turned over before the election, with never-changing wards (and wards whose turnover preserved the characteristic) serving as controls.

G.5 Validation and Results

Tables 6, G3, and G4 report the estimates for 2020 Republicans, Unaffiliated voters, and Democrats, each with three rows per treatment: the main 2020→2024 estimate, a parallel-trends test applying the same (future) treatment to the 2016→2020 outcome change, and a placebo replacing the treatment with bishop changes occurring only after the 2024 election (2025q1–2026q2), which cannot have caused the outcomes. Every main estimate is a precise null, and the pre-trend and placebo rows are null as well, supporting the design. For 2020 Republicans, the 95% confidence interval on the Republican-bishop treatment excludes effects on Republican registration larger than 0.6 percentage points; for 2020 Unaffiliated voters, larger than 0.7 percentage points.

The magnitude benchmark in Section 7.3 follows from these bounds. The peer coefficients in Table 3 are scaled to a change in the ward's partisan share from zero to one. The most such a change could do through the bishop is raise the probability of having a Republican bishop from zero to one, so the bishop channel's contribution to a peer coefficient is at most the effect of acquiring a Republican bishop. Using the upper end of the confidence interval on that effect as its ceiling, the share of the peer effect that bishops could account for is that bound divided by the peer coefficient: for Unaffiliated voters' Republican registration, 0.7 percentage points against 22.8 points, or 3.1%; for Republicans' retention, 0.6 points against 7.9 points, or 7.9%. Two features make these bounds generous. Bishops' partisanship rises only modestly with their wards' (Figure G1c), so a one-for-one response overstates how much a composition shift moves the bishop's party. And the Republican denominator is the effect of Republican exposure on remaining Republican, which is not statistically distinguishable from zero, so the Republican ratio carries wide uncertainty and the Unaffiliated bound is the informative one.

Table G3: Bishop Partisanship and Member Registration — 2020 Unaffiliated voters

	Republican bishop			Democratic bishop		
	ΔRep	ΔDem	ΔUnaff	ΔRep	ΔDem	ΔUnaff
Bishop party (2020 → 2024)	0.001	−0.005*	0.004	−0.009	−0.003	0.012
	(0.003)	(0.003)	(0.004)	(0.010)	(0.010)	(0.011)
Pre-trend (2016 → 2020)	0.001	−0.012	0.011*	−0.008	−0.009	0.017
	(0.005)	(0.008)	(0.006)	(0.014)	(0.027)	(0.024)
Placebo (post-election change)	−0.006	0.002	0.004	0.013	0.002	−0.015
	(0.004)	(0.002)	(0.003)	(0.013)	(0.006)	(0.010)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep. Var. Mean	0.063	0.020	0.918	0.063	0.020	0.918
# Voters	323,984	323,984	323,984	323,984	323,984	323,984
# Wards	3,364	3,364	3,364	3,364	3,364	3,364

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each cell is a separate member-level difference-in-differences with individual and election-year fixed effects (Equation (24)), comparing the change in the registration outcome named in the column header between the 2020 and 2024 elections across wards that did and did not acquire a bishop of the party named in the block header. Bishop party is measured from the bishop's own 2020 voter record. The first row is the main estimate. The second row applies the same 2023–2024 bishop change to the 2016–2020 registration change, which it cannot have caused, as a test of pre-trends. The third row replaces the treatment with bishop changes that occurred only after the 2024 election (2025q1–2026q2), which likewise cannot have caused the 2020–2024 outcome. Standard errors clustered by ward in parentheses. The sample is voters registered with the party named in the caption in 2020 and observed in both elections; bishops themselves are excluded.

Table G4: Bishop Partisanship and Member Registration — 2020 Democrats

	Republican bishop			Democratic bishop		
	ΔRep	ΔDem	ΔUnaff	ΔRep	ΔDem	ΔUnaff
Bishop party (2020 → 2024)	−0.006*	0.000	0.006*	−0.005	0.006	−0.001
	(0.003)	(0.004)	(0.004)	(0.009)	(0.013)	(0.011)
Pre-trend (2016 → 2020)	−0.007	0.005	0.003	−0.005	0.004	0.001
	(0.007)	(0.006)	(0.004)	(0.022)	(0.017)	(0.009)
Placebo (post-election change)	−0.001	0.002	−0.002	0.022	0.000	−0.022**
	(0.004)	(0.005)	(0.004)	(0.016)	(0.014)	(0.009)
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Dep. Var. Mean	0.036	0.921	0.043	0.036	0.921	0.043
# Voters	133,760	133,760	133,760	133,760	133,760	133,760
# Wards	3,357	3,357	3,357	3,357	3,357	3,357

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Each cell is a separate member-level difference-in-differences with individual and election-year fixed effects (Equation (24)), comparing the change in the registration outcome named in the column header between the 2020 and 2024 elections across wards that did and did not acquire a bishop of the party named in the block header. Bishop party is measured from the bishop's own 2020 voter record. The first row is the main estimate. The second row applies the same 2023–2024 bishop change to the 2016–2020 registration change, which it cannot have caused, as a test of pre-trends. The third row replaces the treatment with bishop changes that occurred only after the 2024 election (2025q1–2026q2), which likewise cannot have caused the 2020–2024 outcome. Standard errors clustered by ward in parentheses. The sample is voters registered with the party named in the caption in 2020 and observed in both elections; bishops themselves are excluded.

H Additional Results

This appendix supplements the condensed exhibits of Sections 6–7, contain in order: the full same- versus opposite-demographic decomposition, the retained- versus newly-assigned-peer estimates, the complete individual-characteristic heterogeneity grids, and turnout effects for all three baseline-party subsamples.

H.1 Same- and Opposite-Demographic Peer Exposure: Full Grid

Figure H1 completes the same/opposite-demographic decomposition of Section 7.4 with the two own-party exposure–sample combinations omitted from the main text; the two out-party combinations appear in the main text as Figure 8.

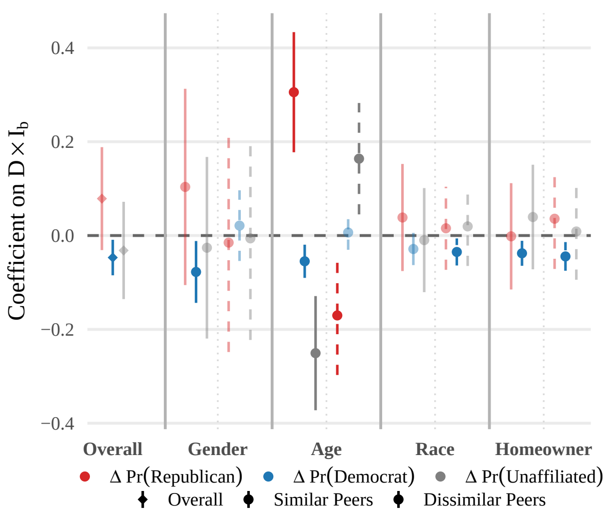

(a) Republican Exposure for 2012 Republicans

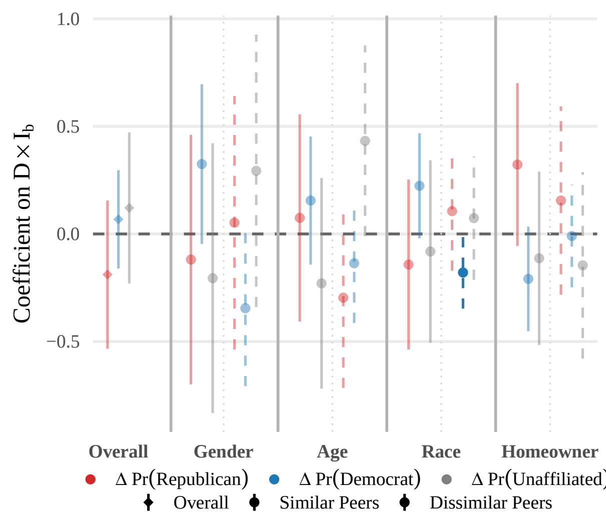

(b) Democratic Exposure for 2012 Unaffiliated voters

Figure H1: Same- and Opposite-Demographic Peer Exposure Effects: Own-Party Exposures. The two out-party combinations, which the main text discusses, appear in Figure 8. Panel construction follows Figure 8.

H.2 Retained and Newly Assigned Peer Exposure: Full Estimates

Table H1 reports the retained-peer and new-peer exposure coefficients from the disruption specification (Appendix D, Equation (18)) for all baseline-party subsamples, exposures, and registration outcomes.

Table H1: Retained vs. Newly Assigned Peer Exposure: Full Grid

		Retained peers			New peers		
Sample	Exposure	ΔRep	ΔUnaff	ΔDem	ΔRep	ΔUnaff	ΔDem
2012 Republicans	Democratic	-0.309^{**}	0.191	0.118^{**}	-0.239	0.077	0.161^{***}
		(0.154)	(0.141)	(0.053)	(0.158)	(0.134)	(0.060)
	Republican	0.209^{***}	-0.136^{**}	-0.072^{***}	0.145^{**}	-0.086	-0.059^{***}
		(0.057)	(0.054)	(0.020)	(0.059)	(0.055)	(0.020)
2012 Unaffiliated voters	Democratic	0.130	-0.329^{*}	0.200	-0.068	0.152	-0.083
		(0.196)	(0.191)	(0.137)	(0.192)	(0.195)	(0.129)
	Republican	0.164^{**}	0.028	-0.192^{***}	0.169^{**}	-0.085	-0.084^{*}
		(0.079)	(0.078)	(0.048)	(0.080)	(0.080)	(0.050)
2012 Democrats	Democratic	0.463	0.189	-0.651^{*}	-0.195	-0.311	0.507
		(0.291)	(0.303)	(0.369)	(0.292)	(0.323)	(0.380)
	Republican	-0.100	-0.296^{**}	0.396^{**}	0.159	-0.235^{*}	0.076
		(0.146)	(0.130)	(0.161)	(0.140)	(0.131)	(0.170)

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors two-way clustered by ward-boundary side and voter in parentheses.

H.3 Individual-Characteristic Heterogeneity

Figures H2 and H3 report treatment effects interacted with the voter's own baseline characteristics (age group, gender, race, rurality, homeownership, home size, prior turnout, and name commonness) for the two main exposure–sample pairs (Democratic exposure among 2012 Republicans; Republican exposure among 2012 Unaffiliated voters). Estimates are largely homogeneous across margins; in particular, effects appear across all age groups.

H.4 Turnout

Table H2 reports the effects of partisan peer exposure on marginal 2020 general-election and primary-election turnout for all three baseline-party subsamples. The general-election estimates are uniformly insignificant, while Republican exposure raises 2020 primary turnout among 2012 Unaffiliated voters (25.4 percentage points, $p = 0.004$) and 2012 Democrats (34.6 percentage points, $p = 0.046$)—the marginal counterparts of the interacted cells in

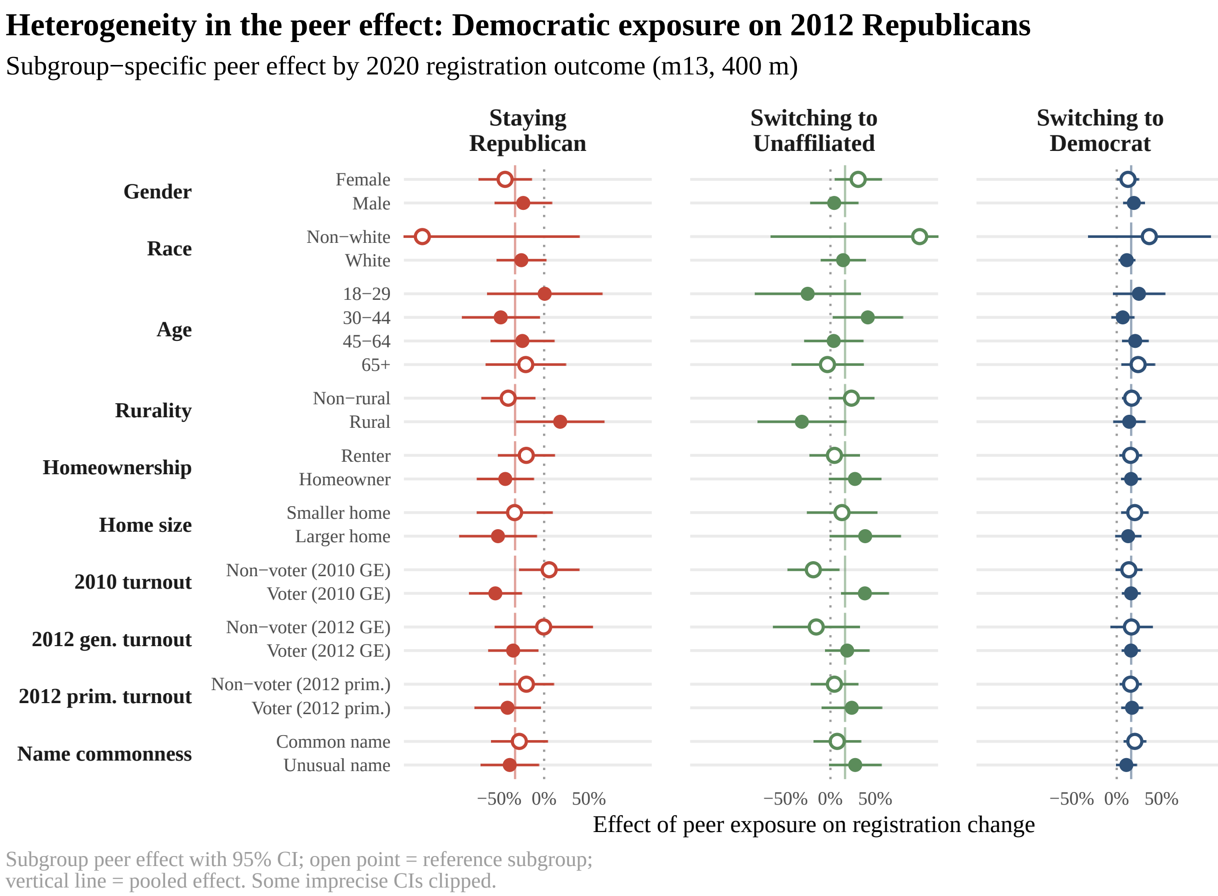

Figure H2: Heterogeneity by Individual Characteristics: Democratic Exposure, 2012 Republicans. Points are per-subgroup treatment effects with 95% confidence intervals; the vertical line marks the pooled estimate.

Section 7.2. Figure H4 plots statewide turnout trends for a balanced panel of voters observed in 2012, 2016, and 2020, by 2012 party registration (panels (a)–(b)) and by 2012–2020 Republican-status transition (panels (c)–(d)). Voters who switched into the Republican Party between 2012 and 2020 raised their primary turnout from 3.7% to 62.0%, against 3.4% to 13.7% among voters who stayed non-Republican; voters who switched between 2012 and 2016 show the same jump in the 2016 cycle (from 3.9% to 35.3%, against 13.8% among those who had not switched).

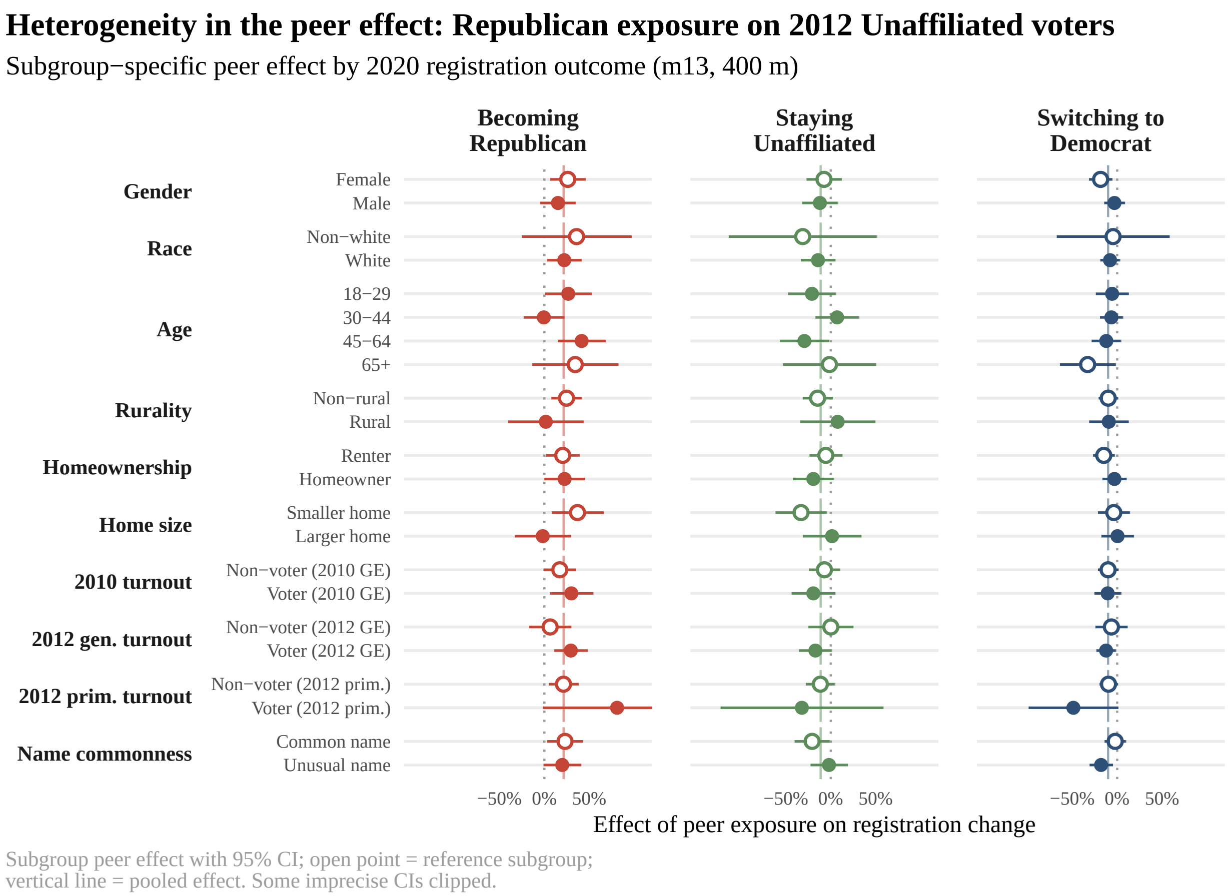

Figure H3: Heterogeneity by Individual Characteristics: Republican Exposure, 2012 Unaffiliated voters. See Figure H2 notes.

Table H2: Peer Partisan Exposure and 2020 Turnout: Full Grid

	2012 Republicans	2012 Unaffiliated voters	2012 Democrats
Panel A: Voted in 2020 general election			
Republican exposure	-0.096	0.224^{*}	-0.098
	(0.075)	(0.128)	(0.184)
Democratic exposure	0.210	-0.348	-0.320
	(0.156)	(0.264)	(0.467)
Dep. Var. Mean	-0.008	0.042	-0.018
Panel B: Voted in 2020 primary election			
Republican exposure	-0.147	0.254^{***}	0.346^{**}
	(0.097)	(0.089)	(0.173)
Democratic exposure	-0.066	-0.094	0.548
	(0.214)	(0.194)	(0.469)
Dep. Var. Mean	0.244	0.227	0.218

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors two-way clustered by ward-boundary side and voter in parentheses.

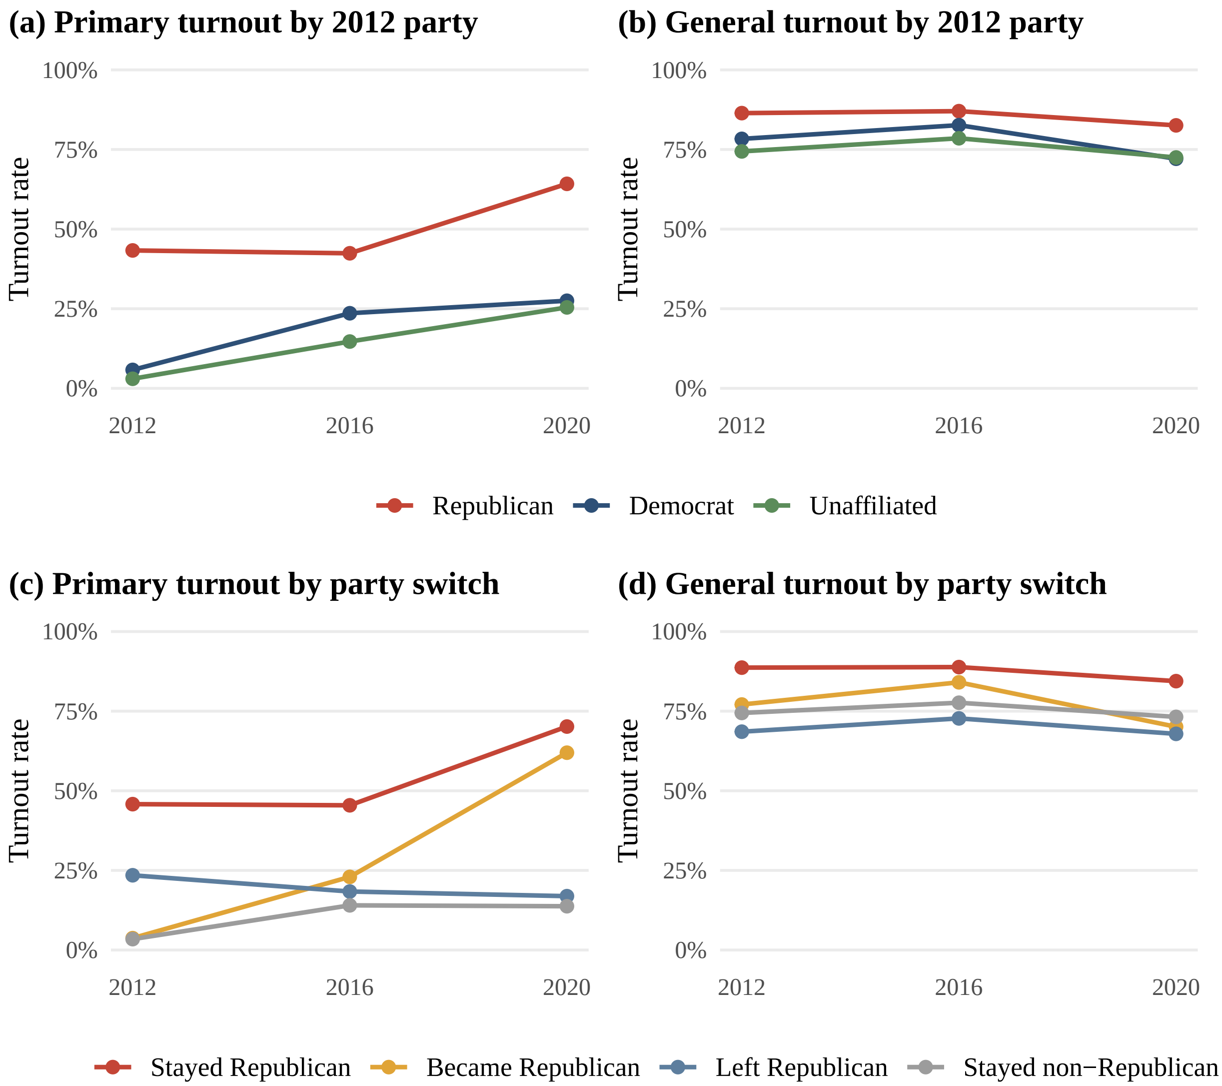

Figure H4: Turnout rates over time, by 2012 party and by 2012–2020 party switch. Panels (a)–(b) plot primary- and general-election turnout across 2012, 2016, and 2020 for a balanced statewide panel of registered voters observed in all three years, split by 2012 party registration. Panels (c)–(d) split the same voters by their 2012-to-2020 Republican-status transition. Turnout is the share of the group recorded as voting in each election. Voters who switch into the Republican Party show a sharp rise in primary turnout (panel c) but no comparable change in general-election turnout (panel d).