

The Geographic Dimensions of Electoral Polarization in the 2004 U.S. Presidential Vote

Ian Sue Wing

Dept. of Geography & Environment, Boston University

Joan L. Walker

Dept. of Civil & Environmental Engineering, U.C. Berkeley

Introduction

The 2004 U.S. presidential election was one of the most divisive in recent history (Pew Research Center 2004). The divisions in the electorate are popularly seen as the culmination of a process of political polarization underway since the 1970s (e.g., Frank 2004), and are epitomized by the now-ubiquitous map of the U.S. which shows swaths of red (i.e., majority Republican) states in the center of the country surrounded by blue (i.e., majority Democratic) states on the east and west coasts and in the north central region. In this chapter we investigate the geographic dimensions of political polarization in the U.S. through the lens of the 2004 election. We elucidate the principal contours of the divisions in the electorate, and characterize the manner in which the effects of the correlates of voting behavior cluster regionally. We take an ecological approach, using spatial econometrics to estimate the interregional divergence in the influences of the characteristics of populations and places on the odds of the Republican vote. To this end we employ aggregated data on 3106 counties in the lower 48 states, which is the finest spatial scale at which both electoral returns and a variety of demographic and contextual variables are readily available.

Our goal is to push the limits of ecological analysis in electoral geography. We first develop a theoretical framework in which geography plays a central role in electoral polarization. Our central hypothesis, which draws on themes in the political science literature (Johnston et al. 2004; Cho and Rudolph 2008), is that a number of social processes that operate at fine spatial scales tend to push individuals voters' views into closer alignment with the ideological preferences of their geographically proximate majority—a phenomenon we call “localized entrenchment”. Drawing on the sociological literature on polarization (Evans et al. 1996; Evans 2003), we circumvent the well-documented handicap of weak correlation between demographic attributes and ideology by employing a richer array of explanatory variables than prior spatial statistical analyses (e.g., O’Loughlin et al 1994). We then apply spatial statistical techniques that exploit the spatial interrelationships among the electoral returns and our set of covariates, and find strong indications of entrenchment. Finally, we employ advanced methods to characterize the spatial heterogeneity in our estimated relationships—rather than re-estimate our aggregate statistical model on different regional sub-samples, we use geographically weighted regression (GWR). This technique enables us to exploit the spatial interdependencies among the entire universe of counties to estimate the fine-scale geographic variation in our covariates’ influences on the 2004 presidential vote, while simultaneously controlling for the underlying spatial distributions of the characteristics of people and places. The patterns of agglomeration in the resulting influences on voting behavior are consistent with our explanation of how local entrenchment might induce polarization of the electorate.

We report three sets of results. First, we construct Local Moran’s I statistics to analyze the spatial clustering of county election returns. We find substantial spatial autocorrelation in voting patterns, evidence that the returns for democrats and Republicans were significantly clustered in different regions of the U.S., and indications of divergence among different sub-populations’ vote distributions based on the spatial clustering of their demographic characteristics. Second, we perform spatial regressions at the aggregate level which identify the demographic and contextual factors that significantly impact the odds of voting Republican. We partition this propensity into direct influences associated with the attributes of counties and their populations, and indirect influences associated with the voting behavior and demographic characteristics of neighboring jurisdictions. The latter effects are particularly large, in many cases outweighing the former, and highlight the importance of the geographically-varying contextual factors that are central to the predictions of our core hypothesis. Finally, our GWR results indicate considerable heterogeneity in the influence of several of our explanatory variables, and, most tellingly, regional agglomeration

in their signs, which suggests that electoral polarization manifests itself as cross-cutting divisions in the U.S. electorate, not between population sub-groups but within sub-groups over space.

Given the nature of our analysis and results, we raise two caveats at the outset. First, when it comes to uncovering the mechanisms through which polarization occurs, we barely scratch the surface. Our more modest objectives are to clarify the irreducible spatial components of the divisions in the American electorate, and to outline their broad contours as the first phase of a program of more rigorous statistical testing. The second caution concerns the ecological fallacy. In particular, we take pains to distinguish what we do find: divergent patterns of spatial clustering in the impacts of the characteristics of counties and their populations on the vote, from what we do not: how the sign and magnitude of the effects of individuals' characteristics on their own voting behavior vary over space. The distinction between these inferences cannot be too sharply drawn (Goodman 1953; Hanushek and Jackson 1974).

The remainder of the chapter is organized into four sections. In section 2 we set the stage by discussing our motivations and framing our inquiry. In section 3 we describe the sources of data used in our analysis, and illustrate the spatial heterogeneity in key variables. We outline our methods of analysis in Section 4 and present and discuss the results in section 5. We conclude in section 6 with directions for future research.

Entrenchment: Geography's Role in Political Polarization

Political polarization is the segregation of the electorate along issue opinion and/or ideological lines, with concentration of voters about opposing extreme positions and concomitant erosion of moderate "centrist" preferences. The phenomenon is illustrated in Fig. x.1, which plots the distribution of preferences in the electorate on a left-leaning (liberal) versus right-leaning (conservative) scale. In panel A, which draws on Fiorina and Abrams (2008), distribution A-I is not polarized, and exhibits the classic "single peaked" preferences of a centrist majority. By contrast, the bimodal distribution A-II, shown by the dashed line, illustrates the polarization of voters into equal opposing factions. The gray Distribution A-III, about which we say more below, is intermediate between A-I and A-II, with fatter tails and a less distinct peak indicating voters' movement away from the center toward the extremes.

The reality of the U.S. electoral landscape is far more complex than this picture suggests, however. Despite considerable heterogeneity in American

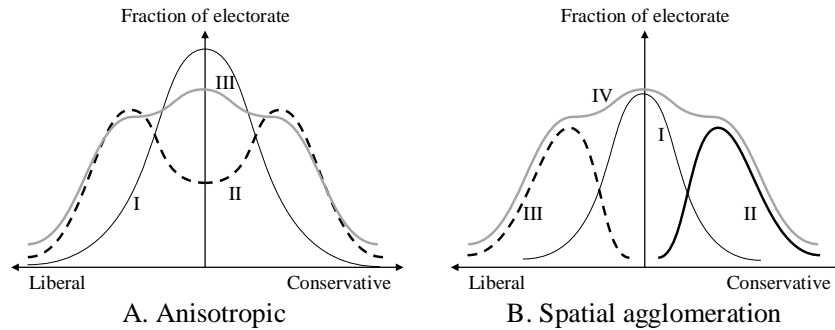
voters' attitudes and beliefs, there is no evidence that the distribution of the electorate is either bimodal, or has recently become substantially more disperse—especially in light of the long view of history (Ansolabehere et al. 2006; Fiorina et al. 2006; Fiorina and Abrams 2008; Klinkner 2004; Klinkner and Hapanowicz 2005). There is, however, abundant evidence that the parties' candidates and activists alike have become increasingly partisan, and have staked out increasingly divergent positions on a range of issues (Bartels 2000; Fleisher and Bond 2004; Mellow and Trubowitz 2005; Poole and Rosenthal 2001; Stonecash, Brewer and Mariani 2002).¹ Also, there are at best only weak indications of a rising intensity of opposing political views amongst the general electorate, and then only for a handful of “hot-button” issues such as abortion or homosexual persons' right to marry (DiMaggio et al. 1996; Evans 2003; Fiorina and Abrams 2008), while the distribution of liberal and conservative views on the broad spectrum of issues appears to have remained fairly even (Fiorina et al. 2006). What has occurred is a public redefinition of the labels “liberal” and “conservative”, which, along with polarized choices offered in the political arena, has served to heighten issues' salience to voters, and has induced them to self-categorize and align more closely with one or the other party despite unchanged underlying preferences (Hetherington 2001; Baldassarri and Gelman 2008; Miller and Hoffmann 1999). These dynamics have led many analysts to conclude that Americans feel more polarized than they in fact are.

Our own view is that while the “red versus blue state” conception of polarization is undeniably simplistic,² to claim that the electorate is not di-

¹ Glaeser et al (2005) develop a theory of strategic extremism which illustrates the incentives political parties have to divide on issues in order to increase their chances of winning at the polls. Partisanship turns on two key elements: among voters, the existence of an intensive margin where the level of support matters (e.g., turnout or donations, as opposed to the extensive margin of voting) and which parties can activate by taking extreme positions that appeal to their respective bases, and the ability of parties to target extreme statements to their own supporters while bypassing those of the opposition, thereby avoiding a backlash. Below, we note that this sort of targeting becomes easier the more the electorate is ideologically segregated along geographical lines.

² This is an example of the modifiable areal unit problem (Openshaw 1984). Differences between the number of Democratic and Republican votes were as large between red and blue counties within some states as they were between some red and blue states. Using counties as the unit of analysis is attractive precisely because, unlike states, congressional districts or electoral precincts, their geographic boundaries are independent of electoral processes relevant to the presidential vote. The consequent absence of selection bias makes us confident in

Fig. x.1 Electoral Polarization: A Conceptual Framework



vided is to deny the essential geographic dimension of the phenomenon. Klinkner (2004) cites a particularly apposite example which captures the essence of the phenomenon. In 1972, *New Yorker* magazine contributor Pauline Kael expressed surprise at Richard Nixon’s re-election as president, saying “Nobody I know voted for him”. The same could be said in 2004. Despite the fact that Republican incumbent George Bush was returned to office with 52 percent of the national electorate, few people in Washington DC knew anyone who voted for him—he gained just 7 percent of the electorate there. Likewise, few people in Idaho’s Madison County knew anyone who voted for Democratic challenger John Kerry, whose record there was similarly dismal. And although Bush won by a margin of 60 percent or greater in 54 percent of counties while Kerry enjoyed a similar margin in only 5 percent of counties, these “landslide” jurisdictions were home to 47 percent of the electorate. Moreover, 38 out of 50 states were carried by one or the other candidate with a margin of 5 percentage points or greater, with a stark divergence in the attitudes and beliefs espoused by the voters in states with Republican and Democratic majorities (e.g., Abramowitz and Saunders 2008: Table 6; Glaeser and Ward 2006: Table 1).

The geographic evidence typically adduced in support of the no-polarization thesis is that county-level returns exhibit variances and indices of dissimilarity that are low and stable, as well as indices of isolation for each party’s turnout that are similar in magnitude and fluctuate with no apparent long-run trend (Glaeser and Ward 2006; Klinkner and Hapanowicz 2005). But these same data indicate substantial geographic clustering

exploiting county characteristics as strictly exogenous covariates in our subsequent analyses.

of voting patterns in recent presidential elections (Kim et al. 2003), a phenomenon which persists into 2004. Democratic and Republican voters were more likely than not to be exposed only to individuals who voted in a similar way, with the result that one fifth of those supporting either party would have needed to relocate for the distribution of votes to be spatially uniform. The latter is the highest percentage since the 1940s (Glaeser and Ward 2006, Fig. 2).³

The statistical manifestation of this sort of division is shown in Fig. x.1B, which illustrates a hypothetical situation in which the electorate is divided among two disjoint regions. Distribution B-I indicates the preferences of centrist voters, whose members are distributed among both regions. A conservative-leaning sub-population of voters with distribution B-II resides in one region, while a liberal-leaning sub-population with distribution B-III resides in the other. It is easy to see that in this society the aggregate preferences B-IV are the same as the intermediate distribution A-III, with zero mean and fair degree of central tendency, but with an electorate that feels—and is—polarized, but along geographic lines.

Our main contention is that this picture describes the 2004 presidential election, not in Kael's literal sense of the regional distributions of electoral returns, but rather in terms of the preferences that DiMaggio et al. (1996), Evans (2003) and others have sought to measure.⁴ Because of the ecological nature of our data, our indicators of preference boil down to the influences of the characteristics of populations and places on the propensity to vote Republican or Democratic. Indeed, we demonstrate that along a number of key dimensions the influence of characteristics on the propensity to vote exhibit substantial spatial agglomeration, with geographic clustering of counties with divergent preferences, as in Fig. x.1B.

In conducting our investigation we take up the gauntlet thrown down by Fiorina and Abrams (2008), demonstrating the strength and stability of the associations between the voting behavior on one hand and the characteristics of populations and places on the other.⁵ Consistent with our interest in

³ These statistics were computed for our sample of 3106 counties in the lower 48 states. Indices of isolation measure the likelihood of Republicans' and democrats' exposure to the opposing group at 55% and 51%, respectively, while non-uniformity in the pattern of votes is given by the index of dissimilarity at 22%.

⁴ Note that B-IV's variance and excess kurtosis are larger than A-II's. These authors test whether these two moments of the distributions of survey respondents' attitudes on a diverse array of social issues have increased over time.

⁵ "...contrasts in individual sociocultural characteristics are not direct indicators of political polarization. Hence, contrasts in such characteristics may or may not constitute evidence of polarization. Analysts must provide additional information about the strength of the links between social characteristics and relevant

the segregation of the electorate over space, and cognizant of the strictures imposed by the cross-sectional data at our disposal, we reinterpret their (temporal) notion of stability to focus on how the correlates of voter behavior vary geographically.

Our first task is to articulate testable propositions about how regionally segregated voter distributions like B-II and B-III might arise. Shifts in the American electorate at broad geographic scales have been well documented, with sorting and clustering of individuals with similar ideological leanings arising as unintended consequence of interstate migration (Frey 2000; Gimpel and Schuknecht 2001), as well as (intentional) ideological realignment in various regions of the U.S.⁶ But it seems unlikely that these forces by themselves are strong enough to generate either the regional homogeneity or intensity of preferences that underpin Fig x.1B.

In our view, the key element is how individuals' social and political values are shaped—and reinforced—by context and agglomeration at finer geographic scales.⁷ Our thesis is that the same social forces that facilitate political participation contribute to ideological reinforcement at the local level. Cho and Rudolph (2008) identify four processes which are relevant in this regard:⁸

1. Residential self-selection, whereby citizens' characteristics jointly predict their residential and ideological preferences, and individuals

political variables, as well as information about the stability of such linkages.” (p. 568)

⁶ For example, southern conservative voters switching from Democratic to Republican (Schreckhise and Shields 2003; Bullock, Hoffman and Gaddie 2005; Valentino and Sears 2005), northeastern voters becoming increasingly liberal (Speel 1998), and the rise of the mountain west as a conservative voting bloc (Marchant-Shapiro and Patterson 1995).

⁷ E.g., Glaeser and Ward (2006: p. 131A): “These differences in beliefs within the United States drive home a central point about how politically relevant beliefs are formed. People in different states have been exposed to quite similar evidence through national media outlets, but they have reached radically different conclusions, and continue to hold these conclusions despite being aware that others disagree. This disagreement requires either different prior beliefs or some other deviation from Bayesian reasoning. One natural alternative model is that people base opinions mostly on the views of those around them. As such, local interactions are critical, and these provide plenty of possibility for wide geographic variation...”

⁸ See also Johnston et al (2004), who develop a slightly different taxonomy.

choose to live near to others who are socially and demographically similar to them, leading to spatial clustering of voting tendencies.⁹

2. Voter mobilization, in which partisan elites selectively target segments of the electorate on the basis of demographic attributes which are spatially clustered, especially in closely contested states or electoral districts. Spatially homogenous preferences facilitate targeting of extreme political statements to demographic groups to which they may have particular salience, catalyzing party alignment and turnout (cf. footnote 5).¹⁰
3. Social interaction, the set of mutually-responsive behaviors adopted by individuals in social networks. Social interactions may amplify ideological divisions because organized networks such as civic associations are a particularly effective mechanism for the exchange of political information (McClurg 2003), but such information tends to be systematical biased due to homophily—the propensity of individuals to interact with others who are similar to them (e.g., McPherson et al. 2001). Social interactions also promote ideological

⁹ Despite anecdotal evidence in favor of this hypothesis (e.g., Bishop 2008; Bishop and Cushing 2004), and indications particular kinds of neighborhood environments influence their residents' ideological leanings, irrespective of demographic composition (Williamson 2008), the political sources and consequences of self-selection have yet to be thoroughly investigated.

¹⁰ E.g., Mutz (2002: pp. 852): “Homogeneous environments are ideal for purposes of encouraging political mobilization. Like-minded people can encourage one another in their viewpoints, promote recognition of common problems, and spur one another on to collective action. Heterogeneity makes these same activities much harder. Participation and involvement are best encouraged by social environments that offer reinforcement and encouragement, not ones that raise the social costs of political engagement.” Also, Williamson (2008: pp. 20-21): “...the spatial sorting of residents by political ideology, once it reaches a sufficiently advanced stage, may help create what Lazarsfeld, Berelson, and Gaudet (1944) termed a ‘reinforcement effect’; not only might residents of a very conservative suburb be less likely to hear a liberal viewpoint from their neighbors but such areas will likely be targeted and contacted frequently by conservative political activists while being relatively ignored by liberal political activists, further reinforcing the relationship between spatial context and individual political outlook.” Community homogeneity facilitates a political campaign’s ability to mobilize voters by reducing the cost of what Lazarsfeld et al. (1944) refer to as activation (“not to form new opinions but raise old opinions over the thresholds of awareness and decision”, p. 74), and reinforcement (“to secure and stabilize and solidify [...] vote intention and finally to translate it into an actual vote”, p. 88).

homogeneity through the process of social learning, with views that are consonant with (dissonant from) those of the majority of network participants receiving positive (negative) reinforcement, leading to closer alignment of preferences within the network.¹¹ Finally, the fact that these effects transpire through direct interpersonal contact (and even non-verbal cues) suggests that the phenomenon of closer individual alignment with the local majority should only persist over a limited spatial domain.

4. So-called “casual observation”, the indirect, often involuntary, social interaction induced by the characteristics of an individual’s environment.¹² A key implication is that the physical attributes of citizens’ action spaces are likely to significantly influence their ideological preferences, independent of neighborhood demographics (cf. Williamson 2008). Non-political, day-to-day social interactions remain a key source of political information for Americans (Klofstad et al. 2006), with the workplace being the principal forum in which they are exposed to dissonant political views (Mutz and Mondak 2006). This suggests that the spatial domain of political influence is not limited to the neighborhood in which an individual resides, and may extend well beyond her commuting distance.

At a minimum, these processes imply that voters’ preferences should be influenced by those of the citizens around them. But, in view of the reinforcing character of the first three processes, we further claim that the likely outcome will be a phenomenon which we term “localized entrenchment”: in the absence of exogenous shocks, communities remain locked in a cycle of reinforcement of the values held by their ideological majorities, with corresponding suppression of the inward diffusion of countervailing viewpoints and ideas, leading to entrenchment of attitudes, beliefs and, ultimately, voting behavior. Our view of entrenchment as closer alignment

¹¹ E.g., Huckfeldt and Sprague (1995). For formal models of this process see Baldassarri and Bearman (2007), Dixit and Weibull (2007) and Glaeser and Sunstein (2007).

¹² E.g., Cho and Rudolph (2008: p. 277): “Casual observation exposes citizens to meaningful information through low-intensity neighbourhood cues such as the display of yard signs, bumper stickers, or simple observations and biases created by how neighbors dress and behave, what types of cars they drive, or how well their garden is groomed. Such low-intensity cues may influence behavior by subtly communicating information about the prevailing norms and sentiments within a community. In particular, they may provide signals about a local community’s political culture and ethic or the nature and distribution of political preferences within that community.”

between the individual vote and the local majority vote is consistent with evidence of increased party identification by voters (Miller and Hoffman 1999), ideological cleavages along geographic lines (Abramowitz and Saunders 2008; fn. 6) and regional concentration of the Democratic and Republican parties' representation in the U.S. Congress (Mellow and Trubowitz 2005).

The econometric consequences of local entrenchment are spatial correlation and endogeneity. These are anticipated by the economic literature on social interactions (Manski 1993, 2002; Glaeser et al. 2003), which suggests that the group of citizens in the zone of political influence around a particular individual will impact her vote decision in three ways. The first is endogenous effects, where the group's average voting behavior affects the individual's vote, which could potentially reflect the influences of any or all of the four processes above. The second is contextual or exogenous effects, where the group's average (exogenous) characteristics affect the individual's vote. This might reflect processes 3 and/or 4, as well as the spatial clustering of citizens with similar characteristics for reasons other than self-selection. The third is correlated effects, where the individual's error term is correlated with the error terms of members of the group because of similar characteristics not observed by the econometrician (process 4), sorting or selection of individuals based on who they are (process 1), or exposure to common shocks (process 2, and the polarized character of choices in the political arena more generally).

An additional consideration is that our ecological data on citizens' characteristics and votes at the level of the county (not the individual) forces us to reinterpret these effects in terms of areal units and their neighboring jurisdictions. We argue that even though this invariably introduces aggregation bias of unknown magnitude and sign, such a reinterpretation is still valid because of exogeneity in the boundaries of our areal units (see fn. 2) and the potentially long spatial reach of processes of casual observation. Moreover, the fact that we know the location of each observation means that the three effects above neatly correspond to the components of different spatial econometric models. Endogenous effects are captured by the coefficient on the spatially lagged county vote in a spatial autoregressive model; contextual effects are captured by the coefficients on spatial lags of the covariates; and correlated effects are indicated by the coefficient on the spatially lagged error term in a spatial error model. Quite likely, all three effects are simultaneously at work in our dataset, which presents a challenge for estimation.

Data

Our dependent variable is the vector of votes cast for Bush as a share of total votes at the county level. We estimate the size of total electorate as the sum of ballots cast for Bush, Kerry and independent candidate Ralph Nader, data for which were downloaded from the CBS News election 2004 website.¹³

For explanatory variables we selected a broad spectrum of demographic characteristics that are likely to have influenced individuals' voting decisions, which we organized into categories similar to those used in prior analyses of political polarization (DiMaggio et al. 1996; Evans 2003).


We employ four sets of demographic variables at the county level. These are the distribution of income, measured over four income categories; housing costs; the distribution of educational attainment, measured over five grades; racial and ethnic composition; and age, sex and national origin. These data were obtained the 2000 U.S. Census and Current Population Estimates data files, and are coded as percentages of either the total or the voting-age population within each county.

We also employ two categories of variables on economic characteristics of places: median household income, unemployment and the composition of employment; and local geographic characteristics such as population growth, the size of the local electorate, whether the county belonged to the core (urban) or outlying (suburban) region of a metropolitan statistical area (MSA), travel time to work and prevalence of commuting outside one's county of residence. Unemployment and wage data were compiled from U.S. Bureau of Labor Statistics Local Area Unemployment Statistics and U.S. Bureau of Economic Analysis Regional Economic Information System data files, respectively, while local geographic variables were collected from the 2000 Census and Current Population Estimates.

We include three additional sets of variables in an attempt to proxy for issues which exit polls indicate played an important role in the election: the war on terror and U.S. military intervention in Iraq, and "moral" or "family" values. Based on social interaction theory, we hypothesize that attitudes toward the former issue among the general population will be most strongly shaped by personal knowledge of—and face-to-face interaction with—individuals who have served or are currently serving in the

¹³ These data are of necessity approximate, not being adjusted for the results of recounts in Ohio and New Mexico. There were additional independent candidates on the ballot in each state, but the numbers of votes cast for them were small. Neither of these factors seems likely to significantly change our main results.

Fig. x.2. Box Plot of Descriptive Statistics of the Dataset



armed forces, and that the diffusion of attitudes will increase with geographic proximity to clusters of this sub-group (e.g., counties which host or immediately surround active military bases). Accordingly, we code for attitudes toward the war on terror using Census data on the fractions of veterans and active military personnel in counties' population.

Like no other set of issues, moral values are the bellwether of electoral polarization as the reflection a so-called "culture war" (Hunter 1992; Miller and Hoffman 1999; Evans and Nunn 1995). Pew Research Center (2004) notes that moral values are not precisely defined, but encompass conservative views on subjects as diverse the appropriate role of religious expression and proselytization in public life, marriage and divorce, wom-

en's fertility and right of access to abortion, and child-rearing in a traditional nuclear family setting. The multivariate and ambiguous character of values, coupled with the fact that they are not directly observable even at the individual level, means that our ability to precisely identify their effects using aggregate data is weak at best.

With regard to religion, a useful indicator is the distribution of adherents to different faiths—particularly evangelical Christians—among the population. We use data from Glenmary Research Center (2004) to construct the distribution of individuals with different religious affiliations by county, which we code as shares of the population.¹⁴ We proxy for attitudes to marriage using data on the fraction of population separated or divorced tabulated by the Census. Although some data are available on rates of teen pregnancy, out-of-wedlock births and abortion rates, they are not disaggregated to the county level, and were not used in our analyses.¹⁵ To proxy for attitudes to fertility and child-rearing practices, we employ Census data on average family size, the average number of children per household, and the percentage of households headed by a female with no husband present.¹⁶

Our final set of covariates captures an important aspect of the debate over values which played out in the 2004 election, namely, the polarization of attitudes toward homosexuals, especially the legalization of same-sex marriage, civil unions or domestic partnership benefits. There is a dearth of data on either the geographic distribution of either gay persons or general attitudes toward them. However, during 2004 eleven states enacted ballot initiatives to ban same-sex marriages.¹⁷ In an attempt to capture the effect of related attitudes on the vote, we treat these initiatives as an exogenous

¹⁴ Campbell and Monson (2008) note that this database suffers from a number of problems, principally non-response bias in survey questionnaires, omission of non-denominational churches—which account for an increasing share of religious participation, and an inability to track the number of residents of one county who attend church in another.

¹⁵ These data are available online from the Alan Guttmacher Institute. Preliminary regressions indicated that the state-level incidence of abortion and teen pregnancy were not significant predictors of the odds of voting Republican, in part because of their collinearity with state fixed effects.

¹⁶ Our use of the proportions of divorced persons and households headed by single females is admittedly crude. In particular, it is hard to know whether the statistical effect of these variables on electoral outcomes is driven by the voting behavior of people in these groups or by morally conservative voters' negative reactions to the former.

¹⁷ E.g., Charisse Jones, "Gay marriage on the ballot in 11 states", USA Today, Oct. 14, 2004. The states are: Arkansas, Georgia, Kentucky, Michigan, Mississippi, Montana, North Dakota, Ohio, Oklahoma, Oregon and Utah.

shock, and construct a dummy variable for gay marriage bans (GMB), coding the counties in these states as ones and the remaining counties in our sample as zeros. Following Campbell and Monson (2008), we include the interaction between this dummy and the percentage of Evangelical Christians in the population as a proxy for the potentially galvanizing influence of the ballot initiatives on turnout by evangelical voters for the Republican party.

We restrict our analysis to the contiguous counties of the lower 48 states, dropping counties in Alaska (for which disaggregate election returns are not tabulated) and Hawaii. Remaining counties for which one or more variables were missing were also dropped. Our final sample consists of 3106 observations (denoted below by N), which we geo-coded using the county centroids from the Census 2000 gazetteer files.

Fig. 2 presents the distributions of the variables as box plots to facilitate comparison. A few covariates, such as the percentages of Asian Americans or persons on active duty in the armed forces, have small magnitudes and exhibit very little variation. Conversely, other variables, such as the percentage of evangelicals and non-adherents, persons of solely Caucasian background, the average number of children per family, and the share of counties' electorates voting for Bush, all vary substantially across counties. In the working paper version of the article (Sue Wing and Walker 2005) we provide additional descriptive statistics that show that these variables exhibit significant interregional heterogeneity.

Our aim is to identify the association between the dependent variable and independent variables above, and then characterize the spatial variations in these relationships. To do this we turn to our spatial econometric toolkit.

Methods

Our analysis proceeds in three phases, the algebraic details of which can be found in an appendix to the preliminary version of this chapter (Sue Wing and Walker, 2005). Our first task is to characterize the degree of spatial polarization in the vote by examining the intensity and geographic scope of spatial clustering in county-level returns. Following Kim et al. (2003), we compute the vector of Local Moran's I statistics for Bush's share of the electorate in each county. Rather than use their method of employing county-to-county commuting flows as a spatial weighting variable, we construct a symmetric spatial weighting matrix (W) based on a simpler dis-

tance-based scheme.¹⁸ We use the results of this calculation to generate maps the regions of statistically significant spatial clustering of votes.

Our second task is to test the predictions of our local entrenchment hypothesis by investigating the effects of the explanatory variables in section 3 on the odds of voting Republican at the national level. We estimate the following linear-in-logarithms logistic model:

$$Y = X\beta + u, \tag{x.1}$$

in which the dependent variable is an $N \times 1$ vector of the log-odds ratios of each county voting Republican,

$$y_c = \log p_c - \log(1 - p_c), \quad c \in N \tag{x.2}$$

where the subscript c indicates counties and p_c is the probability of c 's Republican vote, estimated by Bush's share of the total votes cast, X is an $N \times k$ matrix of covariates given by the logarithms of the continuous independent variables in Figure 1, as well as the dummy variables and interaction terms described in the previous section. We interpret the coefficient β as the vector of elasticities of the odds of voting Republican with respect to the k covariates at the county level.

The overall explanatory power of the basic model in eq. (x.1) was good, but (not surprisingly) tests of the residuals indicated that the disturbance vector u exhibits significant spatial autocorrelation (without state dummies, the Moran's I standard deviate = 57.55, $p < 0.01$). The likely culprits are spatial sorting and selection of voters on the basis of demographics, as well the omission of contextual variables, common shocks—especially congressional and gubernatorial elections which were simultaneously being held in each state, and the endogenous effects of surrounding counties' votes.

To test how much of the spatial dependency in the errors could be explained by omitted contextual factors, we included fixed effects for each state. This dramatically improved the fit and mitigated the degree of spatial autocorrelation, but tests of the residuals still indicated problems (Moran's I standard deviate = 33.44, $p < 0.01$). Moreover, Lagrange multiplier tests of the residuals for an omitted spatially-lagged dependent variable (LM_p)

¹⁸ Consistent with our discussion of the prominent role of social interactions, we defined the neighborhood of each county as a radius of 200 km, which is approximately twice the distance traveled at the highest state-mandated speed limit (75 mph) for the maximum average commute time in Fig. x.2. The advantage of this scheme is that every row in W has at least one non-zero off-diagonal element, which allowed us to row-standardize the resulting matrix of distances without having to worry about divide-by-zero errors.

and spatially autocorrelated errors (LM_λ) led us to reject the linear model, with or without state dummies.¹⁹

Encouraged by these results, we turned to more sophisticated estimators capable of capturing the effects of interest: the spatial lag model

$$Y = \rho WY + X\beta + \varepsilon, \quad (x.3)$$

and the spatial error model, which augments eq. (x.1) with:

$$u = \lambda Wu + \varepsilon. \quad (x.4)$$

The variable ε is an $N \times 1$ vector of i.i.d. errors, ρ is a spatial lag correlation parameter, λ is a spatial error correlation parameter, and W is the $N \times N$ matrix of spatial weights described above. The models corresponding to eqs. (x.3) and (x.4) were estimated by maximum likelihood.

There is little a priori guidance as to which of these models is more appropriate. Spatial lag models are more common in the political science literature, and assume that the effects of a county's attributes on the odds of voting Republican are influenced by neighboring counties votes (i.e., endogenous effects), via the parameter ρ . On the other hand, the spatial error model assumes that spatial autocorrelation can be explained by aggregation bias, sorting and selection, or spatially-varying omitted variables (i.e., correlated effects), captured by the parameter λ .

Our estimates of the two parameters indicate a high degree of spatial dependency in the data ($\rho = 0.40$ and $\lambda = 0.96$, both $p < 0.01$), and the challenge we faced was to apportion this dependency among the three effects above. To this end, we employed Anselin et al's (1996) Lagrange multiplier tests of the spatial lag and spatial error specifications being mutually contaminated by each other, but both the test for error dependence in the possible presence of a missing lagged dependent variable (LM_λ^*), and the test for a missing lagged dependent variable in the possible presence of spatially correlated disturbances (LM_ρ^*), had power against each other. In both tests the null was rejected for the basic as well as the fixed-effects models,²⁰ but the test of robustness of the spatial error model against contamination by a spatially lagged dependent variable saw rejection of the null at the higher level of significance, apparently favoring the spatial error model. The log-likelihood and AIC statistics supported this conclusion, in-

¹⁹ For the basic model, $LM_\rho = 2695.41$ ($p < 0.01$) and $LM_\lambda = 1044.49$ ($p < 0.01$), while for the fixed effects model, $LM_\rho = 523.08$ ($p < 0.01$) and $LM_\lambda = 278.27$ ($p < 0.01$).

²⁰ For the basic model $LM_\rho^* = 172.99$ ($p < 0.01$) and $LM_\lambda^* = 1823.91$ ($p < 0.01$), while for the fixed-effects model $LM_\rho^* = 73.96$ ($p < 0.01$) and $LM_\lambda^* = 318.77$ ($p < 0.01$).

dicating that the spatial error model has the better fit to the data, however the extent of autocorrelation in the spatial error model remained a concern, especially since λ subsumed both endogenous and contextual effects.

In light of McMillen’s (2003) critique that spatial dependence in the error term might simply indicate misspecification—especially given our potential omission of spatially correlated right hand side variables, we decided to pursue a third alternative: the unconstrained spatial Durbin model:

$$Y = \rho WY + X\beta + WX\gamma + \varepsilon. \quad (\text{x.5})$$

This model nests both our lag and error specifications through the restrictions $\gamma = 0$ and $\gamma = -\rho\beta$, respectively (Anselin 2002). The latter “common factor hypothesis” (Burrige 1981) is decisively rejected by a likelihood ratio test (LR = 336.3, $p < 0.01$), suggesting that residual spatial autocorrelation in the error term of (x.3) arises as a consequence of omitted spatial lags of the covariates (i.e., contextual effects). Accordingly, we relied on the results of eq. (x.5) for our insights regarding the aggregate-level correlates of voting patterns in 2004, subject to the caveat that our results likely overstate endogenous and contextual effects while giving short shrift to correlated effects.

Our third task is to bring the results of the previous phases together to elucidate the implications of local entrenchment for the polarization of the U.S. electorate. In the preliminary phase of our analysis we re-estimated eq. (x.5) on contiguous subsamples of counties defined by the nine U.S. census divisions. The parameter estimates varied markedly among regions in magnitude, sign and significance, indicating that the national-level estimates mask substantial spatial heterogeneity. But given the hypothesized importance of local environmental influences for counties’ voting behavior, we sought a way to systematically characterize how the parameters of (x.5) vary over fine geographic scales.

Accordingly, to capture the full extent of spatial non-stationarity in our data we re-estimated our model as a geographically weighted regression (GWR), a nonparametric technique that generates intercept and slope parameters for every county by running a sequence of locally linear regressions on a sub-sample of data from nearby counties (Brunsdon et al. 1996; Fotheringham et al. 1997, 2002). The GWR model can be written:

$$Z_c Y = Z_c X \theta_c + v_c, \quad (\text{x.6})$$

in which Z_c is the matrix of local spatial weights centered around the c^{th} county,²¹ and θ_c is a spatially-varying $k \times 1$ vector of parameters associated with observation c . The latter allows us to map and analyze the spatial variation and clustering in our aggregate results. The fact that the GWR method estimates an intercept for each county drastically diminishes the ability of state dummies and spatial lags to capture unobserved contextual effects, and in any case, the computational exigencies of estimating many additional parameters overwhelmed our computing resources.²² We therefore used GWR to estimate only our basic linear model. Our final step was to test for polarization by examining whether the resulting vector of local odds elasticities θ_c exhibited significant spatial clustering along each of its dimensions (indicating entrenchment), and whether the clusters gave rise to distributions of effects similar to Fig. x.1B.

Results

The Spatial Clustering of Votes and Covariates

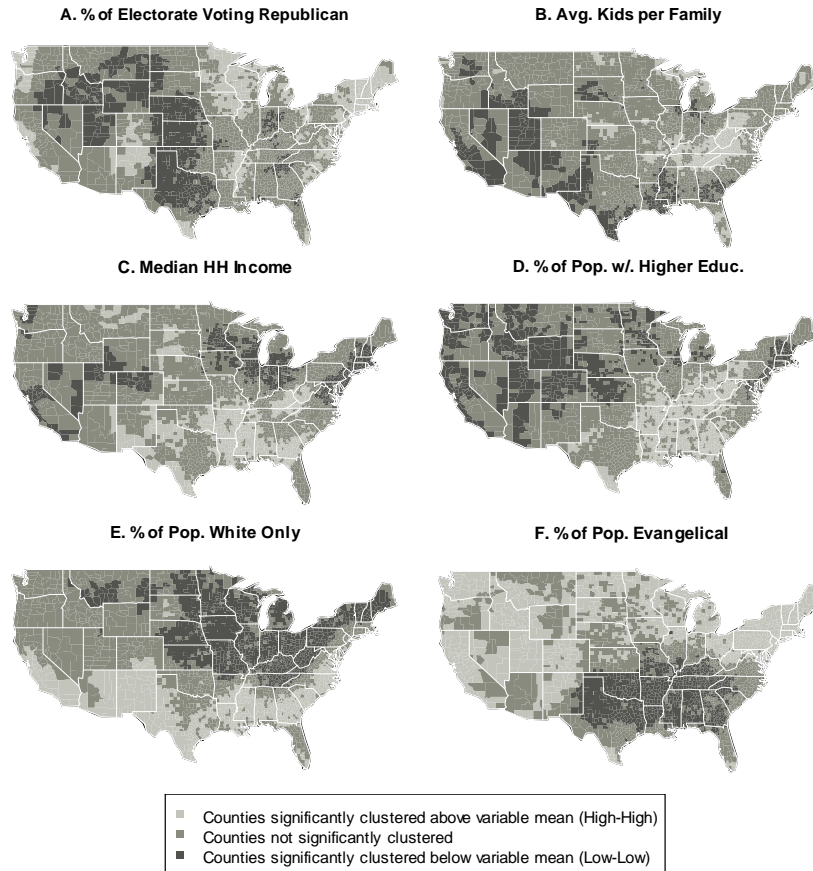
Applying Moran's test to county vote returns reveals significant global spatial autocorrelation in the election results (Moran's I standard deviate = 109.68, $p < 0.01$). We compute local Moran's I statistics for the Republican share of the vote and key independent variables, and plot the results as a series of significance maps, shown in Fig x.3. A two-tailed test of significance ($p < 0.05$) allowed us to classify each observation as one which exhibited significant spatial clustering of voting returns for Bush above (dark grey) or below (light grey) the sample means.

Significant clustering in the share of the electorate voting Republican, shown in panel A, is comparable to that found by Kim et al. (2003: p. 749, Fig. 2B), with clustering above the national average in large swaths of the Midwest, West Central and upper Mountain regions, as well as pockets in Appalachia, and clustering below the average in the Northeast and North Central regions, as well as in pockets along the Pacific coast and in south-

²¹ Specifically, $Z_c = \text{diag}[z_{1c}, \dots, z_{Nc}]$ is an $N \times N$ diagonal matrix of c 's distance-based weights expressed as a local kernel, $z_{jc} = \exp\left(-0.5(d_{jc}/h)^2\right)$, in which d_{jc} is the distance between c and other counties (j), and the spatial interaction radius is given by a fixed bandwidth parameter, h , that we estimate using a crossvalidation procedure.

²² All our analyses were performed using the spatial packages for the R statistical language (Bivand 2006; Bivand and Brunstad 2006).

Fig. x.3 Local Moran's I Significance Maps of Votes and Key Covariates.

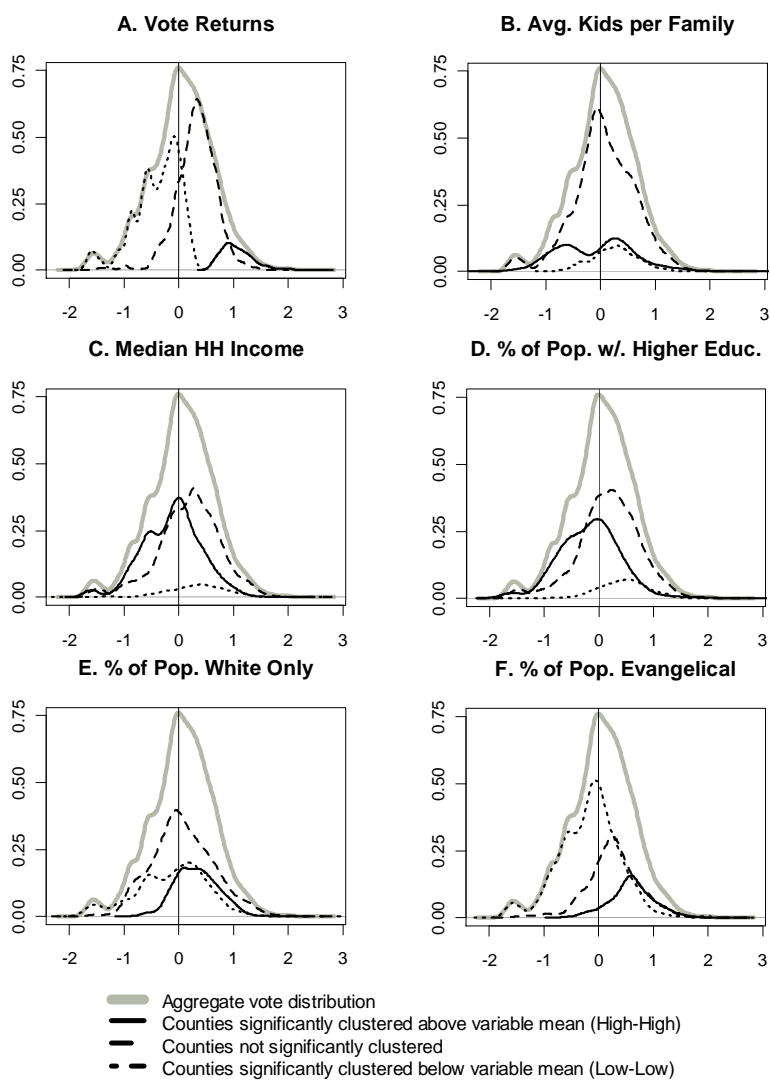


ern Texas and Florida. Such agglomeration is precisely what one would expect to be associated with an electorate that is polarized over space.

It is natural to inquire into the factors on which such clustering might depend. For example, the ideological sorting of populations could be based on any number of factors such as income, race or religion. Preliminary insight into this question can be gained by visually inspecting the patterns of spatial clustering of the explanatory variables.

Panels B-F show the results of computing Local Moran's I statistic for a subset of the covariates in Fig. x.2. The average number of children per family is clustered above the mean in pockets around the Great Lakes and

Fig. x.4 Log-Odds of Voting Republican by County Clusters



Notes: Kernel density estimates, weighted by counties' shares of the national electorate. The log-odds ratio of voting Republican is on the horizontal axis, with probability density on the vertical axis.

across the South and Southwest, and clusters below the mean in the Ohio river valley and pockets in Florida and the Midwest. Median household in-

come exhibits significant positive clustering in the northeast and upper Midwest, as well as in pockets in the mountain west and along the Pacific coast, with significant negative clustering in pockets throughout the Mississippi river valley, the west and the southeast of the country. The proportion of the population with post-secondary education is clustered above the mean in the northeast and in swaths across the upper Mountain West and coastal California, and clusters below the mean across the South. The share of Caucasians in the population is positively clustered in the North and East and negatively clustered across the South, particularly in the Southwest, the share of African Americans is positively clustered in the Southeast and negatively clustered in the North Central region, while the share of Evangelicals is positively clustered in the South and negatively clustered in the Northeast, Mountain and Western regions.

The relationships between the clustering of voting returns and the covariates are not obvious. To shed light on these associations, in Fig. x.4 we follow Ansolabehere et al. (2006) and plot the distributions of the log odds ratio at the county level (eq. x.1) for different subsets of the data based on the clustering of the variables in Fig. x.3. We examine how the propensity to vote Republican is distributed across counties which exhibit significant spatial clustering above or below the means of our subset of variables by constructing separate kernel density estimates, weighted according to the distribution of the electorate across the counties in each sample. A rightward (leftward) shift of the distributions thus indicates citizens' propensity to vote Republican (Democratic). To facilitate comparison we superimpose the plots of the densities on the reference vote distribution of the national electorate, shown in gray, whose unimodality at zero is often taken as prima facie evidence against polarization (e.g., Ansolabehere et al. 2006).

In Panel A, the positive clustering of the vote in the center of the country indicates strongly Republican preferences in the less populous counties there, with counties that do not exhibit significant clustering leaning slightly Republican, and the more populous counties that cluster negatively on the east and west coasts with exhibiting preferences that are moderately Democratic but with a negatively skewed distribution. A very different picture emerges if we segment the electorate according to the spatial clustering of fertility, however. In panel B, the large mass of non-clustered counties mirrors the shape of the aggregate vote distribution, while regions with families that have less than the average number of children tend to lean slightly Republican. Surprisingly, clusters of counties with greater-than-average numbers of kids per family have a bimodal distribution, with similar numbers of voters leaning Republican and Democratic. A similar pattern is exhibited by the influence of Caucasian populations (panel E),

with a higher propensity to vote Republican in clusters of less racially diverse counties, centrist preferences in counties that are not clustered, and a bimodal distribution in clusters of counties with smaller-than-average white populations.

Segmenting the electorate based on the clustering of household income and educational attainment (panels C and D) yields similar results. Counties belonging to low-income and low-education clusters seem to have fairly strong Republican leanings, non-clustered counties show a slight propensity in this direction, and the preferences of counties in high-income and high-education clusters are largely centrist, with slight Democratic leanings. In panel F, clustering of counties with larger-than-average proportions of evangelical adherents is strongly associated with voting Republican. This influence is less strong but still substantial for counties where clustering is not significant, and the large number of voters in clusters of counties with lower-than-average populations of evangelical Christians seem to have centrist or slightly Democratic preferences. Qualitatively similar patterns emerge on the basis of indicators of local geographic context (not shown). The vote distributions for urban and non-urban counties are markedly different, with rural, and especially suburban, contexts exhibiting a strong propensity to vote Republican (cf. McKee 2007, 2008; Williamson 2008).

Overall, the spatial agglomeration in both voting returns and selected covariates is broadly consistent with the predictions of our local entrenchment hypothesis. But even though agglomeration appears to be somewhat related to divisions in voting behavior, the precise association is not readily discernable. With the exception of panel F, Fig. x.4 clearly indicates the dominance of the non-clustered subsample's influence on the aggregate vote distribution. Thus, although entrenchment might well be occurring, there is no polarization of electoral returns across easily observable demographic segments of counties' populations, as anticipated by Fiorina and Abrams (2008).²³ But the key issue is whether, and if so, how, entrenchment might be affecting the propensity to vote, controlling for demographic characteristics. Our ability to draw inferences in this regard is limited by the univariate character of Fig. x.4's distributions, which fails to capture the simultaneous influences of multiple spatially clustered variables on the vote distribution. To address this issue we turn to our regression model, which rigorously establishes the statistical associations between the vote

²³ E.g., as in racially polarized voting, where whites and non-whites have divergent ideological preferences which push their vote distributions in opposite directions away from the mean, like B-II and B-III in Fig x.1B.

distribution and all of our covariates, controlling for the myriad patterns of spatial clustering in the data.

Aggregate-Level Regression Results

Our aggregate-level estimation results are summarized in Table x.1. For our preferred specification, a Lagrange multiplier test of the spatial Durbin model's errors did not indicate significant residual spatial autocorrelation (LM = 2.6, $p > 0.10$), which in our opinion vindicates our statistical approach. The spatial autoregressive parameter is positive and significant ($\rho = 0.3$, $p < 0.01$), and its magnitude suggests that the spatial clustering of voting behavior is accompanied by substantial endogenous effects, even after demographic and contextual influences are controlled for.²⁴ Moreover, with few exceptions the contextual influences associated with spatial lags of the covariates share the same sign as their direct counterparts. Thus, although the various characteristics of populations and places influence county-level voting returns in different directions, these effects are almost uniformly amplified by their geographic context. These, we argue, are powerful pieces of evidence in support of the local entrenchment thesis.

Income, Income Distribution and Housing: Counties with higher median household incomes are significantly more likely to vote Republican, consistent with Kim et al. (2003). Simultaneously, however, having high proportions of families with moderately high and particularly middle incomes (\$35-150,000) significantly lowers the odds of voting Republican, as does proximity to larger populations of the poorest families (< \$15,000). Larger shares of low-income households (\$15-35,000) have the reverse effect. The fact that we drop the proportion of families with the highest income (> \$150,000) from the regression to avoid collinearity then suggests that the propensity to vote Republican varies with income according to a U-shaped

²⁴ Our results suggest that the "social multiplier" associated with voting in the 2004 U.S. presidential election is around 1.4. This is substantially smaller than the values found by Glaeser et al. (2003) for the peer effects of college roommates, criminal behavior in cities, or the human capital spillovers in urban labor markets. This outcome is not surprising given that ballots are secret, and that even with early voting, individuals are only exposed to the influence of neighbors' self-announced behavior for at most three weeks. (Although more prolonged exposure might likely result from proximity to intensely partisan voters.) It therefore seems more plausible that ρ is picking up the influence of correlated effects associated with counties' common exposure to political campaigns, and the reflection of that stimulus in their residents' everyday social interactions.

Table x.1. Spatial Durbin Model Results

	Direct Effects (β)			Spatial Lag Effects (γ)		
% Fam. Inc. < \$15k	-0.010	(0.034)		-0.543	(0.233)	**
% Fam. Inc. \$15-35k	0.193	(0.058)	***	0.592	(0.371)	
% Fam. Inc. \$35-75k	-0.363	(0.063)	***	-1.300	(0.413)	***
% Fam. Inc. \$75-150k	-0.095	(0.033)	***	0.184	(0.241)	
% < 9th grade	0.064	(0.023)	***	0.198	(0.126)	
% Some High School	0.255	(0.035)	***	-0.211	(0.212)	
% High School Grad.	0.512	(0.059)	***	0.404	(0.352)	
% Some College	0.473	(0.052)	***	0.565	(0.324)	*
% Bachelor's Degree	0.091	(0.031)	***	0.332	(0.192)	*
% White Only	0.953	(0.037)	***	-0.322	(0.153)	**
% Latino	0.009	(0.010)		0.086	(0.045)	*
% Foreign-Born	-0.006	(0.009)		-0.159	(0.049)	***
% Fem. H. H. No Husb.	-0.646	(0.037)	***	-0.040	(0.217)	
Avg. Family Size	1.478	(0.380)	***	-0.989	(2.326)	
Avg. Kids per Family	0.426	(0.122)	***	1.364	(0.781)	*
% Veterans	-0.068	(0.039)	*	0.217	(0.214)	
% Evangelical	0.003	(0.005)		0.080	(0.029)	***
% Mainline Protestant	0.006	(0.006)		-0.153	(0.041)	***
% Catholic	-0.004	(0.002)	*	0.016	(0.022)	
% Unclaimed	0.027	(0.012)	**	0.232	(0.063)	***
Median HH Inc.	0.383	(0.092)	***	0.314	(0.498)	
% Unemployment	-0.079	(0.018)	***	0.012	(0.085)	
% Workforce Agric.	0.026	(0.008)	***	-0.106	(0.052)	**
% Workforce Mfg.	0.038	(0.010)	***	0.071	(0.053)	
% of Nat'l Electorate	-0.005	(0.009)		-0.227	(0.056)	***
% Work Outside Cnty.	-0.025	(0.012)	**	-0.293	(0.081)	***
Avg. Travel Time	-0.276	(0.040)	***	-0.474	(0.258)	*
Pop. Change 2000-03	0.646	(0.166)	***	1.314	(0.907)	
Suburban County	-0.070	(0.017)	***	-0.236	(0.183)	
Rural County	-0.052	(0.015)	***	-0.099	(0.141)	
GMB	-0.063	(0.182)		-1.735	(0.396)	***
GMB x % Evangelical	0.473	(0.079)	***	0.797	(0.438)	*
ρ	0.297	(0.063)	***			
LR test	18.759		***			
Log likelihood	-43.62					
ML residual variance	0.060					
AIC	405.23					
LM resid. autocorrelation	2.622					

Notes: the dependent variable is the log odds of voting Republican (eq. x.2); asymptotic standard errors in parentheses; * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

distribution. Once the spatial dependence in the data is accounted for, housing values do not have a significant effect.

Education: Surprisingly, larger proportions of persons at all levels of educational attainment are significantly associated with higher odds of voting Republican. As before, we interpret this result in light of the fact that we drop the proportion of highest-attaining persons (those with postgraduate training) to avoid collinearity. The suggestion is that the propensity to vote Republican varies with education according to an inverted U-shaped distribution, with larger proportions of very low and very high attaining individuals substantially reducing the propensity to vote Republican. Interestingly, contextual influences amplify these forces in both directions, with the odds of voting for Bush reinforced by clustering of individuals with some post-secondary education and attenuated by clustering of college graduates.

Race and Ethnicity: The influences of Asian- and African-American populations were tiny and not statistically significant, which led us to drop these variables from the model. Proximity to higher proportions of persons of Latin American origin is associated with increased odds of voting Republican. The proportion of persons reporting purely Caucasian origins has a similar influence, but its magnitude is an order of magnitude larger. Interestingly, the coefficient on the spatial lag of this sub-population has the opposite sign. The likely reason is that core metro counties, which on average had larger minority populations, tended to vote for Kerry in significantly higher numbers relative to their surrounding suburban counties, which had a significantly higher proportion of white residents.²⁵

Age, Sex and National Origin: The percentages of the elderly and voting age females in the population did not appear to significantly influence the 2004 vote. However, proximity to clusters of persons born outside the U.S. had a small negative impact on the odds of voting Republican.

Moral/Family Values: The effect of the share of divorced or separated individuals in the population was not significant, in line with findings of the broad acceptance of this social phenomenon (e.g., Thornton and Young-DeMarco 2001). The proportion of households headed by a single female was strongly associated with lower odds of voting for Bush, while the corresponding spatial lagged variable is not significant. This result appears less consistent with the “values-voter” hypothesis than with interest group behavior by poor single mothers, who, as the principal beneficiaries of the American welfare system (Gensler 1996) were directly impacted by

²⁵ Kruskal-Wallis rank-sum tests indicated significant differences between suburban and core metro counties’ distributions of the vote and the proportion of the population self-identifying as white only.

Republican-initiated conservative social policies such as accelerated welfare-to-work transitions (see, e.g., Allard 2007). The direct effects of higher average fertility and, especially, larger family sizes were significant and positive. The odds elasticities for family size and proximity to large populations of children are both particularly large, and the similarity in their magnitudes is not surprising given these variables' high correlation.

Iraq/War on Terror: Our proxies for the spatial distribution of attitudes to U.S. foreign policy perform poorly. The proportion of active duty personnel in the population is not significant, while the effect of the proportion of veterans is positive and significant, but small. Thus, bearing in mind the significant limitations of our data, we find little evidence that in 2004 security concerns trumped values in influencing voter behavior (cf. Hillygus and Shields 2005).

Religious Affiliation: We do not find that the shares of adherents to various religious denominations and substantially increase the odds of voting Republican. The share of Catholics and the spatially lagged percentage of mainline Protestants in the population both have small, negative and significant effects, while the spatial lag of the proportion of Evangelicals is significant, positive and not as large, and the fraction of persons with no religious affiliation is significantly positive in the spatial lag and spatial error models. These results are consistent with previous evidence of low turnout among conservative protestants (Manza and Brooks 1997; Woodberry and Smith 1998), as well as the conclusion that religious issues on their own made little difference to the outcome of the election (Hillygus and Shields 2005; Campbell and Monson 2008).

Employment: High unemployment rates are associated with significant reductions in the odds of voting Republican (consistent with Kim et al. 2003), while the fractions of the workforce in agriculture and manufacturing both have the opposite effect. As well, the coefficient on the spatial lag of agricultural employment is negative, which appears to reflect the fact that suburban counties have significantly lower agricultural employment than their surrounding rural counties without significant differences in their vote distributions.²⁶

Local Geographic Factors: The rate of population increase has a large, positive and significant effect of the odds of voting for Bush, whereas travel time to work, the fraction of population commuting outside the county, being a suburban or rural county, or neighboring a more populous county all have smaller negative impacts. The association between a higher pro-

²⁶ Kruskal-Wallis tests indicated significant differences between suburban and rural counties' distributions of the proportion of jobs in agriculture, but not their voting patterns.


propensity to vote Republican voting behavior and rapid growth of the population reflects the effect of migration on spatial sorting along ideological lines, and is consistent with Gimpel and Schuknecht's (2001) finding that interstate in-migration has aided Republicans whereas out-migration has aided democrats. The influence of commute time may simply reflect the fact that voters in the highly urbanized and strongly Democratic areas of the northeast and the west coast live closer to where they work and thus enjoy shorter commutes. But it also suggests a higher probability of counties' residents being exposed to social contexts that are potentially different from their own neighborhoods, with consequent inward diffusion of a diversity of political ideas and beliefs (cf. our discussion of process 4 in the previous section), which would tend to mitigate local entrenchment. The sign of the effects of suburban and rural dummies is at odds with previous findings (McKee 2007, 2008; Williamson 2008), indicating that once contextual and endogenous factors are controlled for, the environment in these types of locales does not appear to strongly increase—or, for that matter, reduce—the propensity to vote Republican.

Gay Marriage Bans: Surprisingly, we find that at the national level, the GMB dummy has a consistently negative, but indirect, impact on the odds of voting Republican.²⁷ Nevertheless, the interaction between the GMB dummy and the evangelical population has strongly positive direct and indirect effects. These two results, taken together with our prior finding that the concentration of evangelical adherents appears to have little effect on its own, are consistent with Campbell and Monson's (2008) conclusion that the state ballot initiatives to ban same-sex marriage served to increase turnout among moral conservatives. If one compares the magnitudes of the two sets of coefficients, the most interesting feature is the suggestion that at the aggregate level the ballot initiatives may have provoked a backlash which was big enough to compensate for its positive (direct and indirect) influences on evangelical turnout.

These results confirm both the existence and importance of effects that are consistent with local entrenchment. Furthermore, the fact that the prevalence of minority populations and our proxies for voters' economic status and orientation on moral values end up having the strongest influence suggests that these are key dimensions along which there was significant segmentation of the American electorate in the 2004 election. Nevertheless, our conclusions are tempered by the fact that the elasticities in Table x.1 are national averages that do not indicate how these divisions might

²⁷ Given our coding of GMB as a state dummy variable, the significance of the indirect spatial lag (as opposed to the direct) coefficient is to be expected, as it is by definition a wide-area effect.

Fig. x.5 Geographically Weighted Regression Results.



Obs.	3106	σ^2 (ML)	0.049
Bandwidth	281.27	AICc	481.461
Effective no. of parameters	604.83	AIC	-120.04
Effective degrees of freedom	2501.17	Resid. sum of squares	151.48

Notes: the dependent variable is the log odds of voting Republican (eq. x.2).

have played out spatially. To shed light on this question we turn to the results of our geographically weighted regression analysis.

Geographically Weighted Regression Results

Our GWR results are summarized in Fig. x.5. Estimation of the model was plagued by multicollinearity between the average number of kids and the

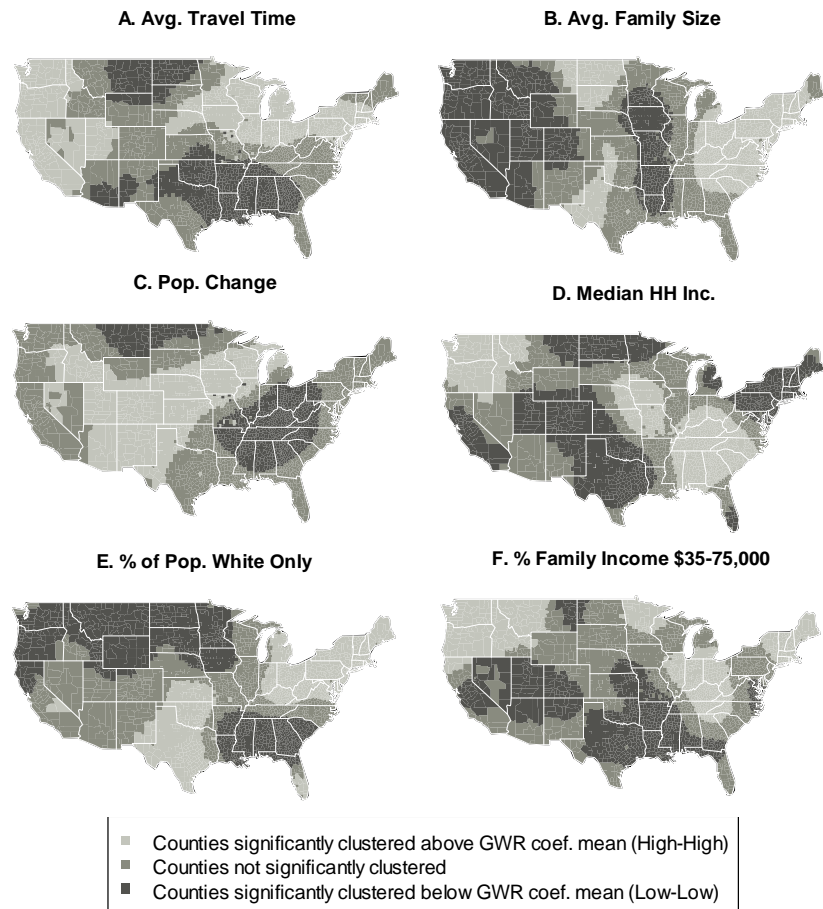
average family size by county, which led us to drop the former variable from our specification. The residual variance and the AIC statistic indicate an improvement in the fit over the spatial Durbin model above, and the distribution of local R^2 values, which ranges from 0.47-0.97 with a mean of 0.73, suggests that the GWR model's overall ability to account for the local spatial variation in the dependent variable is quite good. Our optimal bandwidth estimated through crossvalidation is substantially larger than the one used to compute the weights in the spatial Durbin model, with the result that the global values of the GWR parameters ($\bar{\theta}$) differ slightly from the odds elasticities of the previous section.²⁸ Even so, the overall results are basically the same. With the exception of family size, all of the parameters are less than one in absolute magnitude, which suggests that on average most covariates do not exert overwhelmingly large effects on the electoral returns.

Looking beyond averages, we see that the most important features of the county-level results is the variance of the parameter distributions, and the considerable spatial heterogeneity in the magnitude and sign of the influences of our covariates on the election returns. Our estimates can be classified into three types: (i) variables whose overall impact is large enough to be definitively signed, especially population change, household income and family size, (ii) those whose global impact is negligible but whose spatial variation is large, such as the proportions of veterans and very poor households, and GMB and its interaction term, and (iii) those whose average impact and its cross-county variance are both small—the category into which the majority of variables falls.

None of the covariates whose average estimates are significantly positive or negative exhibits tight clustering of their effects on individual counties' propensity to vote. Furthermore, once we control for the influences of other attributes, GWR does not produce not a simple relationship between the cross-county variation in a particular factor (i.e., column of X) and the variation exhibited by that factor's odds elasticities (i.e., the corresponding column of θ_c in eq. x.6). These results are consistent with our previous findings, and reinforce the point that the electorate is not polarized along easily observable demographic lines. While characteristics such as median household income, average commute time, and the proportions of Cauca-

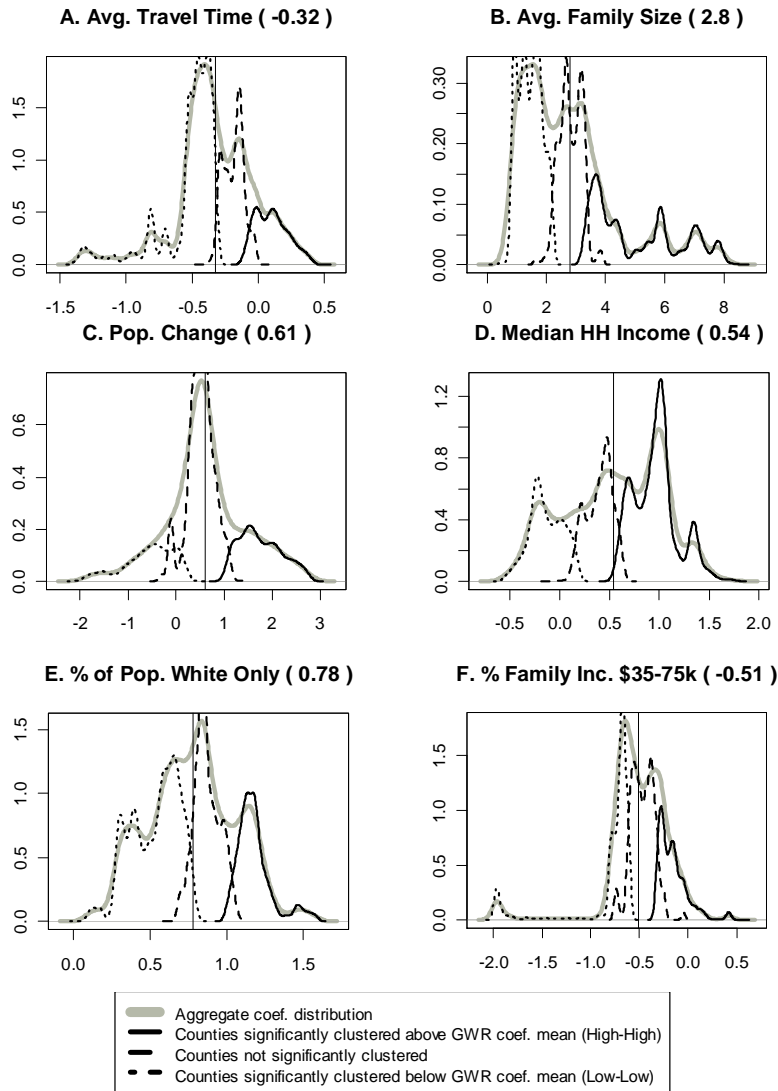
²⁸ Globally, the signs and relative magnitudes of the estimates are similar. However, the magnitudes of almost half of the estimates shrink while the rest increase. The median values of the parameter distributions are in closer agreement with the signs of our spatial Durbin estimates, though slight differences in their magnitudes persist.

Fig. x.6 Local Moran's I Significance Maps of GWR Odds Elasticities.



sians, persons in middle income families, and households headed by single females come closest in this regard, with average effects that are large and either significantly positive or negative in at least three-quarters of our sample, even they exhibit substantial non-stationarity. Category-(i) covariates (above) which seem most likely to shift the vote in a particular direction exert the opposite impact in a substantial minority of counties, and category-(ii) covariates exhibit effects of similar intensity but opposite sign in equally large numbers of counties. Due to the local character of the regression, these attributes are the ones for which the dependence of coun-

Fig. x.7 GWR Odds Elasticities of Voting Republican by County Clusters



Notes: Kernel density estimates, weighted by counties' shares of the national electorate. Odds elasticities of voting Republican are on the horizontal axis, with probability density on the vertical axis. Electorate-weighted mean GWR elasticities are in parentheses.

ties' odds of voting Republican on similar propensities among their neighbors will be the most obvious.

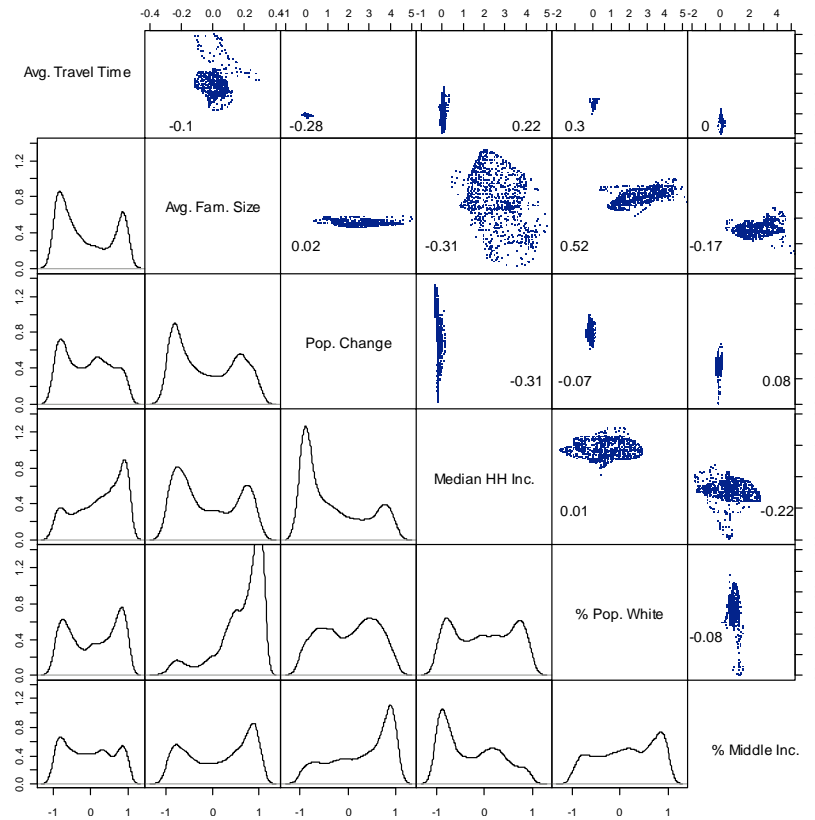
To test this proposition we examine the patterns of agglomeration in the odds elasticities with large spatial variation. Our strategy is to once again compute local Moran's I statistics, this time for the vector of parameters associated with each covariate (θ_c), and, as in Fig. x.3, display the results as a series of significance maps.²⁹ The results are shown in Fig x.6, which illustrates the striking spatial trends of agglomeration in our estimates. The effect of each of the six characteristics on the propensity of a particular county to vote Republican or Democratic depends on how that attribute influences the ideological leaning of its neighbors, which exert an amplifying effect within the clusters. While our GWR model is not able to identify either the precise channels through which these feedbacks operate, or how their signs and intensities vary, it is nonetheless clear that the clustering in Fig x.6 reflects the influences of the spatially lagged variables seen in the previous section.

It is a challenge to even describe—not to mention intuitively accounting for the origins of—these patterns (Sue Wing and Walker 2005). Accordingly, we move directly to comparing the distributions of the elasticity values for counties within and outside the various clusters in Fig. x.6. The latter results are summarized in Fig. x.7 as kernel density estimates of the distributions of odds elasticities that are weighted by counties' shares of the national electorate and segmented according to their propensity to cluster significantly above or below their respective means. The distributions bear a striking resemblance to Fig. x.1B, especially the influence of family size (panel B), for which 54% of the electorate tended to cluster below the mean and 17% clustered above, and median household income (panel D), for which the corresponding proportions are 24% and 48%. We note that these variables have multimodal aggregate distributions with fairly large variances, especially the effect of family size, with its long upper tail. The distributions of spatially-clustered elasticities for commute time and the proportion of Caucasians (panels A and E) are also bimodal, with 18% (21%) of the probability mass in the former (latter) case clustered above the mean and 55% (46%) clustered below. By contrast, the distributions for population change and the proportion of middle income families (panels C and F) exhibit a greater degree of central tendency, with the majority of voters in residing counties that are not significantly clustered in one way or the other.

These results support our hypothesis that local entrenchment is associated with polarization of the electorate. We uncover similar evidence in

²⁹ In conducting these analyses we employ our original 200 km bandwidth kernel.

Fig. x.8 GWR Odds Elasticities: Global and Local Correlations



Notes: upper panels show scatter plots of GWR parameters, with global correlation coefficients in parentheses; lower panels show local correlations among parameters (see Fotheringham et al. 2002) as unweighted kernel density estimates.

the distributions of spatial clustering of other, less wide-ranging odds elasticities, but we leave the elaboration of these details to future work. Our findings underscore the subtle point that polarization is an inherently multidimensional phenomenon. Stepping back, it is clear that the overall picture is not as simple as the one articulated in popular discourse—we show that despite the fact that votes for different parties cluster regionally, they are not concentrated in disjoint subsets of the electorate. Rather, along dimensions such as race, income and indicia of family values the U.S. appears to be divided into disjoint swaths of geographically contiguous coun-

ties, with the same attribute amplifying the propensity to vote Republican in one set of regions while simultaneously exerting the opposite influence in another.

Given this, it is easy to see why a simple pattern of red and blue states does not arise: it is not the case that the same counties cluster above or below the mean along all, or even most, dimensions of the space of characteristics. A particular county's odds elasticities might be in a "high-high" cluster in some dimensions, while in other dimensions they might be in a "low-low" cluster—or not belong to a cluster at all. The county's ultimate propensity to vote one way or another is the scalar product of these varying local odds elasticities and its actual characteristics (i.e., $X_c\theta_c$), which generally differs from the influence that an individual characteristic (given by the relevant element of θ_c) might have. The implication is that in the U.S. context, electoral polarization should be thought of as a series of cross-cutting divisions that manifest themselves not between population sub-groups but within individual sub-groups over space.

We close by qualifying this conclusion with an important caveat. Wheeler and Tiefelsdorf (2005) find that the GWR algorithm can potentially induce spatial bias in the local parameter estimate that is sufficiently large to invalidate their meaningful interpretation.³⁰ The considerable spatial dependence in our GWR results might then lead one to question the extent to which we are able to trust the local values of the odds elasticities, and in particular their patterns of clustering which are the basis of our inferences about polarization. To address this issue, we follow Wheeler and Tiefelsdorf's (2005) recommendations and examine the extent of local and global correlation among the six sets of odds elasticities in Figs. 6 and 7. The results of our robustness tests are summarized in Fig. 8 as scatterplots of the coefficient estimates (upper panels) and distributions of their local correlations (lower panels). The odds elasticities are not globally correlated, but there are indications of correlation at the local level, particularly between the effects of average family size and the proportion of Caucasians in the population, and population change and median household income. Of course, these tests are not conclusive, but in the absence of strong prima facie evidence of bias we are confident that our results stand. In any event,

³⁰ These authors find that a simple model with two independent variables, the coefficients associated with each covariate may exhibit collinearity even if the underlying exogenous variables in the data generating process are uncorrelated, and a high degree of spatial correlation between two covariates increases the potential for the two sets of coefficients to exhibit interdependent, spatially opposing patterns of effects. In both cases the upshot is spurious spatial trends in the GWR estimates.

there is no easy way to remedy the effects of spatial multicollinearity within the analytic framework developed here. Quite likely, efforts in this regard will require an entirely separate program of analysis and testing (e.g., along the lines of Wheeler 2007). Therefore, the best we can do given the constraints of available space is to flag this issue as a priority for future research.

Conclusion

This paper sheds new light on the fundamental role of geography in determining both the outcome of the 2004 U.S. presidential election, and the polarized character of the American electorate more generally. Our guiding hypothesis is that polarization of the U.S. electorate has occurred over space and is attributable to a process of local entrenchment, whereby a variety of social forces amplify county populations' propensity to vote Republican or Democratic.

Analyzing data from a large sample of counties in the lower 48 states, we find the influences on voting behavior associated with contextual and endogenous factors to be broadly consistent with the predictions of our thesis, with considerable spatial clustering in both electoral returns and the characteristics of populations and places clearly pointing to the amplification of the effects on the vote of the attributes of populations and places. A much richer picture emerges when we explicitly account for the geographic variations in these estimates. At the global level, our GWR odds-elasticities basically agree with the results of our aggregate-level analyses, while at the local level exhibiting substantial heterogeneity in both magnitude and sign, and strong spatial trends. The latter imply that in the U.S. context electoral polarization is not synonymous with segmentation across population sub-groups following observable demographic characteristics. Rather, polarization appears to be a phenomenon which occurs within individual sub-groups across space. Furthermore, geography matters in ways that are crucial, but not easily explained using aggregate data analysis. It is not simply the case that the spatial distribution of population characteristics drives the interregional differences of voting patterns observed in the 2004 presidential elections. Rather, the latter emerge from the reinforcing influence of the local social context on the effects of the racial composition, income, and, less tangibly, social values of counties' populations.

Our hope is that this study will motivate geographers and political scientists alike to employ disaggregate individual data to account for the detailed social mechanisms that give rise to these broad spatial trends.

References

- Abramowitz A, Saunders K (2008) Is Polarization a Myth? *J Polit* 70:542-555
- Allard, SW (2007). The Changing Face of Welfare during the Bush Administration. *Publius* 37:304-332
- Ansolabehere S, Rodden J, Snyder JM (2006) Purple America. *J Econ Perspect* 20:97-118
- Anselin L (2002) Under the hood: Issues in the specification and interpretation of spatial regression models. *Agr Eco* 27:247-267
- Anselin L, Bera AK, Florax RJGM, Yoon MJ (1996) Simple diagnostic tests for spatial dependence. *Reg Sci Urban Econ* 26:77-104
- Baldassarri D, Bearman P (2007) Dynamics of Political Polarization. *Am Sociol Rev* 72:784-811
- Baldassarri D, Gelman A (2008) Partisans without constraint: Political polarization and trends in American public opinion. *Am J Sociol*, in press
- Bartels LM (2000) Partisanship and Voting Behavior, 1952-1996. *Am J Polit Sci* 44:35-50.
- Bivand R (2006) Implementing Spatial Data Analysis Software Tools in R. *Geogr Anal* 38:23-40
- Bivand R, Brunstad R (2006). Regional growth in Western Europe: detecting spatial misspecification using the R environment. *Pap Reg Sci* 85:277-297.
- Bishop W (2008) *The Big Sort: Why the Clustering of Like-Minded America Is Tearing Us Apart*. Houghton-Mifflin, New York
- Bishop W, Cushing RG (2004) Response to Philip A. Klinkner's "Red and Blue Scare: The Continuing Diversity of the American Electoral Landscape". *The Forum* 2(2):Art 8
- Brewer MD (2005) The Rise of Partisanship and the Expansion of Partisan Conflict Within the American Electorate. *Polit Res Quart* 58:219-229
- Brunsdon C, Fotheringham AS, Charlton ME (1996). Geographically weighted regression: a method for exploring spatial non-stationarity. *Geogr Anal* 28:281-298
- Bullock CS, Hoffman DR and Gaddie RK (2005) The Consolidation of the White Southern Congressional Vote. *Polit Res Quart* 58:231-243
- Burridge P (1981) Testing for a common factor in a spatial autoregression model. *Environ Plan A* 13: 795-800
- Campbell DE, Monson JQ (2008) The Religion Card: Gay Marriage and the 2004 Presidential Election. *Pub Opin Quart* 72:399-419
- Cho WKT, Rudolph TJ (2008) Emanating Political Participation: Untangling the Spatial Structure behind Participation. *Brit J Polit Sci* 38:273-289
- DiMaggio P, Evans J, Bryson B (1996) Have Americans' social attitudes become more polarized? *Am J Sociol* 102:690-755
- Dixit AK, Weibull JW (2007) Political Polarization. *Proc Nat Acad Sci* 104:7351-7356
- Evans JH (2003) Have Americans' Attitudes Become More Polarized?—An Update. *Soc Sci Quart* 84:71-90

- Evans JH, Nunn LM (2005) The Deeper "Culture Wars" Questions. *The Forum* 3(2):Art 3
- Fiorina MP, Abrams SJ, Pope JC (2006) *Culture War? The Myth of a Polarized America*, 2nd edn Longman, New York
- Fleisher, R, Bond J (2004) The Shrinking Middle in the U.S. Congress, *Brit J Polit Sci* 34:429-451.
- Fotheringham AS, Brunsdon C, Charlton M (2002) *Geographically Weighted Regression: The Analysis of Spatially Varying Relationships*. Wiley, Chichester
- Fotheringham AS, Charlton ME Brunsdon C (1997). Measuring spatial variations in relationships with geographically weighted regression. In: Fischer MM, Getis A (eds) *Recent developments in Spatial Analysis*. Springer-Verlag, London, pp 60-82
- Frank T (2004). *What's the Matter with Kansas? How Conservatives Won the Heart of America*. Metropolitan Books, New York
- Frey WH (2000) *Regional Shifts in America's Voting-Aged Population: What Do They Mean for National Politics?* Population Studies Center Report No. 00-459, Institute of Social Research, University of Michigan
- Gensler, H (ed) (1996). *The American Welfare System: Origins, Structure, and Effects*. Praeger, Westport
- Gimpel JG, Schuknecht JE (2001) Interstate Migration and Electoral Politics. *J Polit* 63:207-231
- Glaeser EL, Ward BA (2006) Myths and Realities of American Political Geography. *J Econ Perspect* 20:119-144
- Glaeser EL, Sacerdote BI, Scheinkman JA (2003) The Social Multiplier. *J Eur Econ Assoc* 1:345-353
- Glaeser EL, Sunstein CR (2007) *Extremism and Social Learning*. NBER Working Paper No W13687
- Glenmary Research Center (2004) *Religious Congregations and Membership in the United States: 2000*. Nashville
- Goodman L (1953) Ecological regression and the behavior of individuals. *Am Sociol Rev* 18:663-64
- Hanushek EA, Jackson JE, Kain JF (1974) Model Specification, Use of Aggregate Data, and the Ecological Correlation Fallacy. *Polit Method* 1:87-107
- Hetherington MJ (2001) Resurgent Mass Partisanship: The Role of Elite Polarization, *Am Polit Sci Rev* 95:619-631
- Hillygus DS, Shields TG (2005) Moral Issues and Voter Decision Making in the 2004 Presidential Election. *PS: Polit Sci Pol* 38: 201-210.
- Huckfeldt R; Sprague J (1995) *Citizens, Politics, and Social Communication: Information and Influence in an Election Campaign*. Cambridge University Press, New York
- Hunter JD (1992) *Culture Wars: The Struggle to Define America*. Basic Books, New York
- Johnston R, Jones K, Sarker R, Propper C, Burgess C, Bolster A (2004) Party support and the neighbourhood effect: spatial polarisation of the British electorate, 1991-2001. *Political Geography* 23:367-402

- Kim J, Elliott E, Wang DM (2003) A spatial analysis of county-level outcomes in U.S. Presidential elections: 1988-2000. *Elect Stud* 22:741-761
- Klinkner PA (2004) Red and Blue Scare: The Continuing Diversity of the American Electoral Landscape. *The Forum* 2(2):Art 2
- Klinkner PA, Hapanowicz A (2005) Red and Blue Déjà Vu: Measuring Political Polarization in the 2004 Election. *The Forum* 3(2):Art 2
- Klofstad C, McClurg SD, Rolfe M (2006) Family Members, Friends, and Neighbors: Differences in Personal and Political Networks. Paper presented at the 2006 Annu Mtg Midwest Polit Sci Assn
- Lazarsfeld P, Berelson B, Gaudet H (1944) *The People's Choice: How the Voter Makes Up His Mind in a Political Campaign*. Columbia Univ Press, New York
- Manski CF (1993) Identification of Endogenous Social Effects: The Reflection Problem. *Rev Econ Stud* 60: 531-542
- Manski CF (2002) Economic Analysis of Social Interactions. *J Econ Perspect* 14:115-136
- Manza J, Brooks C (1997) The religious factor in US presidential elections, 1960-1992. *Am J Sociol* 103:38-81
- Marchant-Shapiro T, Patterson KD (1995) Partisan Change in the Mountain West. *Polit Behav* 17:359-378
- McClurg SD (2003) Social Networks and Political Participation: The Role of Social Interaction in Explaining Political Participation. *Polit Res Quart* 56: 449-464
- McKee SC (2007) Rural Voters in Presidential Elections, 1992-2004. *The Forum* 5(2): Art. 2
- McKee SC (2008) Rural Voters and the Polarization of American Presidential Elections. *PS* 41: 101-108
- McMillen DP (2003) Spatial autocorrelation or model misspecification? *Intl Reg Sci Rev* 26:208-217
- McPherson M, Smith-Lovin L, Cook JM (2001). Birds of a Feather: Homophily in Social Networks." *Annu Rev of Sociol* 27:415-444
- Mellow N, Trubowitz P (2005) Red versus blue: American electoral geography and congressional bipartisanship: 1898-2002. *Polit Geogr* 24:659-677
- Miller AS, Hoffmann JP (1999) The Growing Divisiveness: Culture War or a War of Words? *Soc Forces* 78:721-745
- Mutz DC (2002) The Consequences of Cross-Cutting Networks for Political Participation. *Am J Polit Sci* 46:838-55
- Mutz DC, Mondak JJ (2006) The Workplace as a Context for Cross-Cutting Political Discourse. *J Polit* 68:140-155
- O'Loughlin J, Flint D, Anselin L (1994) The Geography of the Nazi Vote: Context, Confession and Class in the Reichstag Election of 1930. *Ann Assoc Am Geogr* 84:351-380
- Openshaw S (1984) *The Modifiable Areal Unit Problem*. Geo Books, Norwich
- Pew Research Center (2004) Moral Values: How Important? Voters Liked Campaign 2004, But Too Much "Mud-Slinging". Pew Research Center for the People and the Press Report, 11 November

- Poole KT, Rosenthal H (2001) D-NOMINATE After 10 Years: A Comparative Update to Congress: A Political-Economic history of Roll-Call Voting. *Legis Stud Quart* 26:5-29
- Schreckhise WD, Shields TG (2003) Ideological Realignment in the Contemporary U.S. Electorate Revisited. *Soc Sci Quart* 84:596-612
- Speel RW (1998) *Changing Patterns of Voting in the Northern United States: Electoral Realignment, 1952-1996*. Penn State Univ Press, Philadelphia.
- Stonecash JM, Brewer MD, Mariani MD (2002) *Diverging Parties: Social Change, Realignment, and Party Polarization*. Westview Press, Boulder
- Sue Wing I, Walker JL (2005) *The 2004 Presidential Election From A Spatial Perspective*. Mimeo, Boston University
- Thornton A, Young-DeMarco L (2001) Four Decades of Trends in Attitudes Toward Family Issues in the United States: The 1960s Through the 1990s. *J Marriage Fam* 63:1009-1037
- Valentino NA, Sears DO (2005) Old Times There Are Not Forgotten: Race and Partisan Realignment in the Contemporary South. *Am J Polit Sci* 49:672-688
- Wheeler DC (2007) Diagnostic tools and a remedial method for collinearity in geographically weighted regression. *Environ Plann A* 39:2464-2481
- Wheeler DC, Tiefelsdorf M (2005) Multicollinearity and correlation among local regression coefficients in geographically weighted regression. *J Geogr Syst* 7:161-187
- Williamson T (2008) *Sprawl, Spatial Location, and Politics: How Ideological Identification Tracks the Built Environment*. *American Politics Research*: in press
- Woodberry RD, Smith CS (1998) Fundamentalism et al: Conservative Protestants in America. *Ann Rev Sociol* 24:25-56