Habit Formation in Voting: Evidence from Rainy Elections

TOM VOGL **(Princeton University, BREAD, and NBER)**

---

3 December 2015, 11:15-12:30, bld. 72, room 465

Abstract: We estimate habit formation in voting—the effect of past on current turnout—by exploiting transitory voting cost shocks. Using county-level data on U.S. presidential elections from 1952-2012, we find that rainfall on current and past election days reduces voter turnout. Our estimates imply that a 1-point decrease in past turnout lowers current turnout by 0.6-1.0 points. Further analyses suggest that habit formation operates by reinforcing the direct consumption value of voting and that our estimates may be amplified by social spillovers.

Habit Formation in Voting: Evidence from Rainy Elections*

Thomas Fujiwara† Kyle Meng‡ Tom Vogl§

October 2015

Abstract

We estimate habit formation in voting—the effect of past on current turnout—by exploiting transitory voting cost shocks. Using county-level data on U.S. presidential elections from 1952-2012, we find that rainfall on current and past election days reduces voter turnout. Our estimates imply that a 1-point decrease in past turnout lowers current turnout by 0.6-1.0 points. Further analyses suggest that habit formation operates by reinforcing the direct consumption value of voting and that our estimates may be amplified by social spillovers.

*We thank the Editor Matthew Gentzkow, three referees, Ethan Kaplan, Ilyana Kuziemko, Doug Miller, Stefano DellaVigna, and seminar participants at Columbia University, EESP-FGV, EPGE-FGV, Harvard University, Insper, MIT, Princeton University, UC Berkeley, UC Davis, UCLA, UCSD, and University of Toronto for comments; James Campbell, Wolfram Schlenker, and James Snyder for sharing data; and Sarah Weltman and Réka Zempleni for excellent research assistance.

†Princeton University, BREAD, CIFAR, and NBER. E-mail: fujiwara@princeton.edu

‡UC Santa Barbara. E-mail: kmeng@bren.ucsb.edu

§Princeton University, BREAD, and NBER. E-mail: tvogl@princeton.edu

1 Introduction

Voting is the cornerstone of democracy. However, social scientists, philosophers, and policymakers have struggled to explain why citizens vote and why turnout varies extensively within and across countries.[1] Because pivotal-voting models fail to provide satisfying explanations for non-negligible turnout in large elections (the "paradox of voting"), researchers have turned to theories based on intrinsic motivation. Early contributions expanded the "calculus of voting" framework to include a consumption value of turning out, alternatively known as "expressive utility" or "civic duty" (Riker and Ordeshook 1968). More recent theories explore how ethics, prosociality, and social pressure may imbue the act of voting with consumption value (Harsanyi 1977; Feddersen and Sandroni 2006; Benabou and Tirole 2006, Ali and Lin 2014). These theories find support in experimental studies showing that altruism (Fowler 2006; Fowler and Kam 2007; Dawes et al. 2011) and concerns about social image (Gerber et al. 2008; DellaVigna et al. 2014) play a role in driving voters to turn out. Despite the importance of these values for a robust democracy, existing research offers limited insight into how they develop.

We ask if voting is habit-forming, in the sense that past acts of voting raise the probability of voting in the future. In addition to speaking to theories of political participation, the answer to this question has important policy implications. If sizable, habit formation could alter the cumulative turnout benefit of programs like get-out-the-vote campaigns, mandatory voting, paid election days, and improved access to polls, shedding light on a potential mechanism behind the long-term effects of turnout interventions previously explored in the empirical literature.[2] Furthermore, habit formation may influence the optimal age for targeting citizens with these programs.[3]

This question has long intrigued economists and political scientists, partly for its importance and partly for its challenging nature. At least since Brody and Sniderman (1977), researchers have been aware that voter turnout is persistent: voting today is associated with voting in the future.

[1] Feddersen (2004) surveys these issues and notes that "it is unsettling that there is no canonical rational choice model of voting in elections with costs to vote."

[2] Prior empirical research in economics has focused predominantly on the contemporaneous effects of pivot probabilities (Agranov 2013; Hoffman et al. 2013), voting costs (Charles and Stephens 2013), and the media (Stromberg 2004; Gentzkow 2006; DellaVigna and Kaplan 2007; Enikolopov et al. 2011; Gentzkow et al. 2011; Drago et al. 2014; Falck et al. 2014) on turnout.

[3] Taking this argument even further, Lodge and Birch (2012) propose to "make electoral participation compulsory for first-time voters only," since "introducing an obligation for new electors to turn out once would... go a significant way toward breaking the habit of non-voting" and "could have a substantial and lasting impact on turnout."

But while this persistence may reflect habit formation, it may also reflect stability over time in the benefits and costs of voting. Empirically disentangling habit formation from other channels of persistence requires a source of variation in turnout that meets stringent statistical conditions. Not only must it be uncorrelated with the baseline determinants of turnout, but it also cannot have a *direct* effect on the future determinants of turnout.

We address this empirical challenge by exploiting unexpected and transitory shocks to voting costs due to rainfall on election day. Following previous studies documenting that rain decreases turnout (Gomez et al. 2007; Hansford and Gomez 2010; Fraga and Hersh 2010), our test for habit formation amounts to asking whether Election-Day rainfall decreases voter turnout not only in the current election but also during future elections. To ground the analysis conceptually, we present a framework for studying habit formation based on a simple "calculus of voting" model, in which rainfall is a transitory shock to the cost of voting. We use the framework to clarify what is required to identify habit formation, and we discuss why Election-Day rainfall fits such requirements, not only because it is orthogonal to voters' characteristics, but also because it is unexpected (not leading voters and other agents to adapt their behavior prior to election day) and transitory (affecting current but not future voting costs).

Matching daily weather data with county-level U.S. presidential election returns from 1952 to 2012, we find that both contemporaneous and lagged Election-Day rainfall reduce voter turnout. Our main estimates imply that a 1 percentage point decrease in past turnout lowers current turnout by 0.6-1.0 percentage points. Based on detailed exploration of the data, our preferred model includes year fixed effects, county fixed effects, and county-specific quadratic trends, allaying concerns about unobserved heterogeneity or confounding trends. Turnout shows no relation to rainfall on the day of the next presidential election, and it also shows no relation to daily rainfall within the two weeks before and after the current Election Day. These results confirm that only rainfall that fell precisely on the current and previous Election Days matters for current turnout. Additionally, subsamples with stronger contemporaneous effects also exhibit stronger lagged effects. For example, the effects of both contemporaneous and lagged rainfall are larger in poorer and more rural areas, where the costs of inclement weather may be greater. In contrast, a supplementary analysis of midterm elections finds no turnout effects of either contemporaneous or lagged rainfall, except in uncontested elections, which exhibit low turnout and possibly higher rain sensitivity.

We explore two dimensions of the mechanisms underlying these county-level results. First, we note that policies and other shocks that affect aggregate turnout can have persistent impacts due to both individual-level habit formation and social interactions during and between elections. Because rainfall is a county-level shock, our approach is well-suited for capturing the joint impact of these channels. A comparison of our results with existing estimates of individual persistence in voting behavior suggests a county-level social multiplier (Glaeser and Scheinkman 2002) as large as 1.7, implying that for every percentage point increase in turnout resulting from individual habit formation, county average turnout rises 1.7 percentage points. While other factors (e.g., different populations being affected) may account for the difference between previous estimates and ours, we argue nevertheless that the large implied social multiplier is noteworthy. Second, guided by our theoretical framework, we assess which determinant of voting underlies our main result. Drawing on several additional analyses, we argue that it is unlikely to be driven by persistent changes in voting costs (including automatic de-registration of non-voters), by updates to voters' beliefs about their probability of being pivotal, or by changes in voters' preferences over election outcomes. As a consequence, the results suggest that habit formation may be driven by an increase in the consumption value of voting, as in classic economic models of habit formation in consumption (Pollak 1970; Becker and Murphy 1988).

Our attempt to disentangle habit formation from other causes of persistence in the costs and benefits of voting builds on two previous studies.[4] Gerber et al. (2003) and Meredith (2009) both exploit plausibly exogenous variation in past voting to identify the persistent effects of shocks to turnout. Gerber et al. carry out a randomized get-out-the-vote intervention, while Meredith implements a regression discontinuity design based on age thresholds for voter eligibility. Relative to their research designs, ours has both benefits and drawbacks. On the positive side, rainfall is perhaps less likely than their sources of turnout variation to have direct effects on the future determinants of turnout. In Gerber et al.'s experiment, the canvassing procedure included messages appealing

[4]Two other papers use instrumental variables methods that rely on debatable identifying assumptions. Green and Shachar (2000) estimate models where past turnout affects current turnout, including a specification where past turnout is predicted using lagged demographic controls and opinions. Denny and Doyle (2009) estimate similar models using the number of locations a respondent lived while age 16-23 as an instrument for voting in their first eligible election. In other related work, Franklin and Hoboldt (2011) show that Europeans whose first eligible election is a (low-turnout) European Parliament election vote less in national elections, while Atkinson and Fowler (2013) report that saint's day fiestas depress current and future turnout in Mexico. These papers also require added assumptions for a habit formation interpretation.

to a subject's sense of civic duty, political competition, or neighborhood solidarity; in Meredith's study design, barely eligible voters had time to acquire information in the lead-up to election day. We argue that a transitory and unexpected shock in the cost of voting, such as rainfall, may be better suited for estimating habit formation as it is less likely to alter underlying voter preferences or knowledge. Additionally, our sample covers the entire continental U.S. over 60 years, during which all counties experienced rainfall on at least one Election Day. Gerber et al. find effects of a get-out-the-vote campaign preceding the 1998 midterm election on turnout in a 1999 local election in New Haven, CT, while Meredith's results are based on young Californians in the 2000-2006 period.[5] On the negative side, however, our county-level research design cannot isolate individual-level habit formation from the amplifying effects of social interactions.

The paper also relates to three other strands in the literature. First, it speaks to the empirical literature on the determinants of turnout. Several papers study the impacts of media exposure (Stromberg 2004; Gentzkow 2006; DellaVigna and Kaplan 2007; Enikolopov et al. 2011; Gentzkow et al. 2011; Drago et al. 2014; Falck et al. 2014), but they exploit *persistent* variation in media exposure and hence are not able to address the impacts of a transitory shock to turnout. Our results complement this literature by suggesting that the long-run effects of media exposure on turnout may be partly driven by habit formation. Other subsets of the literature do focus on transitory shocks and their persistent effects. For example, Madestam et al. (2013) and Madestam and Yanagizawa-Drott (2011) use rainfall on Tax Day and Independence Day, respectively, to estimate the effect of participating in Tea Party protests and Independence Day celebrations on political preferences and behavior. A voluminous literature within development economics also uses weather shocks as a source of exogenous variation in agricultural productivity and income.[6] Relatedly, Kaplan and Mukand (2011) find persitence from other shocks, showing that citizens registered to vote short after the September 11, 2001 are more likely to be registered as Republicans even half a decade after the terrorist attacks. In research that speaks to possible psychological mechanisms underlying our results, Mullainathan and Washington (2009) show that the act of voting for a candidate leads

[5] In both contexts, political competition was low and Democrats dominated federal elections.

[6] Dell et al. (2014) provide an extensive survey of economics papers that estimate the effects of weather, covering studies in both developed and developing contexts. The vast majority of papers in this literature deals with weather shocks over periods longer than a day, such as a year or agricultural season. Other papers using daily weather shocks, as we do, include studies of the effect of race riots on urban development in the U.S. (Collins and Margo 2007), of political protests on policy changes in France (Huet-Vaughn 2013), and of extreme temperature on mortality (Deschennes and Moretti 2009).

to improved opinions of that candidate, consistent with cognitive dissonance theory. Many of their arguments regarding the choice of candidate can apply to our study of the turnout decision.

Second, our results add empirical evidence to a recent theoretical literature exploring aggregate turnout when past voting experiences influence future voter participation. Building on an earlier paper by Kanazawa (1998), Bendor et al. (2003) model the behavior of voters who guide their turnout with rules of thumb over past turnout decisions and election outcomes. Their model predicts substantial equilibrium turnout, even in large electorates, thus providing a potential solution to the paradox that citizens vote in large numbers despite having little chance of individually swinging the election. While our framework differs from their model (which does not include a "calculus of voting"), our results corroborate features of their theory. Given our focus on the development of voting habits, our results lend particular support to Fowler's (2006) extension of their theory to incorporate habitual voters who always turn out. Finally, our results speak to a broader literature on habit formation in other aspects of economic activity.[7]

2 Motivation: Age Patterns in Voting

To motivate our interest in habit formation, Figure 1 displays U.S. federal election turnout as a function of age using data from the Current Population Survey (CPS) Voter Supplement, 1980-2010. The figure presents two panels, one including all ages from 18 to 80 and one focusing on the first decade of voter eligibility. Both panels plot age-specific means and local linear regressions with bandwidths of 2 years, separately for Presidential and midterm elections. Two aspects of the age patterns are suggestive of habit formation.

In Panel A, which spans the lifecycle, turnout increases monotonically in age through the late 60s, at which point it gradually declines, perhaps due to the onset of old-age disability. This pattern is striking because the opportunity cost of time—wages, employment, childrearing—follows a similar age profile. Hence, over most of the lifecycle, turnout increases with age despite a rising cost of voting.[8] The natural implication is that the perceived benefits of voting increase with age more

[7] For example, habit formation has drawn interest for its potential to resolve puzzles related to asset markets (Constantinides 1990), economic growth (Carroll et al. 2000), monetary policy (Fuhrer 2000), and trade (Atkin 2013).

[8] One possibility is that voters learn how to minimize the costs of voting—faster transportation to the polls, more practical times to vote—as they age. But this argument is difficult to square with the fact that turnout rises with age even in late middle age, when individuals have been eligible to vote for more than two decades.

rapidly than does the opportunity cost of time. Although this implication has several potential explanations, habit formation may play an important role.

In fact, one can glean some evidence of habit formation from these age profiles alone. To highlight this evidence, Panel B of Figure 1 zooms in on ages 18-27, showing scatter plots of the deviations from the fitted relationship for all ages, presented separately by previous presidential election eligibility.[9] The scatter plots display clear jumps in turnout from age 19 to age 20 in midterm elections and from age 21 to age 22 in presidential elections, exactly matching the age pattern of eligibility for one previous presidential election. Similar jumps are evident at the age cutoffs for eligibility for two previous presidential elections: 24 in midterm elections and 26 in presidential elections. The four jumps average 2.1 (S.E. = 0.7) percentage points. Since presidential elections tend to involve high turnout, these discontinuous increases in age-specific turnout suggest habit formation: past voting experiences increase the likelihood of future voting. The evidence is similar to that of Meredith (2009), who studies age patterns in voting using more finely-grained age data from California. However, as we discuss in the next section, although it *suggests* habit formation, one needs additional assumptions—which may fail to hold—to interpret the effect of past eligibility as the effect of past voting *per se*.

3 Identification: Insights from the Downsian Framework

For our purposes, "habit formation" means that the act of voting today, holding constant voters' characteristics, affects voting decisions in the future. Our central contribution is to separate "habit formation" from "persistence" in general, which can be explained by serial correlation in the benefits and costs of voting. For instance, those with interest in politics or a strong sense of civic duty will turn out often, while those with low levels of these variables will rarely vote. A regression of current turnout on its lagged values is thus a poor test of habit formation, since persistent unobserved heterogeneity may explain any serial correlation in voting.

In this section, we draw on the "calculus of voting" framework to pinpoint the conditions necessary to identify habit formation. Within this framework, we discuss previous research designs to estimate habit formation and explain why they may fall short of these conditions. As an alternative

[9] All birth cohorts in Figure 1, Panel B, became eligible to vote at age 18 under the 26th Amendment of 1971.

source of identifying variation, we propose Election-Day rainfall. We take care to list both the benefits and the limitations of our approach, as well as to raise interpretation issues arising from the fact that rainfall affects entire communities, rather than individuals.

3.1 Downsian Framework

To be explicit about the identification problem, we consider habit formation within the "calculus of voting" framework of Downs (1957), Tullock (1967), and Riker and Ordeshook (1968). Citizen i has probability P_{it} of being the pivotal voter in period t's election: with probability P_{it}, her preferred candidate wins if and only if she votes. She obtains benefit B_{it} if her preferred candidate wins the election in period t, regardless of whether she voted, and also enjoys direct utility D_{it} from the act of voting, regardless of the election outcome.[10] The product $P_{it}B_{it}$ is commonly known as the "instrumental utility" of voting, representing the expected policy payoff from the act of voting. In contrast, D_{it} is the direct consumption value the citizen gains from the act of voting, also known as the "expressive utility" of voting. It represents benefits from carrying out a civic duty, adhering to an ethical standard, or complying with social pressure. The citizen incurs cost C_{it} from voting, also regardless of the election outcome. She votes if and only if her net utility of voting $P_{it}B_{it} + D_{it} - C_{it}$ is positive. Denote the voting decision as V_{it}, which equals 1 if the citizen votes, 0 otherwise.

We wish to identify whether $V_{i,t-1}$ affects V_{it}, but as mentioned above, an identification problem arises: the model terms $\{P_{it}, B_{it}, D_{it}, C_{it}\}$ may be serially correlated within an individual. As such, we take advantage of a transitory shock ξ_{it} to the net utility of voting.[11] Incorporating this shock into the framework above, the citizen votes if and only if:

$$P_{it}B_{it} + D_{it} - C_{it} + \xi_{it} > 0 \tag{1}$$

In principle, ξ_{it} could work through any term of the Downsian framework, but in practice, our strategy relies on shocks to C_{it}, while existing research relies on shocks to D_{it}.

Whatever term it affects, the shock must satisfy two conditions. First, it must be independent

[10] In the American context, if V_{it}^R is the benefit to citizen i if a Republican candidate wins and V_{it}^R the benefit if a Democratic candidate wins, then $B \equiv \left|V_{it}^R - V_{it}^D\right|$.

[11] We assume that the support of ξ_{it} includes values that change some citizens' voting decisions.

of the baseline determinants of voting in the same period:

$$\{P_{it}, B_{it}, D_{it}, C_{it}\} \perp \xi_{it} \tag{2}$$

Condition (2) allows us to estimate the effect of the shock on contemporaneous turnout. The second condition for the shock is dynamic:

$$\{P_{it}, B_{it}, D_{it}, C_{it}, \xi_{it}\} \,|V_{i,t-1} \perp \xi_{i,t-1} \tag{3}$$

which states that, conditional on the voting decision the last period, the last period's shock is independent of the current determinants of voting. Condition (3) is similar in spirit to the exclusion restriction in a standard instrumental variables setup, implying that $\xi_{i,t-1}$ affects period t voting only through its effect on period $t-1$ voting and *not* by directly affecting P_{it}, B_{it}, D_{it}, or C_{it}. Additionally, because the determinants of voting in period t include both the baseline terms of the Downsian framework and the shock ξ_{it}, condition (3) implies that ξ_{it} cannot be serially correlated. Under these conditions, an association between $\xi_{i,t-1}$ and V_{it} provides evidence of habit formation. In Section 4.2, we discuss how we rely on these conditions to estimate a local average treatment effect of $V_{i,t-1}$ on V_{it}.[12]

Even if ξ_{it} is independent of the baseline benefits and costs of voting before the realization of the shock, it may not satisfy condition (3). For example, consider a randomized intervention that encouraged citizens to vote in period $t-1$. Randomization guarantees that the intervention satisfies condition (2). But depending on its nature, the intervention may *directly* influence a citizen's consumption value or cost of voting for many periods into the future. In this case, $\xi_{i,t-1}$ affects V_{it} through D_{it} or C_{it}, not solely through $V_{i,t-1}$.

3.2 Previous Research Designs

Two important contributions to the literature on voting persistence rely on research designs that satisfy condition (2) but not necessarily condition (3). The first involves a field experiment, while the second exploits a regression discontinuity design.

[12] Together with the assumption stated in footnote 13, conditions (2)-(3) are equivalent to Condition (1) in Imbens and Angrist (1994).

In the first study, Green et al. (2003) report the results of a randomized trial of a get-out-the-vote (direct mail and canvassing) campaign conducted in New Haven, CT, prior to the general election of 1998. They find higher turnout in the treatment group in both the 1998 general election and the 1999 local election, which they interpret as the effect of habit formation. However, this interpretation assumes that the campaign had no direct lasting effect on the benefits or costs of voting. Although plausible, this assumption is far from certain. For example, if the campaign raised voters' perceived benefit of voting, and this effect lasted more than a year, then condition (3) would be violated. In fact, the experimental get-out-the-vote campaign embedded several messaging treatments that appealed to a subject's sense of civic duty, political competition, or neighborhood solidarity. Because they aim to exploit or manipulate a subject's emotions, these messaging treatments may plausibly affect D_{it} in a lasting way. In other words, the shock to D_{it} may not be *transitory*.

A similar logic applies to Meredith (2009), who uses data from California to compare the voting behavior of those whose 18th birthday was just before the 2000 general election to that of those who turned 18 just after. This approach is similar to ours in Figure 1, Panel B, except that it uses more finely-grained age data on a sample from a particular state in a shorter period. Meredith estimates that those barely eligible to vote in 2000 are more likely to vote in 2004. However, to interpret this evidence as habit formation in voting *per se*, one must assume that experiencing a presidential campaign while eligible to vote for the first time has no persistent direct effects on a citizen's tastes and costs. As Meredith notes, citizens who know they will be eligible to vote may pay more attention to media coverage and campaign messages than those who know they will not be eligible. Because those turning 18 around Election Day are likely to be high school students, they may also pay more attention to school-based efforts to increase civic engagement. If exposure to these sources of information during an individual's first eligible election has persistent effects on the perceived benefits and costs of voting, then condition (3) is violated. In other words, the change in voting costs is *expected*, which may lead to exclusion restriction violations.

In summary, although Green et al. (2003) and Meredith (2009) have moved the literature substantially forward, we do not know the extent of possible exclusion restriction violations in their study designs. To identify habit formation, a shock to the costs or benefits of voting must be *transitory* and *unexpected*.

3.3 Identification using Election-Day Rainfall

As an alternative approach to identifying habit formation in voting, we exploit a transitory shock to the cost of voting: Election-Day rainfall. Four important characteristics of this shock justify our choice. First, as we show below (and as previous research has established), rainfall reduces contemporaneous voter turnout. Second, it is outside of the control of voters, candidates, or any other political agent and is orthogonal to the baseline benefits and costs of voting, before the realization of the shock. Third, it is transient and thus affects contemporaneous voting costs without having a direct effect on the future costs or benefits of voting. Fourth, net of the year fixed effects, county fixed effects, and county-specific trends we include in our econometric model (see Section 4.2 for details), the remaining variation in rainfall is extremely difficult to predict long in advance. Given this difficulty, voters and candidates are unlikely to modify their behavior leading up to an election in anticipation of a rainfall shock. We emphasize this point in light of our discussion of Meredith's (2009) results: if a shock to voting costs can be predicted well in advance, voters and political campaigns may adapt their consumption and production of political information respectively in the period leading up to the election, which may lead to a violation of condition (3).

At the same time, our research design is not unsusceptible to potential exclusion restriction violations. For example,, the unpleasantness of voting on a rainy day may influence the affective state that voters associate with the act of voting. In this case, the positive act of voting (rather than the negative act of abstaining) on a rainy day may reduce future voting propensity, so an effect of lagged rainfall need not imply habit formation. However, this hypothesis assumes that voters fail to blame bad weather for the unpleasantness of voting. Given that most voters have experienced many rainy days in the past, we conjecture that such attribution error is minimal, although we acknowledge the possibility that it biases our results.

Supposing that the exclusion restriction holds, Election-Day rainfall can identify the effect of past turnout on future turnout at the local level. However, because rainfall is an aggregate shock, affecting all individuals within a community, this aggregate form of habit formation may differ from the individual-level form of habit formation and is a notable limitation of our approach. In addition to reflecting the individual-level phenomenon, aggregate habit formation may incorporate additional autoregressive effects arising from social interaction effects—for example, if people speak

to their neighbors about positive voting experiences between elections. Given the literature's current emphasis on social influences on voter turnout, this refinement of the parameter of interest may be desirable, although we acknowledge that many readers may be interested in the individual-level parameter. This possible social multiplier creates ambiguity; although exclusion restriction violations may inflate previous estimates, our estimates may be larger yet, due to a social multiplier.

As with other research designs to identify habit formation in voting, ours cannot isolate particular mechanisms. This limitation is common in design-based strategies to disentangle causality. Habit formation may work through $V_{i.t-1}$ affecting P_{it}, B_{it}, D_{it}, C_{it}, or some combination therein. In other words, past acts of voting may change a citizen's perceived influence on the election outcome (P_{it}), her interest in the election outcome (B_{it}), her sense of ethics or civic duty (D_{it}), or her voting costs (C_{it}). Although rainfall cannot by itself disentangle these mechanisms, we draw on other sources of variation to shed light on this issue in Section 7.

4 Data: County-Level Panel

Mid-latitude rainfall systems, as observed over the United States, can be spatially correlated across areas between 2 to 1,000 km wide with spatial extents that do not fit naturally onto political boundaries. This property implies three data requirements. First, the data pixel resolution of the rainfall data must be fine enough to guarantee that most political units cover at least one pixel so that there is variation in rainfall across neighboring units. Second, turnout data must be at the lowest political unit available in order to to reduce measurement error when pixel-level weather data is aggregated to a single value at the level of the political unit. Third, cross-sectional units in the dataset should have the broadest geographical coverage, in this case over the entire continental U.S., to guarantee a sufficient number of independent observations for a given day. Beyond these data requirements, daily rainfall variation have particular statistical properties that may generate spurious relationships. As such, after presenting a dataset that satisfies the requirements listed above, we detail the properties of our variables of interest, which inform our econometric approach.

Our dataset combines political, meteorological, and demographic information over more than 60 years for all counties in the continental United States. For political data, we use county-level presidential election returns for the years 1948-2012 to generate two variables: voter turnout, which

we define as the ratio of votes to eligible voters, and the Republican vote share.[13] In supplementary analyses, we also draw on midterm election vote returns.[14] For weather, we acquire data with the highest spatial and temporal resolution available for our area and period of interest. Daily gridded rainfall data for the continental United States from 1948-2012 come from the NOAA Climate Prediction Center's Unified Gauge-Based Analysis of rainfall.[15] This source provides pixel-level data at a 0.25 degree by 0.25 degree (or roughly 17 mile by 17 mile) resolution, which we aggregate to the county level using area weights.[16] In addition to data on politics and weather, we also draw on several county demographic and socio-economic covariates from the U.S. Census: racial composition, age structure, median income, and population density.[17] Table 1 provides means, standard deviations, and several percentiles for the variables in our analysis. Voter turnout averages at 58 percent, with a fairly symmetric distribution ranging from a 10th percentile of 42 to a 90th percentile of 76. The Republican vote share, too, is fairly symmetrically distributed around a mean of 55. Similarly, county-level covariates appear symmetrically distributed.

Election-Day rainfall exhibits several noteworthy statistical properties. First, it is a relatively infrequent event. Table 1 shows that its distribution is right-skewed, with a median of 0, a mean of 2.5 millimeters, and a 90th percentile of 7.1 millimeters.[18] Second, when Election-Day rainfall does occur, it is typically experienced by many counties at once. Figure A1 plots the share of U.S. counties that experience any rainfall and rainfall between 0 and 4 millimeters on election days across our sample period, showing that the county share exhibits a roughly bimodal distribution that oscillates between low and high values. Third, while extreme rainfall on Election Day is rare, all counties experience rainfall at some point in our sample period. Figure A2 plots the cumulative share

[13] We obtained county-level vote totals for 1948-2000 from James Snyder, which we supplemented for years 2004-2012 using David Leip's Atlas of U.S. Presidential Elections. We obtained estimates of the number of eligible voters from Genzkow et al. (2011) for the years 1952-2004, which we supplemented with our own estimates using similar methods (based on interpolated data from the U.S. Decennial Census) for the years 2008 and 2012. Because the denominator of the turnout rate is estimated with error, estimated turnout rises above 100 in 0.24% of the observations. We include these observations in the reported analyses, but the results are unchanged if we omit them or top-code turnout at 100.

[14] The source for the midterm data for the 1950-1990 period is ICPSR Study 0013, which we salso uplemented for years 1994-2014 using David Leip's Atlas.

[15] Thus far, we have referred to our treatment as rainfall for ease of exposition. Technically, our treatment variable is precipitation, which captures all forms of condensation from water vapor including rainfall, snow, hail, and ice. Most of the precipitation that falls on the continental U.S. during Election Day is rain.

[16] We validated our constructed weather data against historic weather station data from Weather Underground. Results are similar if we use deviations from long-term norms rather than levels.

[17] We obtained these covariates from Haines (2010) and the website http://quickfacts.census.gov/.

[18] The American Meteorological Society (http://glossary.ametsoc.or/wiki/rain) defines rain as "light" when it falls at a rate of 2.5 millimeters per hour or less and "heavy" when it falls at a rate of more than 7.6 millimeters per hour.

of counties that experienced any rainfall over the sample period, indicating that nearly all counties have experienced Election-Day rainfall by 1972, or 20 years into our sample period. This finding implies that our results use variation from all counties. Fourth, variation in Election-Day rainfall differs considerably across U.S. counties. Figure A3 displays the histogram of the standard deviation in Election-Day rainfall across counties and shows a fairly large spread in rainfall variability across counties. Although all counties provide identifying variation, some do more than others.

Election-Day rainfall and turnout also exhibit common county-level trends, a feature of our data that informs our regression specifciation. To demonstrate the presence of these common trends, we separately estimate county-level linear trends in Presidential turnout and Election-Day rainfall, net of year and county fixed effects. Figure A4 plots county trends in turnout against the county trends in Election-Day rainfall. A strong positive relationship is apparent: counties with a positive trend in Election-Day rainfall also tend to exhibit a positive trend in turnout. Interestingly, these trends are specific to Election-Day rainfall and do not characterize rainfall for the days around Election Day. Figure A5, Panel A, displays coefficients from separate regressions of county turnout trend on rainfall trend during Election-Day and days immediately preceding and succeeding it. The turnout trend has a significant positive association with the Election-Day rainfall trend, as already shown in Figure A4, but also has significant negative associations with the trends in rainfall 7 days before and 4-5 days after Election Day. This pattern suggests that daily rainfall exhibits stochastic trends that happen to be correlated with turnout trends. The oscillation of the coefficients between negative and positive values, which is due to the movement of spatially correlated weather systems, implies that the rainfall trend cannot be captured by average daily rainfall over a period around Election Day. Panel B of Figure A5 shows that the same pattern exists for calendar-day rainfall around November 1st (rather than Election Day, which changes from year to year), suggesting that Panel A's result is not an artifact of how Election-Day rainfall is constructed. Taken together, these results suggest that any credible estimation strategy must detrend the rainfall data.

The estimation strategy must also account for spatial correlations in rainfall and turnout in performing statistical inference. Weather systems induce spatial correlation in rainfall; the design of the electoral college and the bundling of presidential and state-level elections induce correlated turnout incentives across counties within a state. To document the consequences, we regress both turnout and rainfall on a year fixed effect, a county fixed effect, and a county-specific linear trend,

and we map the residuals for three example years—1964, 1984, and 2004—in Figure A6. rainfall residuals are clustered over large areas, while turnout residuals tend to cluster within state borders. As a result, in all analyses, we cluster standard errors at the state level, thus allowing for arbitrary error covariance across counties in a state over any period of time.[19]

5 First Stage: Effect of Contemporaneous Rainfall on Turnout

Given the complex properties of the data documented in Section 4, this section studies how to best to specify a regression of turnout on contemporaneous Election-Day rainfall. This regression is equivalent to a first-stage relationship in our setting. Only after selecting an appropriate specification for contemporaneous rainfall do we move on to estimating the effect of lagged rainfall in Section 6. In other words, we select our preferred model because it provides a stable and credible estimate of the effect of contemporaneous rainfall on turnout, *not* because it shows evidence of habit formation.

5.1 Econometric Method

We explore variants of the following specification, for county c in election year t:

$$turnout_{ct} = \beta_0 rain_{ct} + \tau_t + \eta_c + f_c(t) + \epsilon_{ct} \tag{4}$$

τ_t and η_c are year and county fixed effects, respectively, and $f_c(t)$ is a flexible county-level trend. To match our core analysis sample in Section 6, we estimate equation (4) using turnout and rainfall data from Presidential election years during 1952-2012. Two procedures help us assess the robustness of our modeling choices. First, we gauge the sensitivity of the coefficients on Election-Day rainfall to different fixed effects and trends specifications, as well as to the vector of demographic covariates listed in Table 1. Second, we run placebo tests by re-estimating equation (4) including rainfall on days before and after Election Day, which should have no effect on turnout, as additional covariates.

[19] Our conclusions remain unchanged when we use Conley's (1999) non-parametric estimator for standard errors allowing for arbitrary spatial dependence in a 1500 km radius.

5.2 Results

Estimates of β_0 from equation (4) and its variates appear in Panel A of Table 2. In column (1), which includes no covariates, Election-Day rainfall has a significantly negative association with turnout, although the association disappears upon the addition of county and year fixed effects in column (2). With the further addition of county-specific trends in columns (3)-(5), the association becomes significantly negative again, and the magnitude is stable across specifications with linear, quadratic, and cubic trends. The trends specifications suggest that one millimeter of rainfall decreases turnout by 0.05-0.07 percentage points.

To shed light on which of these specifications most credibly estimates the causal effect of precipitatation, we re-estimate them in a series of regressions that separately add rainfall on each day from two weeks before to two weeks after Election Day. As shown in Figure A7, the models with no covariates, with only county and year fixed effects, and with the addition of linear county trends all produce multiple significant placebo effects of non-Election Days. At the 5 percent level, the three models find significant placebo effects on 14, 5, and 6 of 28 non-Election Days, respectively. In contrast, the models with quadratic or cubic trends find only one significant placebo effect, as expected by chance. To summarize these appendix results, Panel B of Table 2 adds average rainfall in the two weeks surrounding Election Day to each specification. Consistent with the results in Figure A7, the two-week average appears to have a significant effect on turnout in columns (1)-(3) but not in columns (4)-(5), which include quadratic or cubic county-specific trends. As a further check, Panel C of Table 2 adds demographic covariates to each specification. Although the specification with no fixed effects or trends is sensitive to the inclusion of these covariates, the trends specifications are not. Taken together, these findings suggest that the models with quadratic or cubic trends isolate the most credible source of variation in Election-Day rainfall. The importance of location specific trends is not unique to the United States. In a study of rainfall and turnout in Norway, Lind (2015a, 2015b) estimates spurious effects of rainfall on days before and after election day when he omits location-specific time trends.[20]

[20] Another possibility would be to include state-year effects. However, the spatial extent of weather systems and the spatial interpolation used by the NOAA leave little meaningful variation net of state-year effects.

6 Habit Formation: Effect of Lagged Rainfall on Turnout

Based on Section 5, we carry out our main analysis of habit formation using a regression specification with county-specific quadratic trends. We first describe the regression specification and how we use it to quantify habit formation. We then present the main results, followed by a set of robustness checks, an exploration of heterogeneity, and a suppplementary analysis of midterm elections.

6.1 Econometric Method

For county c in election year t, the main regression specification is:

$$turnout_{ct} = \beta_0 rain_{ct} + \beta_1 rain_{c,t-1} + \tau_t + \eta_c + f_c(t) + \varepsilon_{ct} \tag{5}$$

where $f_c(t)$ is a county-specific quadratic polynomial in time. The time subscript $t-1$ corresponds to the previous election, four years earlier. A negative estimate of β_1 would indicate that the effect of Election-Day rainfall persists four years later, providing reduced-form evidence of habit formation. Recalling condition (3) in the Downsian framework, this interpretation relies on the assumption that the shock to voting costs is serially uncorrelated. Because we control for current rainfall, the residual variation in lagged rainfall satisfies this assumption by construction.

While this reduced-form evidence is instructive, an auto-regressive model would be more appropriate for quantifying the degree of habit formation:

$$turnout_{ct} = \rho turnout_{c,t-1} + \nu_{ct} \tag{6}$$

As we discussed in Section 3.1, ordinary least squares regression does not identify this model. However, if we define the error term as $\nu_{ct} \equiv \tau_t + \eta_c + f_c(t) + \varepsilon_{ct}$, then we can use estimates of β_0 and β_1 to compute an estimate of the causal parameter ρ: $\hat{\rho} = \frac{\hat{\beta}_1}{\hat{\beta}_0}$ converges in probability to ρ.[21] This ratio can be seen as an instrumental variables (IV) estimator for ρ, in which lagged rainfall serves as an instrument for lagged turnout. We implement this single-equation procedure as our main estimation strategy, using the delta method for variance estimation. Because our weather data start in 1948,

[21] One concern with this approach is that inter-county migration may bias downward $\hat{\beta}_1$ and $\hat{\rho}$. Molloy et al. (2011) report 5-year cross-county migration rates of almost 20 percent, although over half of these flows are within-state. Because counties in the same state share weather patterns, we expect little bias from migration. Indeed, when we run our regressions at the state level instead of the county level, our conclusions do not change.

we estimate equation (5) using turnout in 1952-2012 and rainfall in 1948-2012.

One shortcoming of this strategy is that β_0 represents the effect of rainfall that occurred in the years 1952-2012, while β_1 represents the persistent effect of rainfall that occurred in the years 1948-2008. If the effect of contemporaneous rainfall on turnout changes over time, then this four-year mismatch may lead to biased estimates of ρ. A two-stage least squares estimator, which uses rainfall in the years 1948-2008 to identify both β_0 and β_1, can address this concern. For each year t in 1952-2012, the dependent variable is turnout in year t, the endogenous variable is turnout in year $t-1$, the instrument is rainfall in year $t-1$, and the exogenous control variable is rainfall in year t. Therefore, the first-stage regression relates turnout for the years 1948-2008 to contemporaneous and future rainfall, while the reduced-form regression is equation (5).

Three aspects of these estimators merit further discussion. First, as with other IV estimators, they requires the monotonicity assumption that turnout weakly decreases in rainfall for all units in our sample. If our unit of observation were the individual, this assumption might not hold. For individuals who enjoy outdoor leisure activities or work in industries like construction or tourism, the time cost of voting may fall on rainy days. Alternatively, individuals who particularly dislike congestion at the polls might vote only in rainy elections, which they anticipate will have low turnout. However, we study counties, not people, and the monotonicity assumption is more likely to hold at the county level.[22]

Second, our estimators identify a local average treatment effect (LATE) of past on current turnout, where the relevant population of compliers is made up of citizens on the margin between voting and abstaining: that is, citizens with $P_{it}B_{it} + D_{it} - C_{it}$ close to zero. This point may have important implications for comparisons with existing research. Green et al.'s (2003) experiment—which gives citizens a small push to vote—has similar compliers, but Meredith's research design—which lowers voting costs from infinity to a finite number—includes a broader swath of the electorate among its compliers, and these compliers may have a different LATE. We return to this issue when discussing the magnitudes of our results in Section 7.

Third, as mentioned in Section 3.3, our estimators do not necessarily identify habit formation at the individual level. In the presence of social interactions, β_0, β_1, and ρ are aggregate effects that

[22]Moreover, Table 4 explores the heterogeneity of effects by county characteristics, while the Online Appendix provides multiple estimates based on restricting the counties or years included in the sample. None of these results suggest that the effect of rainfall on turnout is positive for a subpopulation of counties.

may differ from individual effects. In particular, ρ for a county may be larger than the individual-level habit formation parameter. The magnitude of this difference depends on the size of a "social multiplier" (Glaeser and Scheinkman, 2002). We also discuss this issue further, providing a more formal analysis and evidence on the possible magnitude of social multipliers, in Section 7.

6.2 Results

6.2.1 Main Results

Table 3 shows that the turnout effects of rainfall persist to future elections. Column (1) implements our single-equation approach, estimating equation (5) for turnout in election years from 1952 to 2012. Coefficients on both contemporaneous and lagged rainfall are statistically significant at the 1 percent level, with turnout falling 0.063 and 0.058 percentage points per millimeter of contemporaneous and lagged rainfall, respectively. Our finding changes little with the addition of county-level covariates in column (2). In both columns (1) and (2), the implied habit formation parameter ρ is slightly above 0.9, implying that a 1 percentage point rise in period $t-1$ turnout increases period-t turnout by slightly more than 0.9 percentage points.[23] This estimate of habit formation in voter turnout is substantially larger than existing estimates in the literature.

In the alternative, two-equation approach to estimating ρ, equation (5) is the reduced form, while the first stage regresses turnout on current and future Election-Day rainfall for the period that starts and ends four years earlier. Columns (3)-(4) present this first-stage regression, with and without demographic controls. The coefficients on contemporaneous rainfall are similar in magnitude and statistical significance to the coefficients in columns (1)-(2) and in the preferred specifications in Table 2. On the other hand, future Election-Day rainfall does not affect current turnout. This placebo result indicates that any possible residual serial correlation in rainfall does not meaningfully bias our results. In the two-stage least squares procedure, the contemporaneous coefficient in column (3) or (4) represents the first-stage effect of the instrument, while the lagged effect in column (1) or (2) represents the reduced-form effect of the instrument.[24] The resulting two-stage least squares estimator of ρ is slightly above 0.9, just as in the single-equation ratio estimates of columns (1)-(2).

In columns (5)-(8), we restrict the sample to the years 1956-2008, so we can include additional

[23] If we weight observations by county population, we obtain nearly identical estimates of ρ.

[24] As a result, the ratio of the contemporaneous effect in column (3) or (4) to the lagged effect in column (1) or (2) is the Wald estimator for ρ.

leads and lags. We first re-estimate equation (5) in this shorter sample, with results appearing in columns (5)-(6). Both the contemporaneous and lagged coefficients shrink slightly but remain statistically significant. Because the lagged coefficient shrinks more than the contemporaneous, the estimate of ρ also shrinks, by roughly one-third, although it is within the confidence interval of the full-sample estimate. When we include an additional lead term and an additional lag term in columns (7)-(8), the lead term is small and insignificant—just as in columns (3)-(4)—and the additional lag term suggests that the effect of rainfall continues to dissipate as time passes. With the additional lag term, we have three estimators for ρ: $\frac{\beta_1}{\beta_0}$, $\frac{\beta_2}{\beta_1}$, and $\sqrt{\frac{\beta_2}{\beta_0}}$, where β_2 is the coefficient on $rain_{c,t-2}$. We report their average in columns (7)-(8), and it is identical to the estimates of ρ in columns (5)-(6).

As a falsification test, we re-estimated equation (5) separately adding rainfall on each day from two weeks before to two weeks after Election Day, as well as its lagged value. The results are summarized in Figure 2, which plots the estimated effect of current and lagged rainfall in these non-Election days. In support of our results, none of these 28 placebo estimates yield an effect of rainfall, both in current or lagged form, that is as large as the one we estimate for Election Day. Most effects are situated close to zero, with the Election Day parameters being a clear deviation from them.

Because the changes in coefficients across different sample definitions may raise questions about the robustness of the results, the Online Appendix presents several sensitivity analyses. Figure A8 first checks that no single state or year is influential. In 49 estimations that leave out a single state (48 continental states plus Washington, DC) and 16 estimations that leave out a single year, the point estimates and significance levels of β_0 and β_1 vary little. However, the effect of rainfall may still vary over time. Figure A9 tests this hypothesis by estimating β_0 and β_1 in rolling 8-election windows, from 1952-1980 to 1984-2012. Both coefficients are negative and significant for the early periods but become insignificant and small in magnitude for the later periods. The declining importance of weather for voter turnout has several potential explanations, which we leave for future research. For our purposes, the important finding is that we only find lagged effects in periods that also exhibit contemporaneous effects.

Aside from these analyses assessing sensitivity to sample definition, the Online Appendix also reports several specification checks. Table A1 presents estimates from a range of alternative specifications. As in Table 2, if we estimate the model with only fixed effects and not trends, we find

no effect of rainfall. Specifications with linear, quadratic, or cubic county-specific trends all lead to significant effects of contemporaneous and lagged rainfall, with implied ρ values of 0.9. If we control for trends using decade-county fixed effects instead of polynomials in time, the standard errors increase substantially, rendering the estimates of β_0 and β_1 insignificant. However, β_0, β_1, and ρ maintain similar magnitudes, and ρ remains statistically significant. Finally, the results are also robust to specifications that interact the year fixed effects with the demographic covariates. In a separate specification check, Figure A10 shortens the sample further to add lag and lead terms from $t-5$ to $t+2$. The coefficients start at -0.1 for contemporaneous rainfall and its first lag and then shrink, becoming small and insignificantly different from zero at the fifth lag. Similarly to the result in Table 3, the coefficients on the two lead terms are insignificantly different from zero. To assess whether the linear specification is appropriate, Figure A11 re-estimates the regression in Figure A10 in a semi-parametric specification with rainfall binned into 7 categories. The results reveal no noteworthy non-linearities.[25]

6.2.2 Heterogeneity

Who responds to current and lagged rainfall? Our aggregated data do not allow a detailed exploration of this question, but the demographic and socio-economic covariates from the U.S. Census can help shed some light on it. Table 4 reports regressions that interact these covariates with contemporaneous and lagged rainfall.[26] For comparison, column (1) repeats the main estimate of equation (5) from Table 3. Columns (2)-(5) interact rainfall with each of the covariates individually, while column (6) enters all interactions into the same regression. All models control for the main effects of the covariates included in the interaction terms.

The main conclusion that emerges from Table 4 is that richer and more densely population counties respond less to rainfall. For both contemporaneous and lagged rainfall, the interactions with log median income and log population density in columns (4) and (5) have positive and significant coefficients. In the full model of column (6), the coefficients shrink slightly in magnitude and

[25] The results in Figures A10 and A11 are nor directly comparable to ours benchmark estimates given the different sample periods. In unreported results, semi-parametric estimates with shorter lag lengths also revealed no non-linearities.

[26] We enter each pair of interactions into a separate regression because the results become noisy and uninformative when we include all of them in the same regression. We believe this problem arises because the interpolation of the covariates between census years induces correlated measurement errors.

significance, with three of the four maintaining significance at the 5 percent level. However, the broad implications do not change. According to the coefficients from column (6), a move from the 10th to the 90th percentile in the household income distribution weakens the response to a millimeter of contemporaneous or lagged rainfall by 0.07-0.09 percentage points, while similar move in the population density distribution weakens the response by 0.15-0.17 percentage points. These findings match the conventional wisdom that inclement weather imposes greater costs on poorer and more rural voters, due to their limited access to transportation and, in the case of density, longer distances from the polls. At the same time, the interactions of each covariate with contemporaneous and lagged rainfall are not significantly different from each other, so we cannot reject the null hypothesis that the degree of habit formation varies with income or population density.

6.2.3 Midterm Elections

Our analyses so far have focused on turnout in presidential elections, due to its importance and to the existing evidence of its responsiveness to rainfall (Gomez et al. 2007; Hansford and Gomez 2010). But the results naturally lead to questions about whether one can find the same patterns in midterm election data. To this end, Table 5 carries out a supplementary analysis of midterm election turnout. Relative to the baseline regression specification (5), the analysis of midterm turnout requires a modification to deal with the uncontested elections. Because the stakes are low, turnout in congressional elections with a single candidate may be low, and it may also respond differently to rainfall. To address this issue, all regressions in Table 5 include indicators for whether all congressional elections in a county are uncontested in t and $t-1$. Half of the regressions also include interactions between these indicators and Election-Day rainfall.

Columns (1) and (2) of Table 5 report regressions without interaction terms, which show evidence of neither contemporaneous nor lagged effects. Both coefficients are positive and slightly smaller than their standard errors. Electoral competition matters, however. Turnout is 10 percentage points lower if the current election is uncontested and 1 percentage point lower if the last election was uncontested. One interpretation of the second result is that challengers in previously uncontested districts do not typically present a serious threat to the incumbent. Because uncontested elections are fundamentally different from contested elections, they may exhibit different rainfall effects. Indeed, the interaction terms in columns (3) and (4) indicate a much stronger effect of rainfall in currently uncontested

elections. The lagged interaction term, while not statistically significant, is of the same magnitude as the current interaction term. Combined, these estimates shed new light on the types of elections in which rainfall decreases turnout, but they do not provide conclusive evidence on habit formation in midterms.

7 Discussion: Magnitudes and Mechanisms

7.1 Assessing Magnitudes

7.1.1 Comparison with Past Research

Past estimates of habit formation in voting serve as an important and interesting basis for comparison for our findings. We estimate a habit formation parameter ρ between 0.6 and 1.0, with our preferred specification delivering an estimate of 0.9. By comparison, Gerber et al. (2003) place their persistence parameter at 0.5 in a get-out-the-vote experiment, while Meredith (2009) estimates persistence to be 0.075 using a regression discontinuity design based on voting age restrictions.

Meredith's estimate is an order of magnitude smaller than both Gerber et al.'s and ours, but his study design identifies a different estimand that is likely to be small. In both Gerber et al.'s context and our own, always-voters exist, such that the estimation strategies identify the LATE for marginal voters (compliers). In contrast, Meredith strategy does not allow for always-voters; individuals just short of their 18th birthdays cannot vote under any circumstance. As a result, Meredith effectively recovers a treatment-on-the-treated persistence parameter that averages the effect of past on present voting for voters who, were they eligible to vote, would be both marginal *and* infra-marginal. The effect is zero for infra-marginal voters, which justifies Meredith's small (though statistically significant) estimate.

The fact that our estimate exceeds that of Gerber et al. presents a greater puzzle. On the one hand, this result may be driven by sampling error; the difference is not significantly different from theirs at conventional levels. On the other hand, our preferred point estimate is roughly 90 percent larger than theirs, begging an explanation. We propose four possibilities. First, and perhaps most importantly, our empirical strategy may pick up interpersonal spillovers due to social interactions following election day; Gerber et al.'s design does not. Second, Gerber et al. ran their get-out-the-vote campaign just before a low-stakes midterm election and collected follow-up data

on a local election one year later.[27] The effect of voting in a low-stakes midterm election on voting in a subsequent local election may be smaller than the persistent effects of turnout for presidential elections. Third, the sub-populations induced to vote may differ between the two studies. Gerber et al.'s estimate applies to residents of New Haven whereas our study covers the entire country.[28] Finally, Gerber et al. lost 14 percent of their sample to follow-up. Although attrition was evenly distributed across control and treatment groups, the attriters in the treatment group may have differed in unobservable ways from the attriters in the control group, which would undermine the study design. For the remainder of this section, we focus on the possible role of social interactions.

7.1.2 Spillovers and Social Interactions

Individuals may induce others to vote in the future by sharing past voting experiences (Nickerson 2008; Gerber et al. 2008; Bond et al. 2012; DellaVigna et al. 2013). Such social interactions can produce spillovers, implying that our county-level estimate of habit formation captures the combined effects of individual-level habit formation and social interactions. Importantly, from the perspective of evaluating prospective policies intended to boost turnout, these combined effects might be more relevant than the effect of individual-level habit formation in isolation.

Formally, let b denote the effect of a unit of rainfall on an individual's probability of turnout, and let r be the individual effect of past on current turnout. These parameters are potentially distinct from the corresponding county-level parameters β_0 and ρ. Our objective is to understand the mapping between r and ρ. In the absence of social interactions, these parameters are the same. However, in the presence of (positive) social interactions, part of the effect of past rainfall on turnout operates through social interactions, such that ρ exceeds r (Case and Katz 1991; Glaeser and Scheinkman 2002). Following Glaeser and Scheinkman's (2002) approach, we define a county-level social interactions parameter θ as follows: an individual's likelihood of voting increases by θ percentage points for every 1 percentage point increase in the average turnout of other residents of her county. We take θ to capture social interactions occurring after the current Election Day and before the next, which allows us to write the effects of current and lagged rainfall as $\beta_0 = b$ and $\beta_1 = \frac{br}{1-\theta}$, respectively, making the county-level habit formation parameter $\rho = \frac{r}{1-\theta}$.[29] Hence,

[27] In the 1998 midterm election, both federal races that involved New Haven (the site of Gerber et al.'s study) were decided by margins of more than 30 points.

[28] Additionally, the Gerber et al. experiment is specific to 1998-1999, while we study 1952-2012.

[29] These derivations use the fact that there is a large number of voters in each county (Glaeser and Scheinkman 2002).

the strength of social interactions between election days determines the relationship between the individual- and county-level habit formation parameters.

Unfortunately, little evidence exists on the size of θ. Even if we had individual-level data, we could not distinguish r from ρ using our estimation strategy because rainfall varies at a spatially aggregate level and thus produces estimates that include the effects of social interactions. However, if we take the individual-level persistence parameter from Gerber et al. (2003) as a benchmark for individual habit formation, we can recover a value for θ. Although we have already noted separate reasons that our parameters may differ, it is nonetheless a useful benchmark. Combined with our baseline estimate of ρ at 0.93, their estimate of r at 0.51 implies $\theta = 0.45$.[30]

Is this value for θ reasonable? As a way to gauge its plausibility, we compare it with the social interactions parameter implied by individual and county-level *associations* of past and current voting, which we estimate in Table 5. In columns (1)-(3), we use self-reported individual turnout from the 1972-1984 CPS to estimate an autoregressive panel model of current on past presidential election turnout, with varying demographic, geographic, and temporal controls. Across various specifications, we find an individual level persistence parameter of approximately 0.5. Combined with our main county-level habit formation estimate of 0.93, this estimate implies $\theta = 0.46$. However, autoregressive estimates of ρ may be biased for reasons we already noted. In columns (4)-(7) of Table 5, we estimate a similar autoregressive panel model of county-level presidential election turnout, leading to a persistence parameter of roughly 0.8. If the individual and county-level estimates shown in Table 5 are biased by the same *proportion,* then the ratio between the individual and county-level coefficients provides an unbiased estimate of θ. Indeed, the coefficients presented in Table 5 imply $\theta = 0.38$. In summary, all these exercises yield similar estimates for θ, around 0.4, implying a social multiplier of $\frac{1}{1-\theta} = 1.7$.

7.2 Mechanisms

Whether or not our estimates are amplified by social interactions, they are likely to partly reflect individual-level habit formation. Recall from equation (1) that a citizen i votes if $P_{it}B_{it} + D_{it} \geq C_{it}$.

As a refinement, we could also assume that social interactions have short- and long-run components, θ^S and θ^L, with the former term capturing interactions during the current Election Day (hence creating a multiplier for current rainfall). The comparison between county and individual level effects would thus depend on the ratio of these parameters. Since this refinement is not essential to our dicussion, we omit it to simplify notation.

[30] Our "short sample" estimate of 0.63 implies a social multiplier of 0.19.

Conceivably, any of the framework's terms could depend on past voting experiences. In the Online Appendix, we use the data to explore each term's possible role in explaining our results, and we also consider whether partisan politics contribute to the explanation. We summarize our findings in this section, concluding that accumulation in expressive utility D_{it} is the most likely mechanism.

In one explanation for habit formation, a citizen learns over time about her probability P_{it} of affecting the election outcome. Under most forms of updating (e.g., Bayesian updating), P_{it} would increase after voting for the winner or not voting while supporting the loser, and decrease after voting for the loser or not voting while supporting the winner.[31] However, this explanation has two conceptual limitations. First, it is inconsistent with rational expectations and most forms of forward-looking behavior. Under such assumptions, a citizen would use all available information about her probability of being pivotal, to which her past voting experiences are not relevant. Second, it needs to confront the fact that the objective value of P_{it} is virtually zero. Either very small variations in this probability have large consequences, or voters have unrealistic beliefs. Third, the magnitude of the effect we find is inconsistent with updating P_{it}. While the theory may predict more positive updating than negative updating—by design, more voters support the winner than the loser—narrowly-decided elections should result in little habit formation on average; voters who supported the winner are of roughly the same number as voters who supported the loser. Our estimates of ρ are large even though most presidential elections during our sample period were decided by margins of less than 10 points. Corroborating evidence appears in the Table A2, which adds interactions with the national margin of victory in the previous election to our main estimates. Neither the effect of current nor lagged rainfall varies with the national margin of victory. We can also more directly test this theory's divergent predictions for voting for the winner and voting for the loser. Table A2 estimates the effects of rainfall interacted with a measure of whether a county is politically aligned with the winner or loser of the previous election. We do not find evidence that alignment with the winner affects the degree of habit formation.

A separate explanation for our results involves the strength of citizens' political preferences B_{it}. Mullainathan and Washington (2009) argue that, due to cognitive dissonance, the act of voting causes a citizen to further improve her opinion of her chosen candidate. But if a citizens have

[31] Consistent with this logic, Kanazawa (1998) and Bendor et al. (2003) posit reduced-form behavioral models in which voting for the winner increases future turnout, while voting for the loser decreases future turnout.

objective beliefs about her probability of being pivotal, then any effect on B_{it} will likely have limited consequences for her voting decision because it will be multiplied by a number approaching zero. In fact, we can leverage the fact that B_{it} is multiplied by P_{it} to more formally test whether accumulation in B_{it} can explain our results. This fact is key to distinguishing between B_t and D_t in our framework. The act of voting may lead a citizen to change her tastes regarding politics; the distinction is whether these tastes take the form of instrumental value (caring about the outcome, B_t) or expressive value (caring about voting, D_t). If voting in period $t-1$ increases B_{it}, then evidence of habit formation will be stronger when P_{it} is high. However, Table A2 interacts the rainfall variables with an *ex ante* (before rainfall) measure of expected voter pivotality, finding no evidence that the effects of rainfall vary with the objective probability of being pivotal.

The intrinsic cost of voting C_{it} provides another potential mechanism. This mechanism has two potential sources, one personal and one institutional. On the personal side, a voter must incur informational "fixed costs:" learning the location of the polling station and the best way to get there. She may also be uncertain of how much time the act of voting takes; if she is risk averse, she will become more likely to vote once she learns the true opportunity cost of voting. While these hypotheses are plausible, they seem too small to explain our large habit formation parameter.[32] On the insitutional side, election offices have at various points implemented laws that purge inactive voters from the registration rolls. These purges could lead to evidence of habit formation because repeated non-voters lost their registration, raising the cost of future voting. However, if we restrict our sample to elections with no automatic purging based on *The Book of the States* (avaliable at http://knowledgecenter.csg.org), results are similar to our benchmark.

Given our argument that the other Downsian term are unlikely to fully explain our results, expressive utility D_{it}—the consumption value of voting, stemming from civic duty, ethics, or social pressure—is perhaps the most plausible mechanism. This theory conforms with traditional interpretations of habit formation (Pollak 1970; Becker and Murphy 1988) in which "habits" are consumption tastes. However, since the concepts embedded in D_{it} are essentially unobservable, we cannot directly test it. Accumulation in D_{it} can arise from two types of processes: intrinsic and extrinsic. Intrinsic accumulation in D_{it} refers to the individual-level psychological process by which citizens develop

[32] Moreover, if informational fixed costs matter, the lagged effect of rainfall should be smaller in counties with older populations (who have more experience going to the polls), which is not the case in our data.

attachments to the act of voting, independent of social influences.[33] Extrinsic accumulation in D_{it} occurs at the social level, with D_{it} responding to the community's voting history, not the individual's. Increases in aggregate turnout may affect a community's information, attitudes, and norms about future voting. This class of mechanisms is consistent with the evidence of social influences on turnout (Nickerson 2008; Gerber et al. 2008; Bond et al. 2012; DellaVigna et al. 2013).

Aside from the terms of the Downsian model, actions by political elites may play a role in our results, especially if rain-induced decreases in turnout have a partisan bias. If rainfall affects election outcomes, and elected officials can manipulate voter turnout, then the persistent effects of rainfall shocks may have a political explanation. However, the majority of rainfall shocks are not large enough to change election outcomes. The 90th percentile of the rainfall distribution is 7.1 mm (Table 1), which given our estimates lowers turnout by approximately 0.37 percentage points. Most elections are won by substantially larger vote margins, especially in local races.[34] By this line of reasoning, the average effect of rainfall on who is elected (even in local races) is likely too small to plausibly explain the results. Even so, Table A3 estimates the effect of current and lagged rainfall on the Republican vote share in presidential elections. We find a small, marginally significant coefficient on contemporaneous rainfall suggesting that rain benefits Democrats, but it becomes insignificant with the inclusion of demographic covariates. The coefficient on lagged rainfall is insignificant with and without covariates. These weak results cast doubt on explanations involving the role of elected officials. At face value, they contradict the finding by Gomez et al. (2007) that rainfall benefits Republican candidates. However, Gomez et al. fail to account for clustering and omit county fixed effects and trends from their analyses.[35]

8 Conclusion

Social scientists have repeatedly documented that voting behavior is persistent, but they have struggled to isolate the mechanism driving this empirical regularity. This paper identifies the effects

[33] A possible psychological micro-foundation for this effect is cognitive dissonance: a citizen adapts her tastes regarding the importance of voting to create a consonance with her turnout decision.

[34] For example, only 9.3% of U.S. House of Representative elections in the 1948-1998 period had two-party vote share gaps smaller than 0.5 percentage points (Lee 2008).

[35] Hansford and Gomez (2010) use rainfall as an instrument for estimating the effect of turnout on the Democratic vote share. While they include county fixed effects (but not trends), they fail to account for clustering, which biases their standard errors downward; include multiple interactions of turnout with election characteristics, which obscures the average effect of rainfall; and exclude the South, which makes their sample incomparable to ours.

of habit formation, in which the act of voting today directly affects future turnout, as a causal channel for explaining turnout persistence. We use transitory and unexpected voting cost shocks due to Election-Day rainfall to estimate these effects, finding that a 1 percentage point decrease in current turnout reduces future turnout by 0.6-1.0 percentage points. Additional analyses suggest that this effect is unlikely to be driven by persistent changes in voting costs, by the updating of voter beliefs over the probability of being pivotal, or by changes in voters' perceived benefits from election outcomes. The weight of our evidence suggests that habit formation may occur through an accumulation in the consumption value, or expressive utility, citizens gain from voting.

45 years have passed since Riker and Ordeshook (1968) introduced the D_{it} term to the Downsian model as a solution to the paradox of voter turnout. Although many have accepted the idea that voters get consumption value from the act of voting, the precise form of this consumption value and the way it develops have remained elusive. By adding to the evidence base on habit formation in voting (Gerber et al. 2003; Meredith 2009), this paper speaks to the long-run effects of various turnout interventions that have been recently studied in the empirical political economy literature. Our finding should also further interest in the underlying psychological and social determinants of the consumption value voters gain from the act of voting and, as Feddersen (2004) suggests, in its implications for political economy models of strategic voter mobilization and suppression.

References

Agranov, M., J.K. Goeree, J. Romero, and L. Yariv. (2013). "What Makes Voters Turn Out: The Effects of Polls and Beliefs." Mimeo, Caltech.

Ali, S.N., and C. Lin. (2013). "Why People Vote: Ethical Motives and Social Incentives." *American Economic Journal: Microeconomics* 5(2): 73-98.

Atkin, D. (2013). "Trade, Tastes, and Nutrition in India." *American Economic Review* 103(5): 1629-63.

Atkinson, M.D., and A. Fowler. (2002). "Social Capital and Voter Turnout: Evidence from Saint's Day Fiestas in Mexico." *British Journal of Political Science* 46: 838-855.

Becker, G.S., and K.M. Murphy. (1988). "A Theory of Rational Addiction." *Journal of Political Economy* 96(4): 675-700.

Benabou, R. J., and J. Tirole. (2006). "Incentives and Prosocial Behavior." *American Economic Review* 96(5): 1652-1678.

Bendor, J., D. Diermeier, and M. Ting. (2003). "A Behavioral Model of Turnout." *American*

Political Science Review 97(2): 261-280.

Bond, R.M., C.J. Fariss, J.J. Jones, A.D. Kramer, C. Marlow, J.E. Settle, and J.H. Fowler. (2012). "A 61-Million-Person Experiment in Social Influence and Political Mobilization." *Nature* 489(7415): 295-298.

Brody, R.A., and P.M. Sniderman. (1977). "From Life Space to Polling Place: The Relevance of Personal Concerns for Voting Behavior." *British Journal of Political Science* 7(3): 337-360.

Campbell, J.E., S. Ali, and F. Jalalzai. (2006). "Forecasting the Presidential Vote in the States, 1948-2004." *Journal of Political Marketing* 5(1-2): 33-57.

Carroll, C.D., J. Overland, J., and D.N. Weil. (2000). "Saving and Growth with Habit Formation." *American Economic Review* 90(3): 341-355.

Case, A.C., and L.F. Katz. (1991). "The Company You Keep: The Effects of Family and Neighborhood on Disadvantaged Youths." National Bureau of Economic Research Working Paper 3705.

Charles, K.K., and M. Stephens Jr. (2013). "Employment, Wages, and Voter Turnout." *American Economic Journal: Applied Economics* 5(4): 111-143.

Collins, W.J., and R.A. Margo. (2007). "The Economic Aftermath of The 1960s Riots in American Cities: Evidence from Property Values." *Journal of Economic History* 67(4): 849-883.

Conley, T.G. (1999). "GMM Estimation with Cross Sectional Dependence." *Journal of Econometrics* 92(1): 1-45.

Constantinides, G.M. (1990). "Habit Formation: A Resolution of the Equity Premium Puzzle." *Journal of Political Economy* 98(3): 519-543.

Dawes, C.T., P.J. Loewen, and J. Fowler. (2011). "Social Preferences and Political Participation." *Journal of Politics* 73(3): 845-856.

Dell, M., B.F. Jones, and B.A. Olken. (2014). "What do We Learn from the Weather? The New Climate-Economy Literature." *Journal of Economic Literature* 52(3): 740-798.

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

DellaVigna, S., J. List, U. Malmendier, and G. Rao (2013). "Voting to Tell Others." Mimeo, U.C. Berkeley.

Denny, K., and O. Doyle. (2009). "Does Voting History Matter? Analysing Persistence in Turnout." *American Journal of Political Science* 53(1): 17-35.

Deschenes, O., and E. Moretti. (2009). "Extreme Weather Events, Mortality, and Migration." *Review of Economics and Statistics* 91(4): 659-681.

Downs, A. (1957). *An Economic Theory of Democracy.* New York: Harper Collins.

Drago, F., T. Nannicini, and F. Sobbrio. (2014). "Meet the Press: How Voters and Politicians Respond to Newspaper Entry and Exit." *American Economic Journal: Applied Economics* 6(3): 159-188.

Harsanyi, J.C. (1977). "Morality and the Theory of Rational Behavior." *Social Research* 44(4): 623-56.

Hoffman, M., J. Morgan, and C. Raymond. (2013). "One in a Million: A Field Experiment on Belief

Formation and Pivotal Voting." Mimeo, UC Berkeley.

Enikolopov, R., M. Petrova, and E. Zhuravskaya. (2011). "Media and Political Persuasion: Evidence from Russia." *American Economic Review* 101(7): 3253-3285.

Falck, O., R. Gold, and S. Heblich. (2014). "E-lections: Voting Behavior and the Internet." *American Economic Review* 104(7): 2238-65.

Feddersen, T.J. (2004). "Rational Choice Theory and the Paradox of not Voting." *Journal of Economic Perspectives* 18(1): 99-112.

Feddersen, T., and A. Sandroni. (2006). "A Theory of Participation in Elections." *American Economic Review* 96(4): 1271-1282.

Fowler, J. (2006). "Altruism and Turnout." *Journal of Politics* 68(3): 673-83.

Fowler, J., and C.D. Kam. (2007). "Beyond the Self: Social Identity, Altruism, and Political Participation." *Journal of Politics* 69(3): 813-827.

Fraga, B., and E. Hersh. (2010). "Voting Costs and Voter Turnout in Competitive Elections." *Quarterly Journal of Political Science* 5(4): 339-356.

Franklin, M.N., and S.B. Hobolt. (2011). "The Legacy of Lethargy: How Elections to the European Parliament Depress Turnout." *Electoral Studies* 30(1): 67-76.

Fuhrer, J.C. (2000). "Habit Formation in Consumption and its Implications for Monetary-Policy Models." *American Economic Review* 90(3): 367-390.

Gerber, A.S., D.P. Green, and C.W. Larimer. (2008). "Social Pressure and Vote Turnout: Evidence from a Large-Scale Field Experiment." *American Political Science Review* 102(1): 33-48.

Gerber, A.S., D.P. Green, and R. Shachar. (2003). "Voting May Be Habit-Forming: Evidence From a Randomized Field Experiment." *American Journal of Political Science* 47(3): 540-550.

Gentzkow, M. (2006). "Television and Voter Turnout." *Quarterly Journal of Economics* 121(3): 931-972.

Gentzkow, M., J.M. Shapiro, and M. Sinkinson. (2011). "The Effect of Newspaper Entry and Exit on Electoral Politics." *American Economic Review* 101(7): 2980-3018.

Glaeser, E. and J.A. Scheinkman. (2002). "Nonmarket Interactions." In *Advances in Economics and Econometrics: Theory and Applications: Eighth World Congress* (Vol. 1). Cambridge: Cambridge University Press.

Green, D.P., and R. Shachar. (2000). "Habit Formation and Political Behaviour: Evidence Of Consuetude in Voter Turnout." *British Journal of Political Science* 30(4): 561-573.

Gomez, B.T., T.G. Hansford, and G.A. Krause. (2007). "The Republicans Should Pray for Rain: Weather, Turnout, And Voting in U.S. Presidential Elections." *Journal of Politics* 69(3): 649-663.

Haines, M.R. (2010). *Historical, Demographic, Economic, and Social Data: The United States, 1790-2002*. Ann Arbor, MI: Inter-university Consortium for Political and Social Research.

Hansford, T.G., and B.T. Gomez. (2010). "Estimating the Electoral Effects of Voter Turnout." *American Political Science Review* 104(2): 268-288.

Huet-Vaughn, E. (2013). "Quiet Riot: The Causal Effect of Protest Violence." Mimeo, U.C. Berkeley.

Imbens, G., and J. Angrist. (1994). "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 61(2): 467-476.

Kanazawa, S. (1998). "A Possible Solution to the Paradox of Voter Turnout." *Journal of Politics* 60: 974-995.

Kaplan, E., and S. Mukand. (2011). "The Persistence of Political Partisanship: Evidence from 9/11." Mimeo, University of Maryland.

Lee, D. S. (2008). "Randomized Experiments from non-Random Selection in US House Elections." *Journal of Econometrics* 142(2): 675-697.

Lind, Jo Thori. (2015a). "Rainy Day politics: An Instrumental Variables Approach to the Effect of Parties on Political Outcomes." Mimeo, University of Oslo.

Lind, Jo Thori (2015b). "Spurious Weather Effects." CESifo Working Paper No. 5365.

Lodge, G., and S. Birch. (2012). "The Case for Compulsory Voting." *New Statesman.* Online: http://www.newstatesman.com/blogs/politics/2012/04/case-compulsory-voting.

Madestam, A., D. Shoag, S. Veuger, and D. Yanagizawa. (2013). "Do Political Protests Matter? Evidence from the Tea Party Movement." *Quarterly Journal of Economics* 128(4): 1633-1685.

Madestam, A., and D. Yanagizawa-Drott. (2012). "Shaping the Nation: The Effect of Fourth of July on Political Preferences and Behavior in the United States." Mimeo, Harvard University.

Meredith, M. (2009). "Persistence in Political Participation." *Quarterly Journal of Political Science* 4(3): 187-209.

Molloy, R., C.L. Smith, and A. Wozniak. (2011). "Internal Migration in the United States." *Journal of Economic Perspectives* 25(3): 173-96.

Mullainathan, S., and E. Washington. (2009). "Sticking with Your Vote: Cognitive Dissonance and Political Attitudes." *American Economic Journal: Applied Economics* 1(1): 86-111.

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

Pollak, R.A. (1970). "Habit Formation and Dynamic Demand Functions. *Journal of Political Economy* 78(4): 745-763.

Riker, W.H., and P.C. Ordeshook. (1968). "A Theory of the Calculus of Voting." *American Political Science Review* 62(1): 25-42.

Strömberg, D. (2004). "Radio's Impact on Public Spending." *Quarterly Journal of Economics* 119(1): 189-221.

Tullock, G. (1967). *Toward a Mathematics of Politics.* Ann Arbor: University of Michigan Press.

Figure 1: Age Patterns in Voting, CPS Voter Supplement 1980-2010

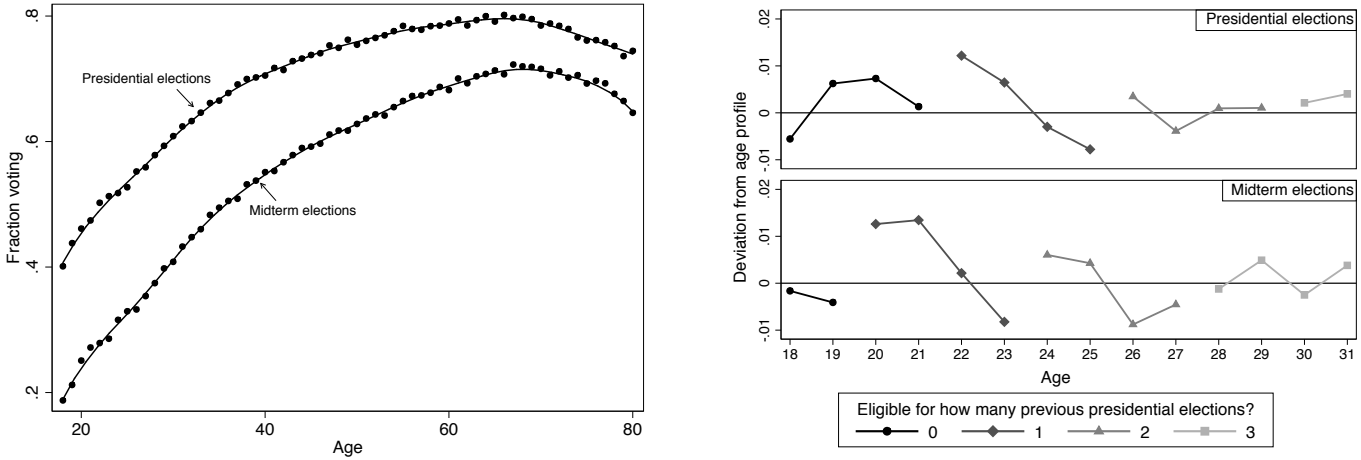

Note: Age profile estimated by local linear regression with a bandwidth of 2 years.

Figure 2: Coefficients on Lagged and Contemporaneous Rainfall, Election Day and Nearby Days

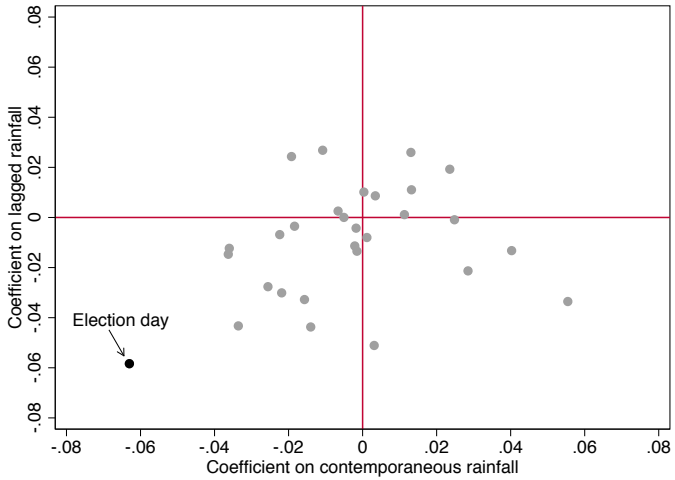

Note: Plot of α_1 on α_0 from regressions of the form:

$$turnout_{ct} = \alpha_0\, other_day_rain_{ct} + \alpha_1\, other_day_rain_{c,t\text{-}1} + \beta_0\, election_day_rain_{ct} + \beta_1\, election_day_rain_{c,t\text{-}1} + \tau_t + \eta_c + f_c(t) + \varepsilon_{ct}$$

where $other_day_rain_{ct}$ is rainfall on a day within 2 weeks of Election Day, and $f_c(t)$ is a county-specific quadratic time trend. Estimates of (α_1,α_0) appear in grey; the estimate of (β_1,β_0) from equation (4) appears in black.

Table 1: Descriptive Statistics

			Percentiles				
	Mean	Std. Dev.	10th	25th	50th	75th	90th
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Politics							
Voter turnout	58.4	13.6	41.8	49.4	58.3	67.4	75.8
Republican vote share	55.3	14.2	36.6	46.3	56.1	65.2	72.9
Weather							
Rainfall on Election Day (mm)	2.4	6.6	0.0	0.0	0.024	1.4	7.0
Demographics							
% white	87.8	15.8	64.2	82.3	95.0	98.6	99.7
% over 65	13.2	4.4	7.8	10.1	12.8	15.9	19.0
Log median household income (2012 $)	10.6	0.3	10.2	10.5	10.7	10.8	11.0
Log population density (people/sq. mile)	3.6	1.6	1.5	2.8	3.6	4.5	5.6

Note: The sample includes 49,594 county-year observations, based on presidential elections from 1952-2012 in 3,108 counties.

Table 2: Effect of Contemporaneous Rainfall on Turnout, Various Specifications

	(1)	(2)	(3)	(4)	(5)
A. Baseline specification					
Election-Day rain	-0.165	-0.013	-0.070	-0.054	-0.055
	[0.043]***	[0.031]	[0.014]***	[0.021]**	[0.022]**
B. Controlling for 2-week avg. rainfall					
Election-Day rain	-0.082	0.014	-0.055	-0.052	-0.052
	[0.038]**	[0.031]	[0.023]**	[0.020]**	[0.020]**
Average rain from 7 days before to 7 days after Election Day	-0.588	-0.296	-0.161	-0.026	-0.027
	[0.279]**	[0.122]**	[0.081]*	[0.071]	[0.071]
C. Controlling for 2-week avg. rainfall and demographic covariates					
Election-Day rain	0.041	0.014	-0.049	-0.049	-0.050
	[0.035]	[0.026]	[0.022]**	[0.019]**	[0.019]**
Average rain from 7 days before to 7 days after Election Day	-0.492	-0.277	-0.138	-0.022	-0.022
	[0.191]**	[0.099]***	[0.075]*	[0.071]	[0.071]
Number of county-years	49,594	49,594	49,594	49,594	49,594
Number of counties	3108	3108	3108	3108	3108
Election years	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012
County and year fixed effects		✓	✓	✓	✓
County-specific linear trends			✓		
County-specific quadratic trends				✓	
County-specific cubic trends					✓

Note: Dependent variable is voter turnout (0-100). Brackets contain standard errors clustered at the state level. County covariates are the white population share, the over-65 population share, log median income, and log population density. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effect of Contemporaneous and Lagged Rainfall on Turnout

	Full sample				Short sample			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Election-Day rain., t+1			-0.001	-0.002			-0.014	-0.008
			[0.021]	[0.021]			[0.028]	[0.028]
Election-Day rain., t	-0.063	-0.060	-0.061	-0.061	-0.052	-0.050	-0.061	-0.058
	[0.023]***	[0.022]***	[0.021]***	[0.022]***	[0.022]**	[0.021]**	[0.026]**	[0.025]**
Election-Day rain., t-1	-0.058	-0.058			-0.033	-0.034	-0.043	-0.043
	[0.021]***	[0.020]**			[0.017]*	[0.016]**	[0.027]	[0.026]
Election-Day rain., t-2							-0.024	-0.026
							[0.020]	[0.019]
ρ	0.93	0.95	0.96	0.93	0.63	0.68	0.63	0.68
	[0.33]***	[0.34]***	[0.36]***	[0.35]***	[0.30]**	[0.31]**	[0.27]**	[0.27]**
Estimation method for ρ	Delta	Delta	2SLS	2SLS	Delta	Delta	Delta	Delta
Number of county-years	49,594	49,594	49,524	49,524	43,400	43,400	43,400	43,400
Number of counties	3,108	3,108	3,108	3,108	3,108	3,108	3,108	3,108
Election years	1952-2012	1952-2012	1948-2008	1948-2008	1956-2008	1956-2008	1956-2008	1956-2008
County covariates		✓		✓		✓		✓

Note: Dependent variable is voter turnout (0-100). Brackets contain standard errors clustered at the state level. All regressions include year fixed effects, county fixed effects, and county-specific quadratic trends. County covariates are the white population share, the over-65 population share, log median income, and log population density. * p<0.1, ** p<0.05, *** p<0.01

Table 4: Interactions with County Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
Election-Day rain, t	-0.063	-0.253	-0.074	-1.34	-0.153	-1.17
	[0.023]***	[0.094]**	[0.042]*	[0.52]**	[0.039]***	[0.57]**
Election-Day rain, t-1	-0.058	-0.136	-0.084	-1.62	-0.162	-1.41
	[0.021]***	[0.097]	[0.034]**	[0.47]***	[0.041]***	[0.46]***
(% white) × (rain, t)		0.0024				0.0017
		[0.0011]**				[0.0010]
(% white) × (rain, t-1)		0.0010				0.0005
		[0.0010]				[0.0008]
(% over 65) × (rain, t)			0.0011			0.0035
			[0.0023]			[0.0025]
(% over 65) × (rain, t-1)			-0.0021			0.0026
			[0.0017]			[0.0017]
(Log median income) × (rain, t)				0.120		0.083
				[0.049]**		[0.059]
(Log median income) × (rain, t-1)				0.147		0.116
				[0.043]***		[0.043]***
(Log population density) × (rain, t)					0.021	0.014
					[0.006]***	[0.006]**
(Log population density) × (rain, t-1)					0.024	0.013
					[0.007]***	[0.006]**
Number of county-years	49,594	49,594	49,594	49,594	49,594	49,594
Number of counties	3,108	3,108	3,108	3,108	3,108	3,108
Election years	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012

Note: Dependent variable is voter turnout (0-100). Sample includes presidential elections from 1952-2012. Brackets contain standard errors clustered at the state level. All regressions include year fixed effects, county fixed effects, county-specific quadratic trends, and the main effects of any demographic variables included in the interaction terms. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Midterm Elections

	(1)	(2)	(3)	(4)
Election-Day rain, t	0.013	0.013	0.020	0.020
	[0.029]	[0.029]	[0.028]	[0.028]
Election-Day rain, t-1	0.021	0.022	0.027	0.028
	[0.040]	[0.039]	[0.037]	[0.036]
Uncontested, t	-10.40	-10.44	-10.13	-10.18
	[0.74]***	[0.72]***	[0.75]***	[0.72]***
Uncontested, t-1	-1.32	-1.35	-1.04	-1.08
	[0.51]**	[0.50]***	[0.50]**	[0.49]***
(Uncontested, t) × (Election-Day rain, t)			-0.080	-0.077
			[0.034]**	[0.034]**
(Uncontested, t-1) × (Election-Day rain, t-1)			-0.084	-0.082
			[0.064]	[0.065]
Number of county-years	45,420	45,420	45,420	45,420
Number of counties	3,107	3,107	3,107	3,107
Election years	1954-2010	1954-2010	1954-2010	1954-2010
County covariates		✓		✓

Note: Dependent variable is voter turnout (0-100). Brackets contain standard errors clustered at the state level. Sample excludes Washington, DC, which has no congressional elections. All regressions include year fixed effects, county fixed effects, and county-specific quadratic trends, as well as controlling for whether (Uncontested, t-1) is missing. Columns (3)-(4) also control for the interaction of lagged rainfall with the missing indicator. County covariates are the white population share, the over-65 population share, log median income, and log population density. * p<0.1, ** p<0.05, *** p<0.01

Table 6: Persistence in Voter Turnout, Individuals versus Counties

	Individuals (CPS 1972-1984)			Counties			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Turnout in last election	0.48	0.51	0.46	0.83	0.82	0.78	0.69
	[0.01]***	[0.01]***	[0.01]***	[0.01]***	[0.02]***	[0.03]***	[0.04]***
Number of observations	315,970	315,970	315,970	49,524	49,524	49,524	49,524
Election years	1972-1984	1972-1984	1972-1984	1952-2012	1952-2012	1952-2012	1952-2012
Covariates		✓	✓		✓	✓	✓
Year FE		✓	✓		✓	✓	✓
State-group/state FE			✓			✓	
County FE							✓

Note: Dependent variable is voter turnout (0-100). * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Individual analysis: Brackets contain SEs clustered at the state-group level. The analysis uses 23 state-groups instead of 50 states because the 1976 CPS does not contain state identifiers. In the subsample from other years, results were identical in estimations with clustering or fixed effects at the state, rather than state-group, level. Covariates include education level, age, age squared, gender, and race.

County analysis: Brackets contain SEs clustered at the state level. Covariates include the white population share, the over-65 population share, log median income, and log population density.

ONLINE APPENDIX – NOT FOR PUBLICATION

Figure A1: Share of Counties with Election-Day Rainfall by Year

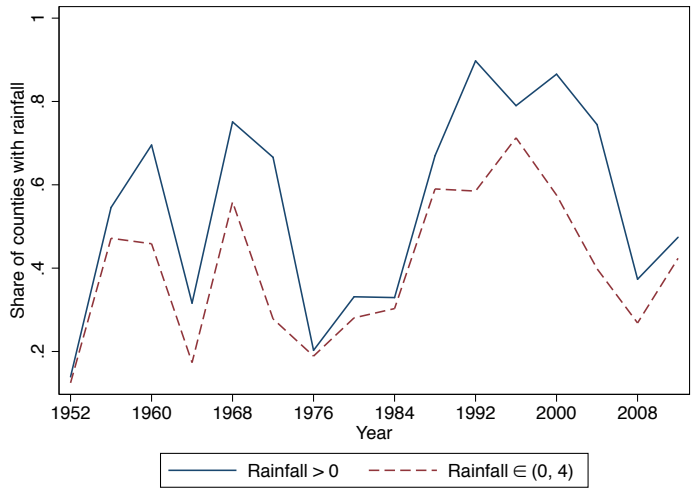

Figure A2: Cumulative Share of Counties with Election-Day Rainfall

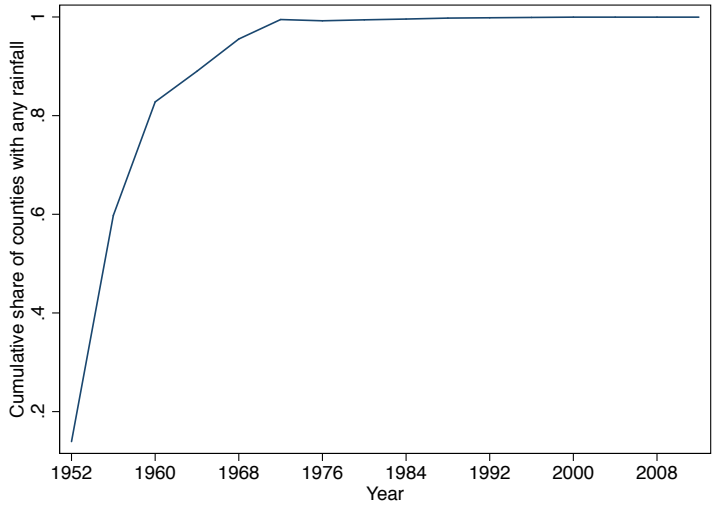

Figure A3: Histogram of Standard Deviation of Rainfall

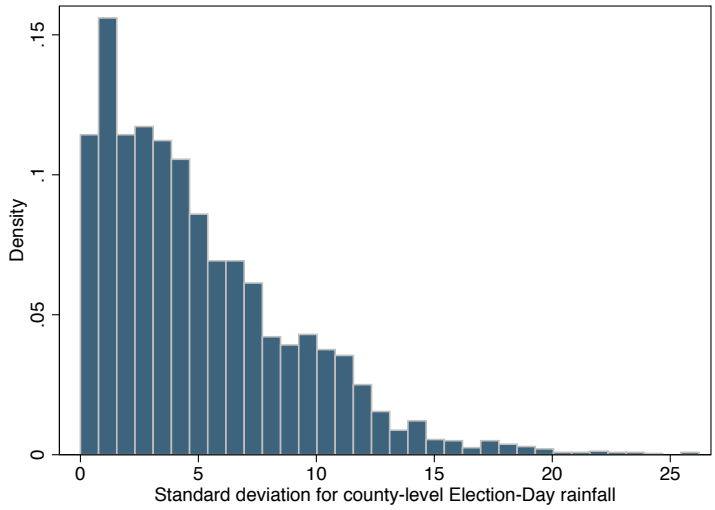

Figure A4: County-Level Trends in Turnout and Election-Day Rainfall

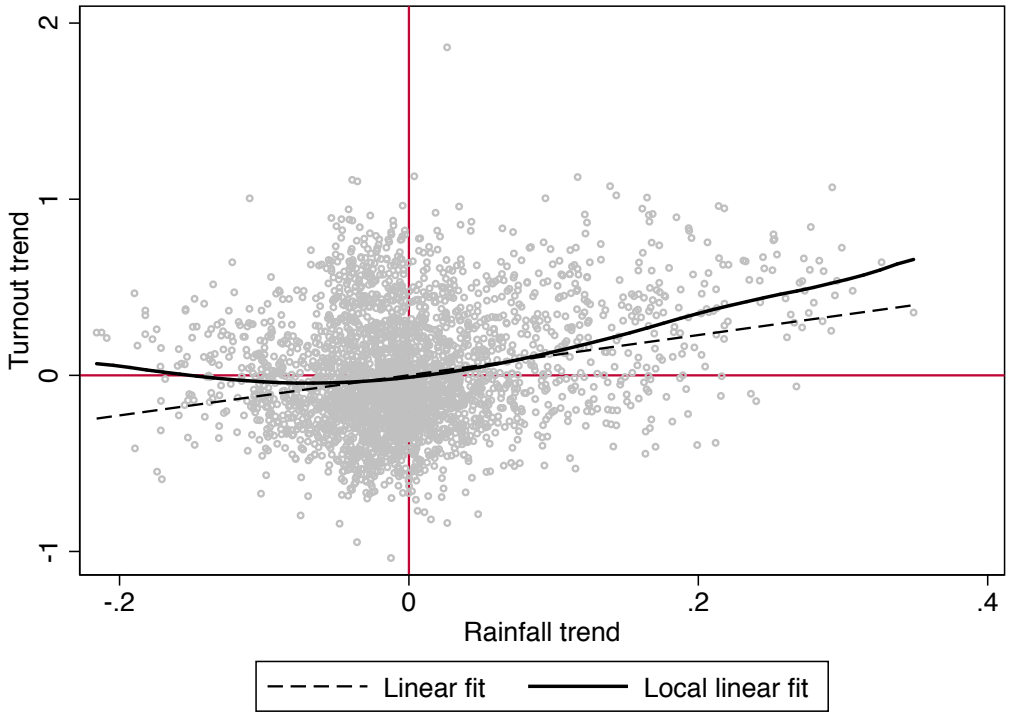

Note: After purging Election-Day rainfall and turnout of county and year effects, we estimated county-specific linear trends in these variables. Each dot corresponds to the coefficient from a regression of the trend in turnout on the trend in Election-Day rainfall. The local linear regression has a bandwidth of 0.1.

Figure A5: Associations of Trends in Turnout and Trends in Rainfall on Alternative Days

Panel A: Days Relative to Election Day

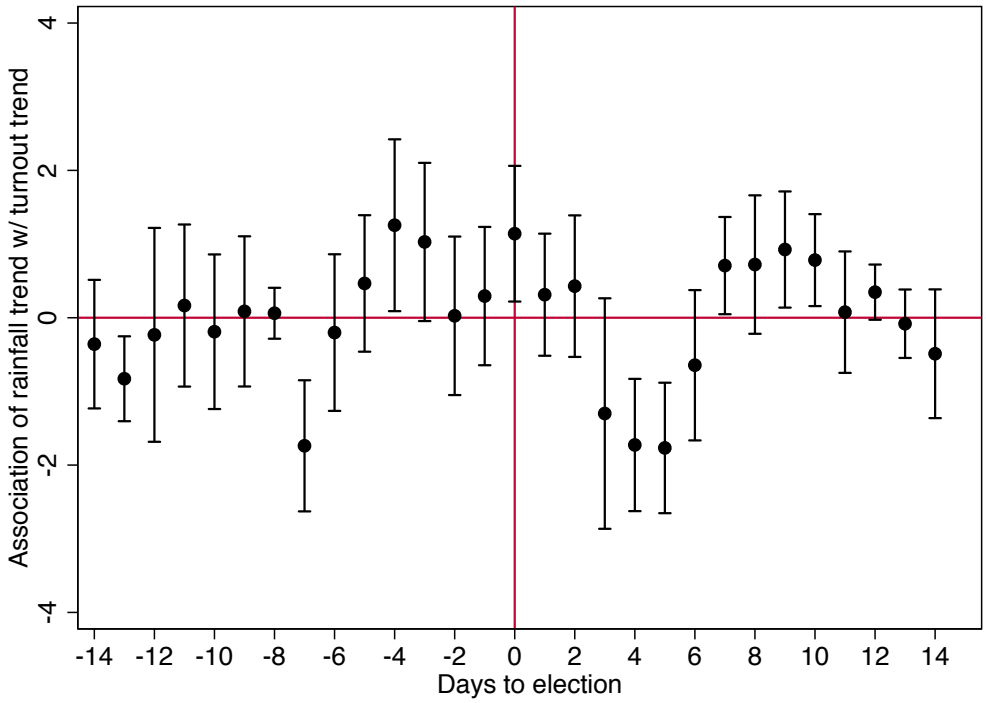

Panel B: Calendar Days

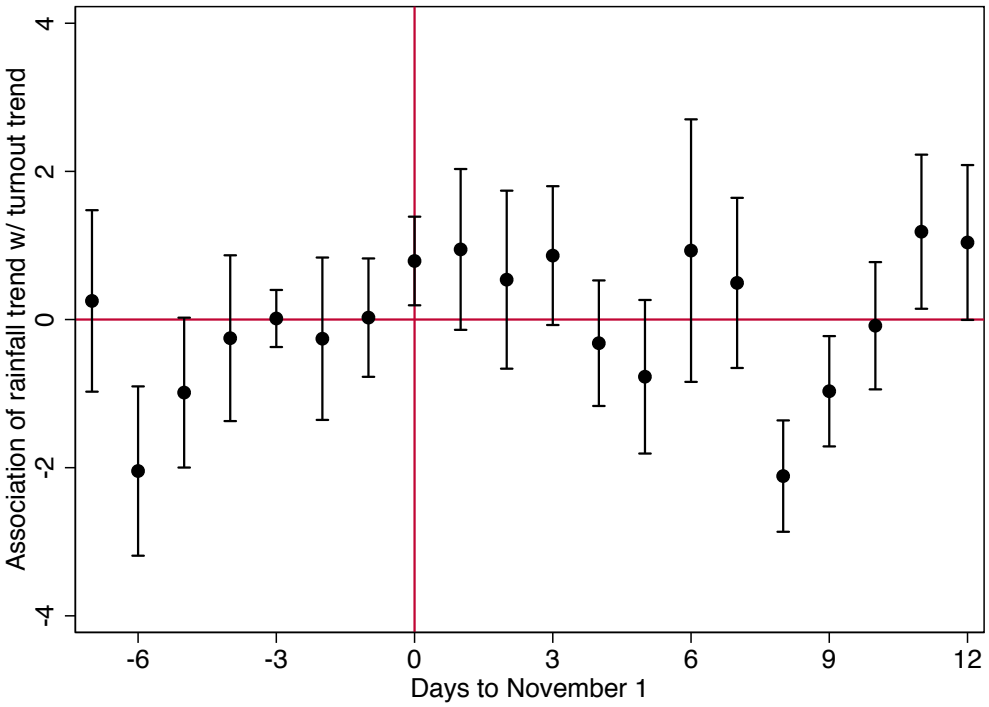

Note: After purging daily rainfall and turnout of county and year effects, we estimated county-specific linear trends in these variables. Each dot corresponds to the coefficient from a regression of the trend in turnout on the trend in rainfall on the specified day. Capped spikes are 95% CIs.

Figure A6: Rainfall and Turnout Residuals, 2004

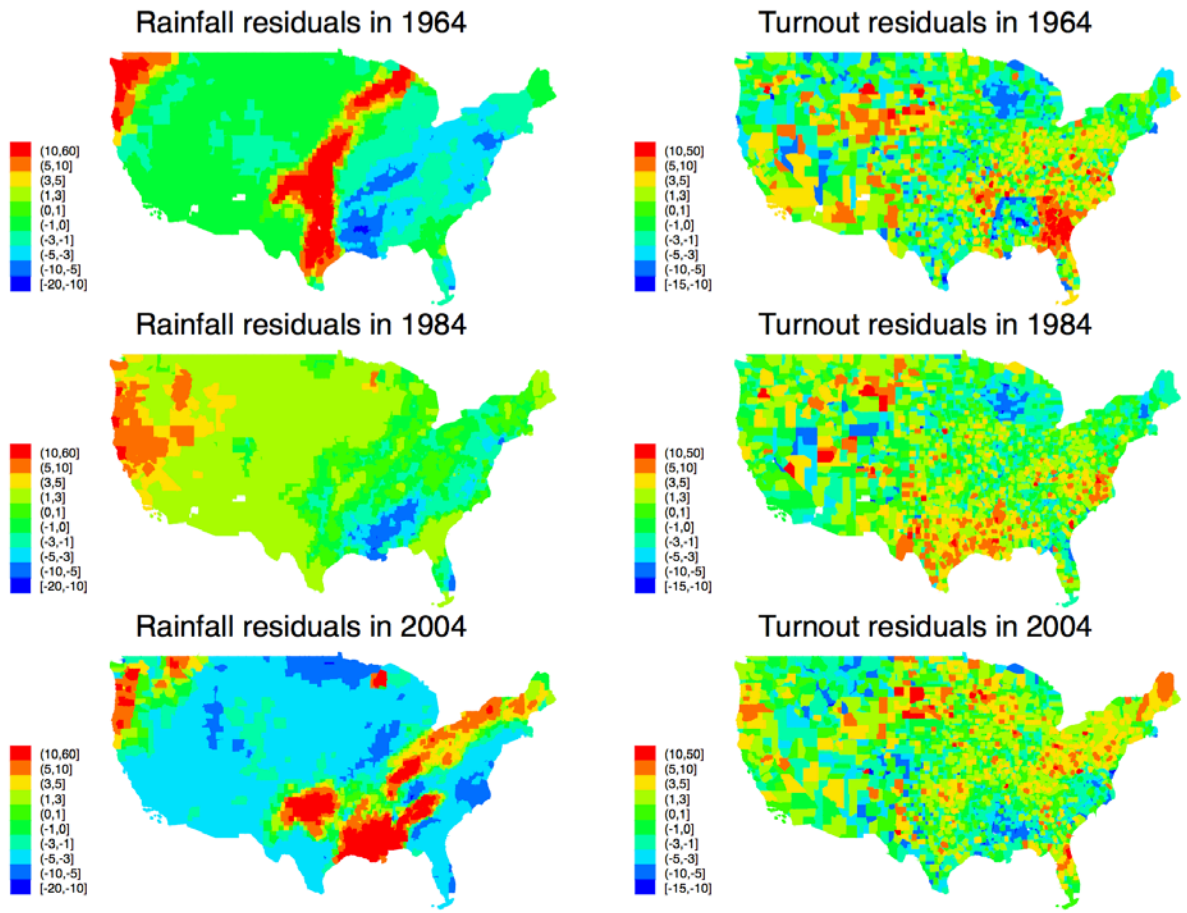

Note: Residuals from regressions of rainfall (mm) and turnout on year and county fixed effects and county trends.

Figure A7: Effects of Rainfall on Election Day and Nearby Days, Different Specifications

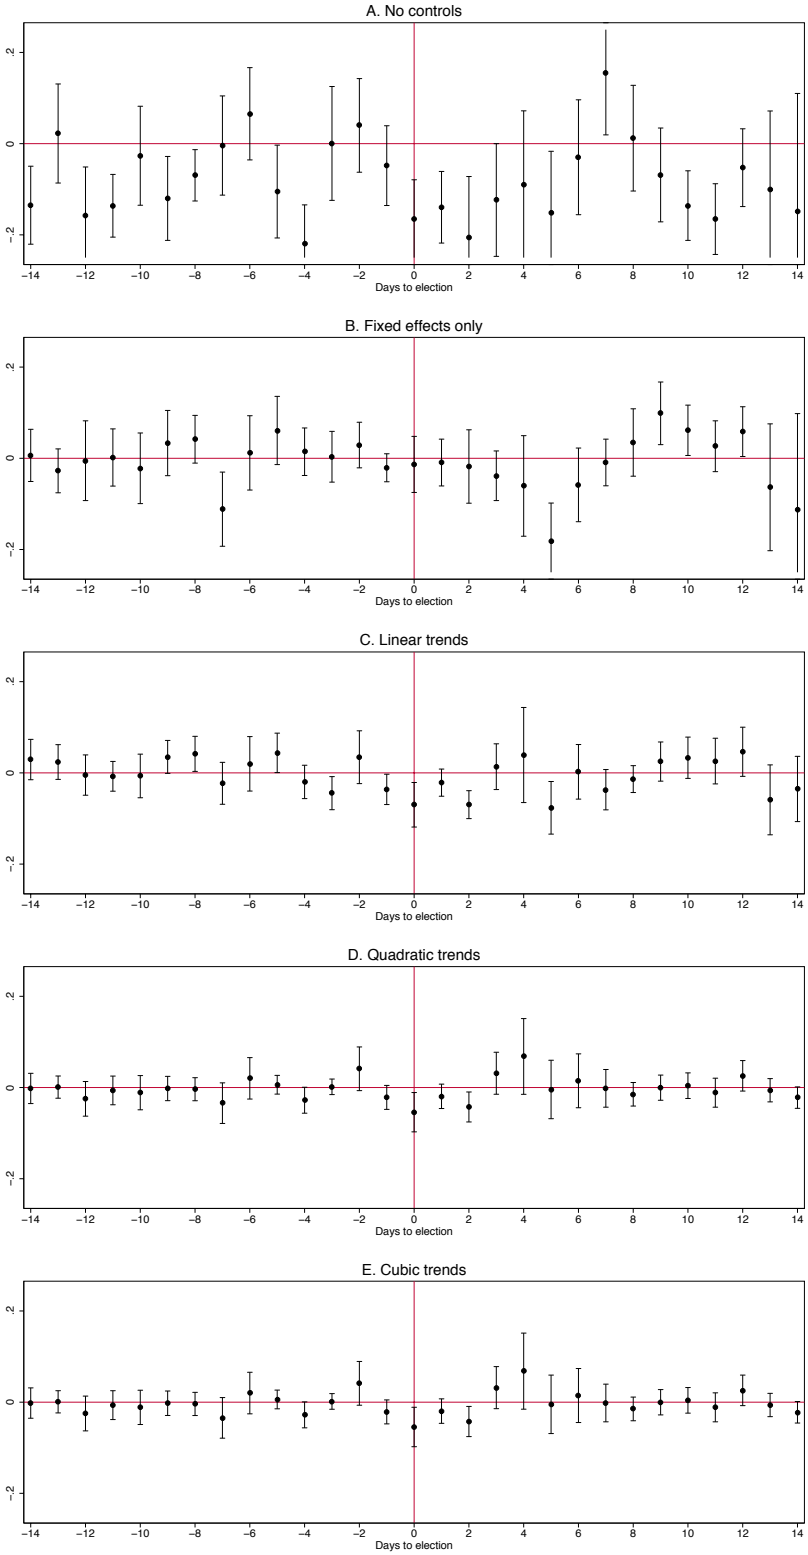

Note: Plot of α from regression: $turnout_{ct}$ = constant + α $other_day_rain_{ct}$ + β $election_day_rain_{ct}$ + e_{ct}, where e_{ct} may contain year/county fixed effects or county trends. α estimated separately for each placebo day. Capped spikes are 95% CIs. The absence of a cap indicates that the CI extends beyond the range of the y-axis.

Figure A8: Leave-One-Out Checks

Panel A: Leave Out One State

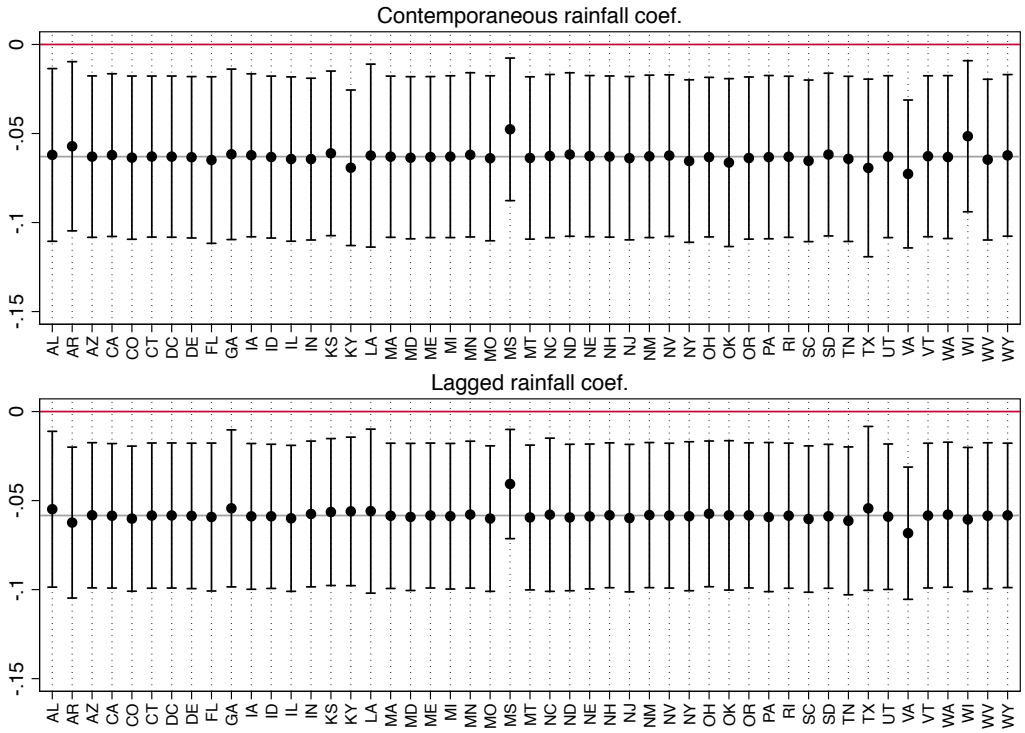

Panel B: Leave Out One Year

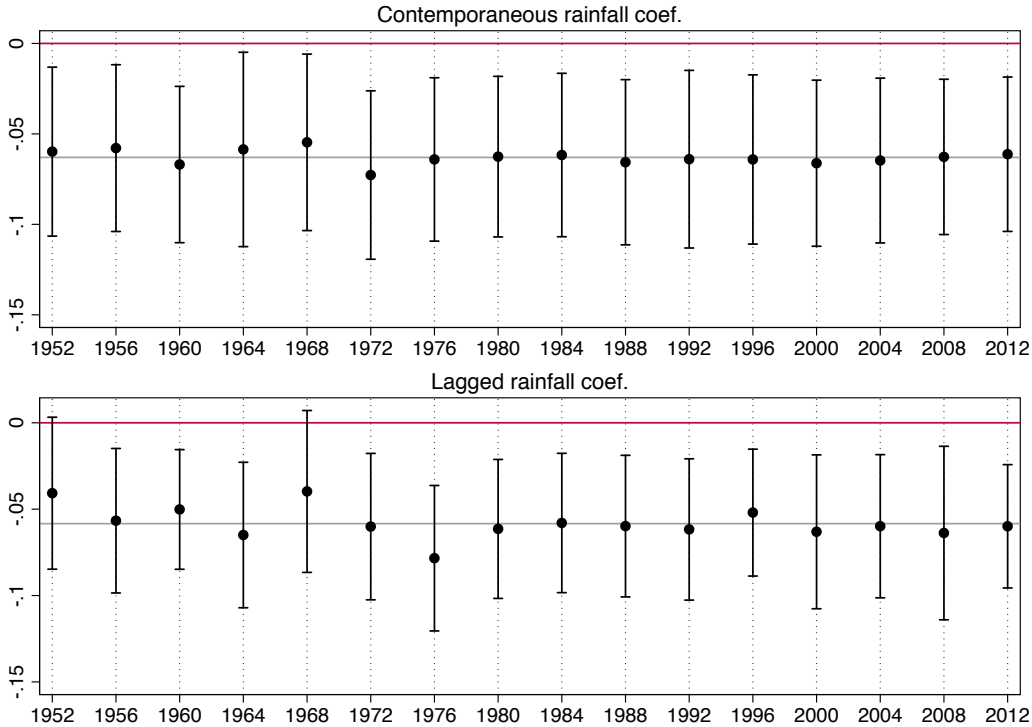

Note: Each estimate is based on a sample that omits the state or year on the x-axis. Dots are coefficients; capped spikes are 95% CIs. Light gray horizontal lines represent full-sample estimates.

Figure A9: Rolling Window Estimates

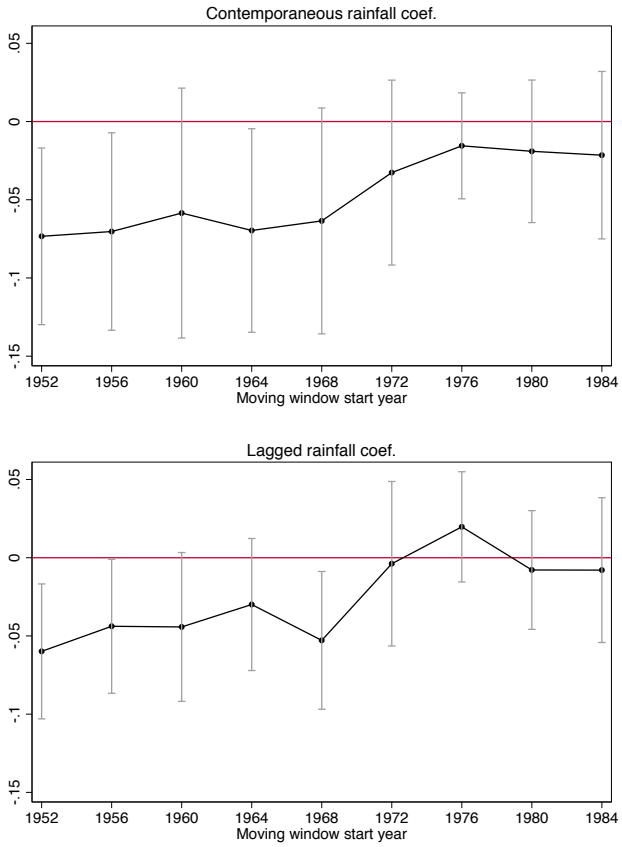

Note: Each estimate is based on a sample with eight elections starting in the specified year. Dots are coefficients; capped spikes are 95% CIs.

Figure A10: Additional Leads and Lags

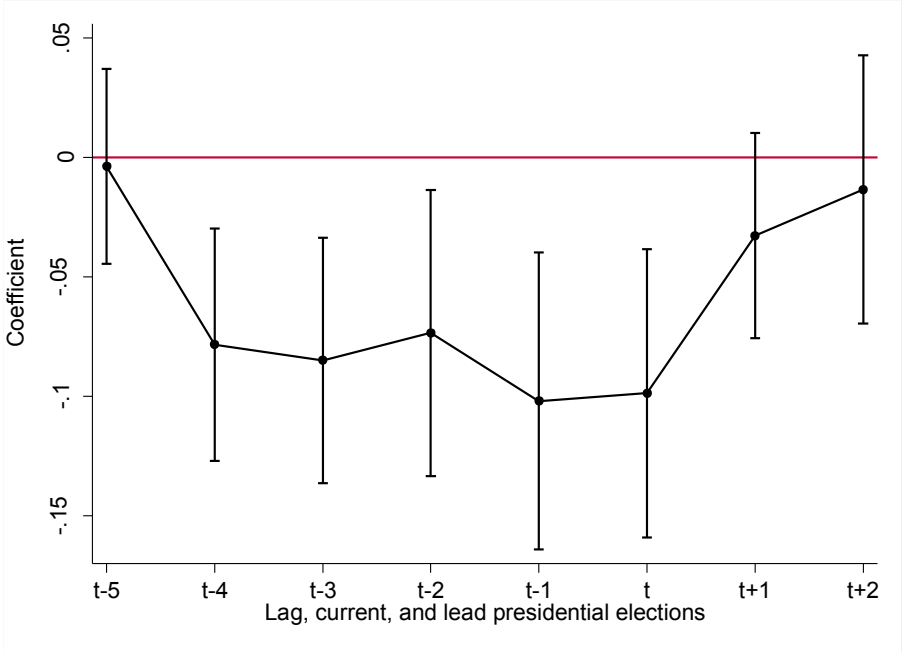

Note: Coefficients and 95% CIs from a model jointly estimating election-day rainfall from period t-5 to t+2. Rainfall effects are modeled linearly. Model includes year fixed effects, county fixed effects, and county quadratic trends.

Figure A11: Checking nonlinearity of response function

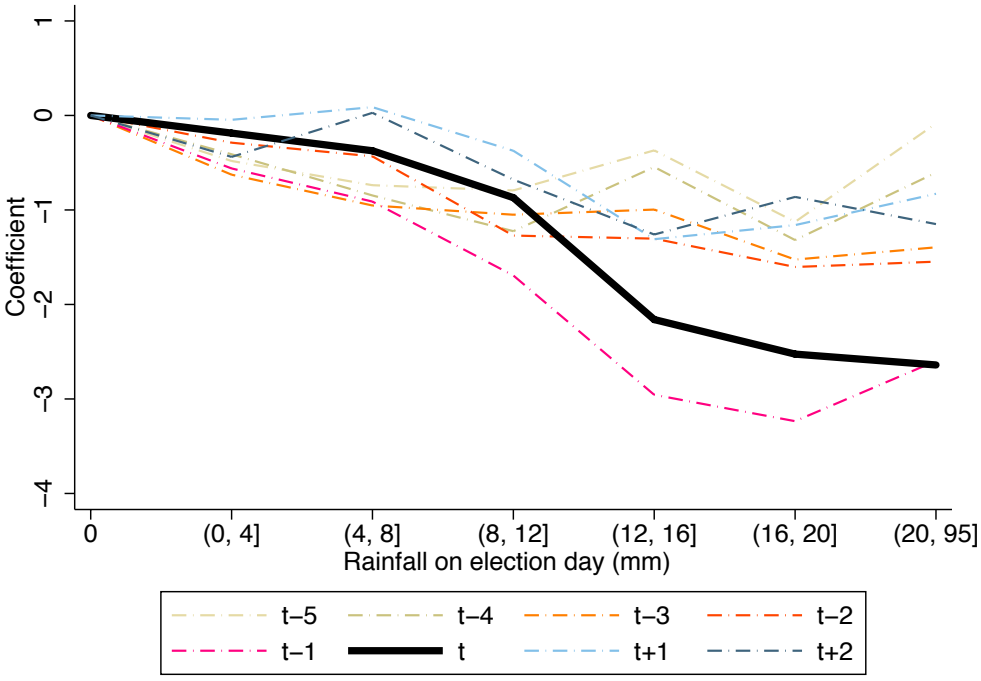

Note: Coefficients from a model jointly estimating election-day rainfall from period t-5 to t+2. Rainfall effects are modeled nonlinearly using discrete bins with dry election days as the omitted category. Model includes year fixed effects, county fixed effects, and county quadratic trends.

Table A1: Effect of Contemporaneous and Lagged Rainfall on Turnout – Alternative Specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Election-Day rain, t	-0.012	-0.079	-0.063	-0.063	-0.071	-0.054	-0.054	-0.035	-0.063
	[0.031]	[0.026]***	[0.023]***	[0.023]***	[0.045]	[0.019]***	[0.021]**	[0.017]**	[0.025]**
Election-Day rain, t-1	0.016	-0.070	-0.058	-0.059	-0.064	-0.053	-0.053	-0.040	-0.063
	[0.021]	[0.026]***	[0.021]***	[0.021]***	[0.040]	[0.017]***	[0.020]***	[0.014]***	[0.021]***
ρ	-1.33	0.89	0.93	0.92	0.90	0.99	0.98	1.10	1.00
	[4.38]	[0.28]***	[0.33]***	[0.32]***	[0.45]**	[0.38]**	[0.35]***	[0.57]*	[0.35]***
Number of county-years	49,594	49,594	49,594	49,594	49,594	49,524	49,524	49,524	49,524
Number of counties	3,108	3,108	3,108	3,108	3,108	3,108	3,108	3,108	3,108
Election years	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012	1952-2012
County and year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
County linear trends		✓				✓	✓	✓	✓
County quadratic trends			✓						
County cubic trends				✓					
Decade-county FE					✓				
Year fixed effects interacted with:						Log median income	Over-65 pop. share	White pop. share	Pop. Density

Note: Dependent variable is voter turnout (0-100). Brackets contain standard errors clustered at the state level. ρ is estimated using the delta method. The variables interacted with year fixed effects are for the first period of the sample (1952). * p<0.1, ** p<0.05, *** p<0.01

Table A2: Interactions with Electoral Characteristics

	Interaction with…				
	Alignment w/ winner, t-1	State pivot prob., t	Nat'l vote margin, t-1	Republican incumbent, t	Incumbent running, t
	(1)	(2)	(3)	(4)	(5)
Election-Day rain, t	-0.058	-0.061	-0.047	-0.077	-0.033
	[0.023]**	[0.029]**	[0.040]	[0.024]***	[0.042]
Election-Day rain, t-1	-0.058	-0.047	-0.067	-0.088	-0.107
	[0.021]***	[0.018]**	[0.014]***	[0.031]***	[0.012]***
(Variable) × (rain, t)	0.0012	-50	-0.018	0.018	0.041
	[0.0008]	[157]	[0.037]	[0.034]	[0.039]
(Variable) × (rain, t-1)	-0.0005	-105	-0.037	0.060	0.078
	[0.0007]	[96]	[0.031]	[0.040]	[0.030]***
Number of county-years	49,393	42,944	49,524	49,524	49,524
Number of counties	3,108	3,108	3,108	3,108	3,108
Election years	1952-2012	1952-2004	1952-2012	1952-2012	1952-2012

Note: Dependent variable is voter turnout (0-100). Sample includes presidential elections from 1952-2012. Brackets contain standard errors clustered at the state level. All regressions include year and county fixed effects, county-specific quadratic trends, and the main effects of any variables included in the interaction terms. Column (1) adds interactions with a measure of whether the county is aligned with winning candidate of the presidential election. To avoid endogeneity, we use a county's Republican vote share two elections ago to ascertain its partisan leaning. *Alignment with winner, t-1* is equal to the county's Republican vote share in *t-2* minus 50 if a Republican won the national election in *t-1*, and is equal to 50 minus the county's Republican vote share in *t-2* if a Democrat won in *t-1*. Column (2) adds interactions with a measure of predicted pivotalness. We use Campbell et al.'s (2006) model to calculate a predicted Democratic vote share, d_{st}, for each state s and election year t. The probability of a randomly drawn voter breaking a state-level tie is $(1/N_{st})\varphi(d_{st} - 0.5/\sigma_{st})$, where $\varphi(\cdot)$ is the standard normal density function, σ_{st} is the standard deviation of d_{st}, and N_{st} is the number of registered voters. Our conclusions do not change if we use predicted closeness rather than predicted pivotalness. The point estimates and standard errors for both the interacted pivotal coefficients are large because the probability of being pivotal is typically on the order of 10^{-4} percent. Column (3) adds interactions with the absolute value of the national vote share difference between the Republican and Democratic presidential candidates. Columns (4) adds interactions with an indicator for whether the incumbent President is a Republican, and column (5) adds interactions with an indicator for whether the incumbent President is running for re-election. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table A3: Effect of Contemporaneous and Lagged Rainfall on the Republican Vote Share

	(1)	(2)
Election-Day rain, t	-0.048	-0.042
	[0.028]*	[0.027]
Election-Day rain, t-1	0.048	-0.041
	[0.033]	[0.031]
Number of county-years	49,511	49,511
Number of counties	3,108	3,108
Election years	1952-2012	1952-2012
County covariates		✓

Note: Dependent variable is voter turnout (0-100). Brackets contain standard errors clustered at the state level. All regressions include year fixed effects, county fixed effects, and county-specific quadratic trends. County covariates are the white population share, the over-65 population share, log median income, and log population density. * p<0.1, ** p<0.05, *** p<0.01