Estimating Habit Formation in Voting*

Thomas Fujiwara† Kyle Meng‡ Tom Vogl§

November 2013

Abstract

We estimate the effect of past on current voting by exploiting transitory voting cost shocks. Using county-level data on U.S. presidential elections from 1952-2012, we find that precipitation on current and past election days reduces voter turnout. Our estimates imply that a 1 point decrease in past turnout lowers current turnout by 0.7-0.9 points. Consistent with a dynamic Downsian framework, current precipitation has stronger effects following previous inclement elections. We find no effect of precipitation on nearby days or future election days, suggesting our result reflects the perpetuation of the effects of voting costs though habit formation.

*We thank Doug Miller and seminar participants at UC Davis for comments; James Campbell, Wolfram Schlenker, and James Snyder for sharing data; and Sarah Weltman for excellent research assistance.

†Princeton University, BREAD, CIFAR, and NBER.

‡Princeton University and UCSB.

§Princeton University, BREAD, and NBER.

1 Introduction

Voter turnout persists over time: a citizen who votes today is more likely to vote in the future (Brody and Sniderman 1977). Although this pattern is widespread, it is poorly understood. On the one hand, an individual may repeat her decision to turn out or abstain because her benefits and costs of voting are stable over time. On the other, persistence in turnout may be evidence of habit formation. When voting is habit forming, the mere act of voting today increases the likelihood of future voting, holding constant voters' characteristics. The distinction between habit formation and turnout persistence in general is critical for the design of policies aimed at boosting political participation. The presence of habit formation would alter the cost-benefit analysis of voter participation programs such as get-out-the-vote campaigns, mandatory voting, paid election days, and improved access to polls. Evidence of habit formation as an important source of turnout persistence would also inform the relative impacts such programs have on citizens of different ages.[1]

Despite these important consequences, efforts to empirically disentangle habit formation from other channels of persistence have met difficulty, for at least two reasons. First, the task requires a source of variation in turnout that is uncorrelated with other baseline determinants of the voting decision. Second, even if such variation exists, it can only identify habit formation if it has no direct effect on future turnout. We address this empirical challenge by exploiting transitory and unexpected shocks to voting costs due to precipitation on election day. Following previous studies documenting that rain decreases turnout (Gomez et al. 2007; Hansford and Gomez 2010; Fraga and Hersh 2011), our test for habit formation amounts to asking whether election-day precipitation decreases voter turnout not only in the current election but also during future elections. Because precipitation on election day affects future turnout *only* through its impact on current voting, we are able to isolate the effect of habit formation.

To ground the analysis conceptually, we propose a theoretical framework for habit formation based on a dynamic extention of the Downsian "calculus of voting" model. In the framework, precipitation is a transitory shock to the cost of voting, propagated through time because the current turnout decision affects the utility of voting in future elections. We use the framework to clarify

[1] 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."

what is required to identify habit formation, and we discuss why election-day precipitation fits such requirements, not only because it is orthogonal to voters' characteristics, but also because it is transitory (affecting current but not future voting costs) and unexpected (not leading voters and other agents to adapt their behavior prior to election day).

Matching daily weather data with county-level U.S. presidential election returns from 1952 to 2012, we find that both current and lagged election-day precipitation reduce voter turnout. Our models include year fixed effects, county fixed effects, and county-specific trends, allaying concerns about unobserved heterogeneity or confounding trends. We provide a series of robustness and falsification tests in support to our interpretation that the effect of lagged rainfall is due to habit formation. First, turnout shows no relation to precipitation on the day of the next presidential election. Second, turnout shows no relation to precipitation two weeks after the current election day or the previous election day. These results strengthen our claim that only precipitation that fell precisely on the previous election day matters. Third, the effects of both current and lagged precipitation is stronger in rural areas, where the costs of inclement weather are greater.

We further explore the mechanisms driving habit formation in voting decisions, testing an additional prediction from our theoretical framework: that past shocks to the cost of voting increase the sensitivity of voter turnout to current cost shocks. Indeed, our results show that precipitation has a greater effect on current turnout if the previous election also involved precipitation. Drawing on several additional analyses, we also argue that our main result is unlikely to be driven by persistent changes in voting costs, 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 past voting experiences reinforce the expressive utility citizens gain from the act of voting: that is to say, the consumption benefit of voting.

By documenting habit formation in voting, the paper relates to three strands in the literature. First, a large number of studies document that individuals who voted in the past are more likely to vote in the future, suggesting that this association reflects habit formation (e.g., Brody and Sniderman 1977; Plutzer 2002). However, these studies typically do not attempt to disentagle habit formation from persistence in the costs and benefits of voting.[2] Notable exceptions include Gerber et

[2] Two papers address this issue using instrumental variables methods that rely on uncertain identifying assumptions. Using a panel of voters, 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)

al. (2003) and Meredith (2009), who 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. We discuss these results in the context of our conceptual framework, highlighting that they require additional assumptions—which do not necessarily hold—to be interpreted as evidence of habit formation.[3] Specifially, both treatments could plausibly affect future turnout even if an individual did not vote in the current election. We explain why the assumptions are more natural in the case of a transitory and unexpected shock in the cost of voting, such as precipitation.

Second, the paper speaks to the literature on the formation of political preferences, especially research establishing how actions feed back into tastes and beliefs.[4] Also exploiting daily precipitation, Madestam et al. (2013) and Madestam and Yanagizawa-Drott (2011) use rainfall on Tax Day and Independence Day to estimate the effect of participating in Tea Party protests and independence day celebrations on political preferences and behavior.[5] Other shocks also have persistent effects; for example, Kaplan and Mukand (2011) show 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 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.

Third, and perhaps more importantly, 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. They highlight that their model predicts substantial equilibrium turnout, even in large electorates, thus providing a potential solution to the paradox that citizens vote in large

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.

[3]Franklin and Hoboldt (2011) show that European voters whose first eligible election is a (low-turnout) European Parliament election are less likely to vote in subsequent national elections, and Atkinson and Fowler (2013) report that the occurence of saint's day fiestas depress turnout in both current and future Mexican elections. Similarly to Gerber et al. and Meredith, these papers require additional assumptions to justify a habit formation interepretion.

[4]Another relevant series of papers tests for habit formation in consumption (Heien and Durham 1991; Dynan 2000).

[5]Other authors have used weather shocks to estimate the effect of race riots on urban development in the U.S. (Collins and Margo 2007) and of political protests on policy changes in France (Huet-Vaughn 2013).

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 some assumptions in 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. These theories offer a promising direction toward the understanding of mass voter turnout, which social scientists have so struggled to explain (see Feddersen 2004), even going so far as to call it "the paradox that ate rational choice" (Fiorina 1990). Our documentation of habit formation in voting provides an empirical basis for that theoretical project.

2 Motivation: Age Patterns in Voting

As motivation for our interest in habit formation, Figure 1 displays U.S. federal election turnout as a function of age using data from the Current Population 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. 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 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.[6] The natural implication is that the perceived benefits of voting increase with age more rapidly than does the opportunity cost of time. Although this rising perceived benefit of voting 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, separating the scatter plot by previous presidential election eligibility.[7] 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

[6] It is possible 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.

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

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.[8]

3 Habit Formation in a Downsian Framework

We consider habit formation in a dynamic extention to the "calculus of voting" framework of Downs (1957), Tullock (1967), and Riker and Ordershook (1968, 1973). 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. 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 known as the "expressive utility" of voting, representing intrinsic satisfaction or "warm glow" (Andreoni 1989) from carrying out a civic duty. The citizen incurs cost C_{it} from voting, also regardless of the election outcome. All four terms are positive, so she votes if and only if:[9]

$$P_{it}B_{it} + D_{it} \geq C_{it} \tag{1}$$

Denote the voting decision as V_{it}, which equals 1 if the citizen votes, 0 otherwise.

3.1 Identifying Habit Formation

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 constantly high levels of B_{it} or D_{it} turn out often, while those with low levels of these variables will rarely vote. A regression of current turnout on its

[8]The evidence in Figure 1, Panel B, is similar to that of Meredith (2009), who studies age patterns in voting using a dataset from California with fine-grained age data. As we discuss below, although the evidence suggests habit formation, it fails to distinguish the effect of past eligibility from the effect of past voting *per se*.

[9]Technically, the citizen is indifferent between voting and abstaining when expression (1) holds with equality. To simplify the exposition, we assume throughout that the citizen votes in this situation.

lagged values is thus a poor test of habit formation, since persistent unobserved heterogeneity may explain any serial correlation in voting.

To identify whether $V_{i,t-1}$ affects V_{it}, we wish to take advantage of a *transitory* lagged shock $\xi_{i,t-1}$ to one of the terms of the model. The shock must satisfy:

$$\{P_{it}, B_{it}, D_{it}, C_{it}\} \,|V_{i,t-1} \perp \xi_{i,t-1} \tag{2}$$

which states that, conditional on the voting decision in period $t-1$, the shock is independent of P_{it}, B_{it}, D_{it}, and C_{it}. Assumption (2) 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}. Under this assumption, an association between $\xi_{i,t-1}$ and V_{it} provides reduced-form evidence of habit formation.

Note that even if ξ_{it} is independent of the baseline benefits and costs of voting before the realization of the shock, it may not satisfy assumption (2). For example, consider a randomized intervention that encouraged citizens to vote in period $t-1$. Randomization guarantees that the intervention is independent of baseline benefits and costs. But depending on its nature, the intervention may directly influence a citizen's sense of civic duty or cost of voting for many periods into the future. In this case, ξ_{it} affects $V_{i,t+1}$, but not necessarily solely through V_{it}.

Two important contributions to the literature on habit formation in voting fit this characterization. In the first, 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, in our framework, this interpretation requires the assumption 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 assumption (2) would be violated. A similar logic applies to Meredith (2009), who uses data from California to compare the voting behavior of those whose 18^{th} birthday was just before the 2000 general election to 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 less generalizable sample. Meredith estimates that those barely eligible to vote in 2000 are more likely to vote in 2004, which he wishes to interpret as evidence that past voter participation affects current voter participation. But this interpretation requires the assumption that experiencing a presidential campaign while eligible to vote has no persistent direct effects on a citizen's tastes and costs. Here again, the assumption is debatable. 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 has persistent effects on the perceived benefits and costs of voting, then assumption (2) is violated.

As an alternative approach to identifying habit formation in voting, we exploit a transitory shock to the cost of voting: election-day precipitation. Four important characteristics of this shock justify our choice. First, as we show below (and as previous research has established), precipitation 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 future voting costs. Fourth, net of the year fixed effects, county fixed effects, and county-specific trends we include in our econometric model (see Section 5 for details), the remaining variation in precipitation 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 precipitation 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 may adapt their consumption of political information in the period leading up to the election, which may lead to a violation of assumption (2).

At the same time, we note a potential exclusion restriction violation in the case of precipitation. 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 on a rainy day (rather than the negative act of *not* voting on a rainy day) may reduce future voting propensity, so that an effect of lagged precipitation does not imply habit formation. However, this potential violation requires that voters fail to fully attribute the unpleasantness of voting on a rainy day to bad weather. Given that most voters have experienced many rainy days in the past, such attribution error is likely to be minimal.

3.2 Modeling Habit Formation

Although the mere identification of habit formation is one of our central contributions, we also propose and test a specific mechanism through which it occurs. Specifically, we hypothesize that habit formation operates primarily through the accumulation of the expressive utility of voting. In a setup similar to Becker and Murphy's adictive goods model (1988), we specify D_{it} as a stock variable that increases with past acts of voting:

$$D_{it} = D_{i,t-1} + \alpha V_{i,t-1} \tag{3}$$

where $\alpha > 0$ captures the extent to which past acts of voting increase the utility flow from future acts of voting. A citizen enters the electorate with initial value D_{i0}, drawn from a population distribution; later values of D_{it} reflect her stock of civic duty or attachment to the democratic process, which is endogenous to past acts of voting. If changes in P_{it}, B_{it}, and C_{it} from $t-1$ to t are unrelated to the voting decision in $t-1$, then this accumulation in D_{it} implies that the act of voting in period $t-1$ increases the likelihood of voting in period t.

To draw attention to the accumulation framework's implications for our empirical work, we focus on its predictions for voter behavior following cost shocks. These predictions require a slightly more precise setup. To this end, let ξ_{it} be a cost shock on $\left[0, \bar{\xi}\right]$, and let $\tilde{C}_{it}$ be the cost of voting prior to the cost shock, so that $C_{it} = \tilde{C}_{it} + \xi_{it}$. Combining terms, let $\tilde{U}_{it} = P_{it}B_{it} + D_{it} - \tilde{C}_{it}$ be the baseline utility from voting, before the realization of the cost shock. Assume that ξ_{it} is independent of $\tilde{U}_{it}$ and that changes in P_{it}, B_{it}, and C_{it} are invariant to V_{is}, $s < t$. Within this setup, we highlight three testable predictions.

The first prediction—which follows directly from our assumption of D_{it} depending on past turnout—is that current and past cost shocks negatively affect turnout. To see this, note that a citizen votes if and only if $\tilde{U}_{it} \geq \xi_{it}$. Some citizens are infra-marginal; those with $\tilde{U}_{it} < 0$ will never vote in period t or in any period thereafter, while those with $\tilde{U}_{it} \geq \bar{\xi}$ are sure to vote in period t and for all subsequent periods. We refer to these citizens as *never-voters* and *always-voters*, respectively. But so long as marginal citizens exist—i.e., so long as $\tilde{U}_{it} \in \left[0, \bar{\xi}\right)$ for some citizens—then $\Pr[V_{it}=1|\xi_{it},\xi_{i,t-1}]$ decreases in ξ_{it}. Furthermore, because voting in period $t-1$ increases D_{it}, $\Pr[V_{it}=1|\xi_{it},\xi_{i,t-1}]$ also decreases in $\xi_{i,t-1}$. As a result, the framework implies that both current and lagged precipitation decrease voter turnout.

The framework also gives rise to a second prediction, involving the interaction of current and lagged cost shocks: a higher draw of $\xi_{i,t-1}$ increases the sensitivity of $\Pr[V_{it}=1|\xi_{it},\xi_{i,t-1}]$ to ξ_{it}. To establish this result, we focus on citizens who are marginal in period $t-1$, with $\tilde{U}_{i,t-1}\in\left[0,\bar{\xi}\right)$. For a citizen of this type, a sufficiently large cost shock $\xi_{i,t-1}$ keeps her from voting in period $t-1$, so that in the following period, $\tilde{U}_{it}$ is smaller than it would have been following a rain-free election. As a result, the minimum cost shock required to keep the citizen from voting in period t is smaller than it would have been following a rain-free election. In other words, large shocks prevent marginal individuals from becoming less marginal in the next election. So the framework predicts that larger lagged precipitation shocks make turnout more sensitive to current precipitation shocks.

A third prediction, which matches Figure 1 but is not relevant to our main empirical work below, is that turnout rises with age if the distributions of P_{it}, B_{it}, and $\tilde{C}_{it}$ are time-invariant. Citizens who are marginal in period 0—i.e., those with $\tilde{U}_{i0}\in\left[0,\bar{\xi}\right)$—will gradually become always-voters in subsequent periods, while those who are initially infra-marginal will remain infra-marginal in all periods. Hence, the framework provides one explanation for why the young vote less than the old. Even so, it is unlikely to be the sole explanation. In Figure 1, turnout rises with age even across ages that do not differ in exposure to federal elections: for instance, 18 and 19 year-olds. 19 year-olds are more likely to vote than 18 year-olds even though citizens of both ages have been exposed to exactly one federal election. While these youths may differ in exposure to local elections, other factors may also play a role in their different turnout rates.

Throughout this discussion, we have assumed that D_{it} does not depreciate over time. One can incorporate depreciation by rewriting equation (3) as $D_{it}=(1-\delta)\,D_{i,t-1}+\alpha V_{i,t-1}$, where $\delta\in(0,1)$ is the depreciation rate. This modification does not change the framework's main prediction: that turnout decreases in both current and past cost shocks. But it does complicate the second and third predictions above, regarding the interaction of current and past cost shocks and the age profile of turnout. Both these predictions hold, however, if the depreciation rate δ is sufficiently small relative to the habit formation coefficient α.

4 Data

Implementation of our empirical strategy requires consideration of spatial scales when combining weather and voting turnout data. Mid-latitude precipitation systems, as observed over the United States, can have characteristic lengths anywhere between 2 and 1,000 km. Thus, the spatial shapes of precipitation systems do not fit naturally onto political boundaries. In the U.S., with the exception of the states with the smallest area, precipitation systems rarely cover an entire state. This insight yields three particular requirements for our data. First, the turnout data must be at the smallest political unit available, to reduce measurement error in the spatially-averaged precipitation variable. Second, the pixel resolution of the raw precipitation data must be fine enough to guarantee that most political units cover at least one pixel. Finally, because daily precipitation may be spatially correlated, the precipitation data must have the broadest geographical coverage, to guarantee a sufficient number of independent observations for a given day.

For politics data, we use county-level presidential election returns for the years 1952-2012 to generate two variables of interest: voter turnout, which we define as the ratio of votes to eligible voters, and the Republican vote share.[10] For weather, we acquire data with the highest spatial and temporal resolution available for the continental United States. Daily precipitation data for the continental United States for days before, on, and after election day from 1948-2012 come from the NOAA Climate Prediction Center's Unified Gauge-Based Analysis of Precipitation. 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.[11] 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.[12]

Table 1 provides means, standard deviations, and several quantiles for the variables in our analysis. Voter turnout averages at 58 percent, with a fairly symmetric distribution ranging from a

[10]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.

[11]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.

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

10^{th} percentile of 42 to a 90^{th} percentile of 76. The Republican vote share, too, is fairly symmetrically distributed around a mean of 55. In contrast, the precipitation distribution is right-skewed, with a median of 0, a mean of 2.5 millimeters, and a 90^{th} percentile of 7.1 millimeters.[13] Elections are therefore usually free of precipitation, and when precipitation does occur, it is typically minor.

5 Empirical Strategy

In our main regression specification, we estimate turnout as a function of current and lagged election-day precipitation, a year fixed effect, a county fixed effect, and a county-specific linear time trend. For county i in election year t:

$$turnout_{it} = \beta_0 precip_{it} + \beta_1 precip_{i,t-1} + \tau_t + \eta_i + \lambda_i t + \varepsilon_{it} \tag{4}$$

Note that $t-1$ corresponds to the previous election, four years earlier. To assess the robustness of this specification, we perform two three checks. First, we gauge the sensitivity of the coefficients on election-day precipitation to the inclusion of the vector of covariates listed in the bottom panel of Table 1, and to the inclusion of higher-order county-specific trends. Second, we run placebo tests by including measures of precipitation two weeks after election day, as well as precipitation on future election days. Third, we verify the stability of the results in samples that omit all observations from a single state or a single year.

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.[14] A combination of two factors make this wide cluster definition appropriate. First, precipitation is spatially correlated. Second, both 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 these facts, we regress both turnout and precipitation on a year fixed effect, a county fixed effect, and a county-specific trend, mapping the residuals for selected years in Appendix Figure 1. Precipitation residuals are clustered over large areas, while turnout residuals tend to cluster within state borders.

[13] 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.

[14] Our conclusions remain unchanged when we use Conley's (1999) non-parametric estimator for standard errors allowing for arbitrary spatial dependence in a 1500 kilometer radius. If anything, the spatially adjusted standard errors are smaller than the clustered standard errors.

This reduced-form regression is instructive, but our focus on habit formation leads to interest in identifying an auto-regressive model:

$$turnout_{it} = \rho turnout_{i,t-1} + \nu_{it} \tag{5}$$

For reasons discussed in Section 3.1, ordinary least squares regression does not identify this model. However, one can use estimates of β_0 and β_1 from equation (3) to compute estimates of the causal parameter ρ: $\hat{\rho} = \frac{\hat{\beta}_1}{\hat{\beta}_0}$ converges in probability to ρ. This ratio is an instrumental variables (IV) estimator for ρ, in which lagged precipitation serves as an instrument for lagged turnout. We estimate its variance using the delta method.

Two aspects of this estimator merit further discussion. First, as with other IV estimators, it requires the monotonicity assumption that turnout weakly decreases in precipitation 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. Second, our estimator does not necessarily identify the degree of habit formation at the individual level, as we discussed in Section 3.1. In the presence of social effects—for example if people speak to their neighbors about a positive voting experience—the aggregate ρ for the county may be larger than the individual-level habit formation parameter. We return to this issue when we discuss the magnitudes of our results in Section 7.

Building on the main regression specification, we also estimate two models with interactions. First, guided by the theory, we fit a model with the interaction of current and lagged precipitation:

$$turnout_{it} = \beta_0 precip_{it} + \beta_1 precip_{i,t-1} + \gamma precip_{it} \cdot precip_{i,t-1} + \tau_t + \eta_i + \lambda_i t + \varepsilon_{it} \tag{6}$$

The theoretical framework predicts that γ should be negative; lagged precipitation shocks increase the sensitivity of turnout to contemporaneous precipitation shocks. Second, we interact precipitation with county demographic and socio-economic characteristics, X_{it}:

$$\begin{aligned} turnout_{it} &= \beta_0 precip_{it} + \beta_1 precip_{i,t-1} + X'_{it}\theta \\ &\quad + precip_{it} X'_{it}\gamma_0 + precip_{i,t-1} X'_{it}\gamma_1 + \tau_t + \eta_i + \lambda_i t + \varepsilon_{it} \end{aligned} \tag{7}$$

The theory does not have specific predictions for these interactions, but they are instructive for assessing which citizens respond to precipitation and which citizens are habit forming.

6 Results

This section presents our main findings. We first demonstrate that precipitation on election day adversely affects contemporaneous voter turnout. This result is robust to a number of modeling choices and placebo tests. Next, we show that precipitation on election day has a strong persistent effect on turnout in subsequent presidential elections, and we confirm the framework's prediction that contemporaneous and lagged precipitation interact. To conclude the section, we explore how the effects of contemporaneous and lagged precipitation vary across several demographic characteristics.

6.1 Effect of Contemporaneous Precipitation on Turnout

For the equivalent of a "first stage" estimate, Table 2 focuses on the contemporaneous effect of election-day precipitation on voter turnout. In Column (1), we estimate equation (4) with no lags and find a statistically significant coefficient (at the 1 percent level) implying that a 1 millimeter increase in precipitation decreases voter turnout by 0.07 percentage points. Column (2) shows that the linear model used in Column (1) is reasonable by estimating a semi-parametric model using dummy variables for increasing bins of precipitation. With just one exception (the 8-12 millimeter category), each successive increase in precipitation is associated with a further decrease in turnout. Relative to the omitted category of 0 millimeters, an intense election-day storm with rainfall totaling more than 20 millimeters causes turnout to fall by 2 percentage points.

As a placebo test, we test whether precipitation before and after election day affected election turnout. In Column (3), we include linear terms for the precipitation 7 and 14 days before and after election day. All four placebo coefficients are substantially smaller and less significant than the main effect. Three are statistically insignificant, while precipitation one week after election day is marginally significant at the 8 percent level. Overall, the evidence in Table 1 points to a robust effect of election-day precipitation on contemporaneous turnout.

6.2 Effect of Lagged Precipitation on Turnout

Table 3 shows that the turnout effects of election-day precipitation persist to future elections. Column (1) estimates equation (4), with both contemporaneous and lagged precipitation terms. Both coefficients are statistically significant at the 1% level, with turnout falling 0.08 and 0.07 percentage points per millimeter of contemporaneous and lagged precipitation, respectively. This finding changes little with the addition of county-level covariates in column (2).[15] In both columns, the implied habit formation parameter ρ is roughly 0.9, implying that a 1 percentage point rise in period $t-1$ turnout increases period t turnout by 0.9 percentage points. This estimate of habit formation in voter turnout is substantially larger than existing estimates in the literature, a matter we discuss in Section 7.

The implied habit formation coefficient changes similarly little in columns (3)-(6), which include various combinations of placebos, leads, and lags. As a placebo check, columns (3) and (4) add contemporaneous and lagged precipitation two weeks after election day to columns (1) and (2), respectively. Neither placebo is significantly associated with turnout, and the coefficients on election-day precipitation are little changed. In columns (5) and (6), we add the lead and second lag of election-day precipitation. The lead serves as another falsification exercise: future election-day precipitation does not affect current turnout. Also of interest is the coefficient on twice-lagged precipitation, which is significantly negative, as the theoretical framework would predict. The magnitudes of the coefficients are unsteady, but the implied estimate of ρ—which we calculate by averaging the ratio of the t coefficient to the $t-1$ coefficient and the ratio of the $t-1$ coefficient to the $t-2$ coefficient—falls only slightly, to 0.8.[16] Notes also that column (6) shows small and insignificant coefficients on all leads and lags of precipitation two weeks after election day.

We conduct two further robustness checks for these main results. First, Appendix Table 1

[15] In unreported results, our estimates are also unaffected by the inclusion of daily temperature as a covariate. The daily temperature data are from Schlenker (2009).

[16] Following the discussion in Section 5, both ratios are consistent estimators for ρ. If turnout were really an autoregressive process of order 1 (AR-1), we could also use a third estimator: the square root of the ratio of the t coefficient to the $t-2$ coefficient. However, the theoretical framework implies that voting in period $t-2$ affects the probability of voting in period t even if an individual does not vote in period t, so the process is not AR-1.

reports estimates from specifications with county-specific trends of different orders. With quadratic or cubic trends, the results are of a similar magnitude and significance level, suggesting that linear trends are sufficiently flexible. The trends are important, however; when we omit them altogether, the results become unstable. Second, Appendix Figure 2 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, β_1, and ρ vary little. We conclude that the results are not driven by outliers.

In the last two columns of Table 3, we verify the theoretical framework's prediction that a past precipitation shock increases the sensitivity of turnout to a contemporaneous precipitation shock. In estimates of equation (6) in columns (7) and (8), the coefficient on the interaction of current and lagged precipitation is significantly negative. Again, when column (8) uses precipitation two weeks after election day as a placebo, neither the main effects nor the interaction effect are significant. In magnitude, the interaction effect is 12 percent of the main effect of contemporaneous precipitation. To illustrate, consider the effects of an intense storm with rainfall totaling 1 inch (roughly the 98^{th} percentile of the precipitation distribution). If the previous election was rain-free, then the storm reduces turnout by 1.4 percentage points. If the previous election had a similar storm, then the current storm reduces turnout by 1.6 percentage points.

6.3 Heterogeneity

Who responds to current and lagged precipitation? 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 shed some light on it. Table 4 reports estimates of equation (7), which interacts these covariates with contemporaneous and lagged precipitation.[17] For comparison, column (1) repeats the main estimate of equation (4) from Table 3 (column 1).

Table 4's most noteworthy result, shown in column (5), is that the effects of lagged and contemporaneous precipitation are weaker in counties with high population density. The main effects and interactions imply that at the 10^{th} percentile of population density, a millimeter of current or lagged precipitation reduces turnout by 0.13 percentage points, while at the 90^{th} percentile of

[17] 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 all the covariates between census years induces correlated measurement errors.

population density, a millimeter of current or lagged precipitation reduces turnout by 0.04 to 0.06 percentage points. This finding matches the conventional wisdom that inclement weather imposes greater costs on rural voters than on their urban counterparts, due to their longer distances from the polls and their access to fewer modes of transportation.

Also interesting in Table 4 is the finding that the interactions with county's over-65 population share are of opposite sign. Although the coefficients on these interactions are not individually significant, the difference between them is significant at the 6 percent level. Because the interaction with contemporaneous precipitation has a positive coefficient and that with lagged precipitation has a negative coefficient, we can infer that counties with more sizable elderly populations are less sensitive to contemporaneous precipitation and more sensitive to lagged precipitation. Recalling our theoretical framework, one interpretation is that the elderly are predominantly always-voters or never-voters, but that the few who are marginal are especially sensitive to precipitation shocks. Consequently, older citizens as a group appear to be especially habit-forming.

7 Assessing Magnitudes

7.1 Comparison with Previous Research

With estimated habit formation parameters ρ between 0.7 and 0.9, our results imply more habit formation than previous research. Gerber et al. (2003) place this parameter value at 0.5 in their get-out-the-vote experiment, while Meredith (2009) places it at 0.075 in his regression discontinuity design based on voting age restrictions. These differences have several possible explanations.

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, one that is likely to be small. In both Gerber et al.'s context and our own, always-voters exist, so that both study designs identify local average treatment effects for marginal voters (the compliers). In contrast, Meredith strategy allows for no always-voters; individuals just short of their 18^{th} birthdays cannot vote under any circumstance. As a result, Meredith effectively recovers a treatment-on-the-treated habit formation parameter that averages the effect of past voting on future voting for marginal *and* infra-marginal voters. The effect is zero for infra-marginal voters, so Meredith's tiny (though statistically significant) estimate is unsurprising.

The fact that our estimate exceeds that of Gerber et al. presents a greater puzzle, which has four possible explanations. First, our estimator may pick up interpersonal spillovers that are not present in Gerber et al.'s design, as discussed in the next subsection. 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.[18] The effect of voting in a low-stakes midterm election on voting in a subsequent local election may be smaller than the effect of voting in a presidential election on voting in a subsequent presidential election. Third, the sub-populations induced to vote may differ between the two studies. Gerber et al.'s estimate applies to residents of New Haven who are sensitive to get-out-the-vote campaigns, whereas ours applies to citizens around the country who are sensitive to inclement weather.[19] Fourth, 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, undermining the study design.

7.2 Spillovers and Social Interactions

If neighbors tell one another about positive voting experiences, or if social norms about voting change following a high turnout election, then the county-level habit formation parameter will exceed the individual-level habit formation parameter. This theory has support from a recent literature on the role of social interactions in voting (Nickerson 2008; Gerber et al. 2008; DellaVigna et al. 2013).

Formally, let ϕ denote the effect of a millimeter of precipitation on an *individual's* probability of turnout, and δ be the share of individuals that are habit forming (i.e., if they do not vote at t, they will not vote at $t+1$). In the absence of social interactions, our estimated current effect of precipitation on turnout (β_0) equals ϕ, and the lagged effect of precipitation (β_1) equals $\phi\delta$. The ratio of the latter to the former is δ, the habit formation parameter. However, under the presence of (positive) social interactions, the effect of an input (precipitation) on a decision (turnout) estimated at a more aggregate (county) level exceeds the effect at the individual level (Case and Katz 1991, Glaeser and Scheinkman 2003), due to the presence of a social multiplier. In our context, if an

[18] 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.

[19] Also note that the Gerber et al. experiment is specific to two elections in 1998-1999, while our study covers the 1952-2012 period.

individual's probability of voting is positively affected by the number of voters in his county, then precipitation will have not only a direct effect on turnout, but also an indirect effect via social interactions.

Applying Glaeser and Scheinkman's (2003) approach to our context, we define the reference groups to be the county (our unit of observation) and a social interactions parameter θ. Its interpretation is that an exogenous 1 percentage point increase in the average turnout of the county raises the probability that an individual resident of that county votes by θ: $V_{ic} = \theta \bar{V}_c$, where $\bar{V}_c$ is average turnout in county c.[20] To use this setup for to study habit formation, we separate the social interactions parameter into two components: θ^S and θ^L. The former captures the short-run social interactions present in the effect of current rainfall on turnout: for example, if voter 1's turnout decision depends on the weather, and voter 2 depends on voter 1 to give her a ride to the polls. In contrast, θ^L represents the long run social interactions captured by the effect of lagged precipitation on turnout: for example, if voter 1 convinces some of her peers, in the four years between elections, about the advantages of voting or not voting. As our examples suggest, a natural assumption is that $\theta^L > \theta^S$, since (given the difficulty in predicting precipitation) θ^L captures the effects of social interactions in the course of four years, while θ^S captures those that occur within a day.

In terms of the newly-defined parameters, the effects of current and lagged rainfall are $\beta_0 = \frac{\phi}{1-\theta^S}$ and $\beta_1 = \frac{\phi\delta}{1-\theta^L}$, respectively. The ratio between these terms is $\rho = \delta\frac{1-\theta^S}{1-\theta^L}$. If short-run social interactions are negligible ($\theta^S = 0$), the individual-level habit formation parameter δ can be obtained by multiplying ρ by $\beta_0 = (1 - \theta^L)$. For example, the ratio between coefficients in our preferred specification is 0.89 (Table 3). If the social interactions parameter is 0.2 (for every four people directly induced not to vote by lagged precipitation, one more will be convinced not to vote given social interactions), then the individual-level habit formation (δ) is 0.71.

8 Alternative Theories of Habit Formation in Voting

Although our habit formation framework provides a compelling explanation for the dependence of current on past turnout, the Downsian model suggests other possible mechanisms as well. Recall

[20] In principle, $\bar{V}_c$ excludes individual i's turnout from the county mean. However, given the size of county populations, this consideration is negligible. Under Glaeser and Scheinkman's (2002) terminology, the social interactions are "global" and can be seen as a reduced form approximation to a more complex pattern of social interactions that accounts for individual i's social network within the county.

from equation (1) that a citizen votes if and only if $P_{it}B_{it} + D_{it} \geq C_{it}$. We have hypothesized that only D_{it} depends on past voting experiences. Conceivably, any of the framework's other three terms could also depend on past voting experiences. In this section, we explore this possibility for each additional term of the framework but conclude that our habit formation framework best matches the data. We begin with P_{it} and then consider B_{it} and C_{it} in turn. The section concludes by discussing why partisan politics (a factor outside the Downsian model) is unlikely to play a major role.

8.1 Political Efficacy (P_{it})

In one theory with relevance for our results, past voting experiences shape citizens' sense of external political efficacy (Campbell et al. 1954): the degree to which they believe their actions to affect political outcomes. A citizen with limited understanding of the electoral system may learn over time about her probability of being pivotal. Suppose she takes that probability to be time-invariant, so that P_{it} represents her latest estimate of the probability, based on her experiences with the electoral system. Under Bayesian updating, P_{it} increases after voting for the winner or not voting while supporting the loser, and it decreases after voting for the loser or not voting while supporting the winner. Consistent with this 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.

Even without further analyses, existing evidence suggests that this theory falls short of explaining our results. First, on a conceptual note, the theory is inconsistent with rational expectations and most forms of forward-looking behavior. Under such assumptions, voters would use all available information about the probability of being pivotal, to which their past voting experiences are not relevant. Second, an explanation based on political efficacy 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 priors for their pivotalness. Third, although 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 because voters who supported the winner are of roughly the same number as voters who supported the loser. Contrary to this prediction, our estimates of ρ are large even though most presidential elections during our sample period were decided by margins of less than 10 points.[21]

We can also test this theory's divergent predictions for voting for the winner and voting for the loser. To do so, Table 5 interacts contemporaneous and lagged precipitation with measures of whether a county is politically aligned with the winner or loser of the previous election. To avoid issues of endogeneity, we use a county's Republican vote share two elections ago to ascertain its partisan leaning. We define two new terms for Table 5. A county is "partisan" if its Republican vote share in $t-2$ was in the top or bottom 10 percent of the vote share distribution, and a county is "aligned" if it is both "partisan" and politically aligned with the winner in $t-1$. A dry day increases the likelihood of voting for the winner in aligned counties and increases the likelihood of voting for the loser in counties that are partisan but not aligned. Therefore, in the model:

$$\begin{aligned} turnout_{it} \quad = \quad & \beta_0 precip_{it} + \beta_1 precip_{i,t-1} + \theta^p partisan_{i,t-2} + \theta^a aligned_{i,t-1} \\ & + \gamma_0^p \left(partisan_{i,t-2} \times precip_{it}\right) + \gamma_1^p \left(partisan_{i,t-2} \times precip_{i,t-1}\right) \\ & + \gamma_0^a \left(aligned_{i,t-1} \times precip_{it}\right) + \gamma_1^a \left(aligned_{i,t-1} \times precip_{i,t-1}\right) \\ & + \gamma_0^r \left(Rwinner_{t-1} \times precip_{it}\right) + \gamma_1^r \left(Rwinner_{t-1} \times precip_{i,t-1}\right) \\ & + \tau_t + \eta_i + \lambda_i t + \varepsilon_{it} \end{aligned} \tag{8}$$

we expect partisan, non-aligned counties to have less pronounced effects of lagged precipitation than non-partisan counties ($\gamma_1^p > 0$); aligned counties to have more pronounced effects of lagged precipitation than partisan, non-aligned counties ($\gamma_1^a < 0$); and aligned counties to have less-pronounced effects of lagged precipitation than non-partisan counties ($\gamma_1^a + \gamma_1^p < 0$). In equation (8), we control for whether a Republican won the previous election ($Rwinner_{t-1}$) because $aligned_{i,t-1}$ is essentially an interaction between $partisan_{i,t-2}$ and $Rwinner_{t-1}$.

Estimates of equation (8), shown in column (1) of Table 5, fail to fully confirm these predictions. Although counties aligned with the winner of the last election exhibit the strongest effects of lagged precipitation (both γ_1^a and $\gamma_1^a + \gamma_1^p$ are negative), we find no evidence that counties aligned with the loser are different from non-partisan counties (γ_1^p is not significantly different from zero). In other words, widespread voting for the winner raises future county turnout, but widespread voting for the loser does not decrease it. While this finding may speak to the psychology of habit formation,

[21] In unreported results, the coefficient on lagged rainfall did not vary with the national margin of victory in the previous election.

it does not support a model in which citizens learn about their probability of being pivotal.

A potential concern with regression specification (8) is that the inclusion of $partisan_{i,t-2}$ violates the strict exogeneity assumption required for fixed effects estimation, due to a relationship between turnout and vote shares.[22] To address these concerns, column (3) uses an alternative measure of $partisan_{i,t-2}$, based on the Republican vote share predicted by a county's demographic and socio-economic characteristics. Specifically, we run a regression of the Republican vote share on the white population share, the over-65 population share, log median household income, log population density, an indicator for location in the South, and—because of the South's unique politics over the second half of the twentieth century—interactions of the South indicator with all other covariates.[23] Using the predicted values from this regression, we define a county as partisan if its predicted vote share is in the top or bottom 10 percent of the predicted vote share distribution. Because this measure of partisanship is a generated regressor, standard errors are block-bootstrapped at the state level. The magnitudes of the coefficients change somewhat, but the conclusion remains unchanged: aligned, partisan counties exhibit a high degree of habit formation, but dis-aligned, partisan counties are not significantly different from non-partisan counties.

To explore the robustness of the vote-for-the-winner effect, columns (2) and (4) control more flexibly for underlying partisanship. Instead of including the "partisan" term and its interactions, these models distinguish between heavily Republican and heavily Democratic areas. This alternative specification cannot test for average differences between partisan and non-partisan counties, but it can estimate the excess sensitivity of partisan, aligned counties relative to partisan, dis-aligned counties (equivalent to γ_1^a in equation(8)). The coefficient on the interaction of aligned with lagged precipitation remains significantly negative.

8.2 Instrumental Utility (B_{it})

A separate explanation for our results involves the strength of citizens' political preferences. If the act of voting causes an individual to care more about political outcomes or to develop more polarized political preferences, then past voting could affect current voting through B_{it}, the benefit to the individual if her preferred candidate wins. An effect of past voting on B_{it} is consistent with

[22] After one controls for year fixed effects, county fixed effects, and county-specific trends, the data show no relationship between turnout and vote shares. We return to this issue below.

[23] The interactions of the South indicator with other covariates are not crucial to the results.

Mullainathan and Washington's (2009) finding that, due to cognitive dissonance in the choice of candidates, the act of voting causes a citizen to further improve her opinion of her chosen candidate. However, if citizens have objective beliefs about the probability of being pivotal, then any effect on B_{it} will likely have limited consequences for the voting decision because it will be multiplied by a number approaching zero.

Additionally, we can leverage the fact that B_{it} is multiplied by P_{it} to test whether accumulation in B_{it} can explain our results. To this end, we extend our framework to:

$$P_{it}(B_{i,t-1} + \alpha_B V_t) + (D_{i,t-1} + \alpha_D V_t) \geq C_{it} \tag{9}$$

where we now allow separate accumulation terms in B_{it} and D_{it}, as captured by α_B and α_D. While the distinction between B_t and D_t is clear in our framework, empirically separating the two mechanisms may be difficult. The act of voting may lead a citizen to change her tastes regarding political participation; the distinction is whether these tastes take the form of instrumental value (caring about the outcome) or expressive value (caring about voting). If B_{it} accumulates ($\alpha_B > 0$), then evidence of habit formation will be stronger when P_{it} is high. Our test thus introduces interactions between precipitation and state-level voter pivotalness, in the following specification:

$$\begin{aligned} turnout_{it} \;=\; & \beta_0 precip_{it} + \beta_1 precip_{i,t-1} + \theta pivotal_{it} \\ & + \gamma_0 \left(pivotal_{it} \times precip_{it}\right) + \gamma_1 \left(pivotal_{it} \times precip_{i,t-1}\right) + \tau_t + \eta_i + \lambda_i t + \varepsilon_{it} \end{aligned} \tag{10}$$

The static Downsian framework predicts $\theta > 0$ and $\gamma_0 > 0$: a higher likelihood of being pivotal increases turnout and offsets the negative effects of precipitation on election day. If the act of voting in the previous period increases B_{it}, we should also observe $\gamma_1 < 0$. To construct a measure for P_{it}, we use the forecasting model developed by Campbell (1992) and extended in Campell et al. (2006) to obtain an *ex ante* (before precipitation) predicted state-level Democratic vote share for elections from 1952-2004. The predicted Democratic vote share and its forecast uncertainty determine the probability that a randomly drawn voter will hold the tie-breaking vote for a given state in a given election year.[24]

[24] We use Campbell et al.'s (2006) model to predict the Democratic vote share, $\hat{d}_{st}$ for state s and election year t. The probability of a randomly drawn voter breaking a state-level tie is $(1/N_{st})\phi\left(\hat{d}_{st} - 0.5/\hat{\sigma}_{st}\right)$, where $\phi(\cdot)$ is the standard normal density function, $\hat{\sigma}_{st}$ is the standard deviation of $\hat{d}_{st}$, and N_{st} is the number of registered voters. Our conclusions do not change if we use predicted closeness rather than predicted pivotalness.

Table 6, which estimates equation (10), shows no evidence of a role for the probability of being pivotal. To ensure sample selection is not affecting our results, column (1) re-estimates our main specification for the sample of election days from 1952-2004 for which we have a measure of state-level pivotalness. The implied ρ of 0.72 is not statistically distinct from our main sample result presented in Table 3. The rest of the table tests for the role of pivotalness. Column (2) estimates equation (10) using a continuous measure of state-level pivotalness, and none of the estimated parameters of interest are statistically significant.[25] At the same time, the effects of uninteracted current and lagged precipitation, which now capture the effects of habit formation when $P_{it} = 0$ and can be due solely to accumulations in D_{it}, remain statistically significant. The implied ρ from accumulations in D_{it} alone is 0.65, a number within the uncertainty of our main result in column (1) of Table 3. In column (3), we explore an alternative specification by constructing a indicator for observations with the pivotal probabilities exceeding the sample median. We again find no evidence of accumulation in in B_{it}. Altogether, Table 6 suggests that habit formation is likely operate through accumulation in the expressive and not instrumental value of voting.

8.3 Voting Costs (C_{it})

The final term of the Downsian framework, C_{it}, may also play a role: past experience with voting may lower the future cost of voting. This mechanism has two potential sources, one personal and one institutional. As an example of personal costs, voters must occasionally incur informational "fixed costs:" learning the location of the polling station and the best way to get there. They may also be uncertain of how much time the act of voting takes; if they are risk averse, they will become more likely to vote once they learn the true opportunity cost of voting. While this hypothesis is plausible, it is unlikely to be the only mechanism driving our results. First, if informational fixed costs matter, one would expect the lagged effect of rainfall to be smaller in counties with older populations (whose voters have more experience going to the polls), which is not the case in our data. Second, individuals who cast a vote for the winner are more likely to form habits, which is difficult to reconcile with a model in which voting lowers informational costs. In other words, if habit formation were mediated entirely by voting costs, then reduction of future voting costs would

[25] The point estimates and standard errors for both the interacted and uninteracted pivotal coefficients are large because the probability of being pivotal is typically on the order of 10^{-4} percent.

have to be particularly salient in partisan counties, only when the preferred candidate wins. We can think of little reason why learning about location, for example, would interact with *ex post* election results in such a manner.

On the institutional side, state and county election offices have at various points implemented laws that purge inactive voters from the registration rolls. These laws varied substantially (both across and within states) with regard to the frequency of the purges and the number of elections a voter must miss to be purged (Mitchell and Wlezien 1995). After the 1993 National Voter Registration Act (NVRA), automatic purges of non-voters has ceased in all states. Nonetheless, while they were in effect, these laws could produce "habit formation," as we have defined it. Non-voters in several consecutive elections would lose their registration, raising the cost of future voting and making them less likely to vote again. However, these laws are unlikely to explain our results. First, the purging of inactive voters from the voter rolls has no obvious connection to the vote-for-the-winner effects reported in Table 5, as there is no reason purging would vary by county partisanship with timing that matches the identity of the election winner. Second, to further test the importance of these laws in our results, we re-estimated our main specification using only the states and elections with no automatic purging of non-voters. Specifically, for every presidential election year, we used *The Book of the States*[26] to code whether, in a given year, a state made registration "*subject to cancellation for failure to vote at certain specified intervals*." Restricting the estimation of 4 to counties in states that had no automatic purging of non-voters at election t yields results very similar to those from our main specification (Column 1, Table 3).[27] . This result is consistent with previous findings that the effects of purging non-voters on turnout are negligible (Wolfinger and Rosenstone 1980, Mitchell and Wlezien 1995).[28]

Note that we have only discussed costs incurred immediately prior to the act of voting. Longer-term costs, like voter registration or learning about party platforms, precede the realization of the precipitation shock. Given the difficulty of predicting deviations from trend long in advance, they

[26]http://knowledgecenter.csg.org/kc/content/book-states-archive-1935-2009

[27]Volumes of *The Book of the States* from before 1960 do not report on automatic purging, so these estimates use only the 1960-2012 elections. 48% of observations in this period are from state-elections without automatic purging (20,799 observations). After 1993, the NVRA has abolished automatic purging of voters, and hence all observations from the 1996 election and afterward enter the estimation. In this subsample, the effect of current precipitation on turnout is -0.069 (S.E. = 0.032), while the effect of lagged precipitation is -0.069 (S.E. = 0.039), implying a ρ of 1.01 (S.E. = 0.33).

[28]Mitchel and Wlezien (1995) state that removing purge laws "*serve to keep people on the rolls who are not very likely to vote*."

are also not amenable to adjustment in anticipation of the precipitation shock.

8.4 What Role for Politics?

Until now, the discussion has treated the voter as an isolated individual, rather than as a participant in an interactive political process. But actions by political elites may play a role, especially in light of existing evidence that rain-induced decreases in turnout have a partisan bias, benefitting Republicans (Gomez, Hansford, and Krause 2007; Hansford and Gomez 2010). If precipitation shocks affect election outcomes, and if incumbents are especially empowered to manipulate voter turnout, then the persistent effects of precipitation shocks may have a political explanation.

Several pieces of evidence suggest that such political explanations are not the primary mechanism. First, in unreported results, the effect of lagged precipitation did not vary with the party of the incumbent or with an indicator for whether the incumbent was running for re-election. Second, regional variation in the effects of precipitation, reported in Table 6, is at odds with theories in which lagged precipitation matters because it changes the party in power. Precipitation decreases turnout in both the South and the Non-South, but it has divergent effects on the Republican vote share: decreasing it in the South, increasing it elsewhere. Our main findings are not due to precipitation advantaging a party that has outsized sway over voter turnout.

9 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 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 precipitation to estimate the effects of voting habit formation on future election turnout. We find that a 1 percentage point decrease in current turnout decreases future turnout by 0.7-0.9 percentage points. Additional analyses motivated by a dynamic Downsian framework 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 occurs through an accumulation in the expressive value 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. We hope our finding of habit formation in voting re-ignites interest in the underlying psychological and social determinants of the intrinsic value voters place on the act of voting. As Feddersen (2004) suggests, our evidence of habit formation should also help inform political economy models of strategic voter mobilization and suppression.

References

Andreoni, J. (1989). "Giving with Impure Altruism: Applications to Charity and Ricardian Equivalence." *Journal of Political Economy* 97(6): 1447-1458.

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.

Bendor, J., D. Diermeier, and M. Ting. (2003). "A Behavioral Model of Turnout." *American Political Science Review* 97(2): 261-280.

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. (1992). "Forecasting the Presidential Vote in the States." *American Journal of Political Science* 35(2): 386-407.

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.

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.

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.

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.

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

Dynan, K.E. (2000). "Habit Formation in Consumer Preferences: Evidence from Panel Data." *American Economic Review* 90(3): 391-406.

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

Fiorina, M. (1990). "Information and Rationality in Elections." In J. Ferejohn and J. Kuklinski, eds., *Information and Democratic Processes.* Urbana: University of Illinois Press.

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.

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., 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. (2003). "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.

Heien, D., and C. Durham. (2003). "A Test of the Habit Formation Hypothesis Using Household Data." *Review of Economics and Statistics* 73(2): 189-199.

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

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.

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. (Forthcoming). "Do Political Protests Matter? Evidence from the Tea Party Movement." *Quarterly Journal of Economics.*

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.

Mitchell, G.E., and C. Wlezien. (1995). "The Impact of Legal Constraints on Voter Registration, Turnout, and The Composition of The American Electorate." *Political Behavior* 17(2): 179-202.

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.

Plutzer, E. (2002). "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood*." American Political Science Review* 96(1): 41-56.

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

Riker, W.H., and P.C. Ordeshook. (1973). *An Introduction to Positive Political Theory.* Englewood Cliffs, NJ: Prentice-Hall.

Schlenker, W. (2009). "Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change." *Proceedings from the National Academy of Science* 106(37): 15594-8.

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

Wolfinger, R., and S. Rosenstone. (1980). *Who Votes?* New Haven: Yale University Press.

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

Panel A: All Ages

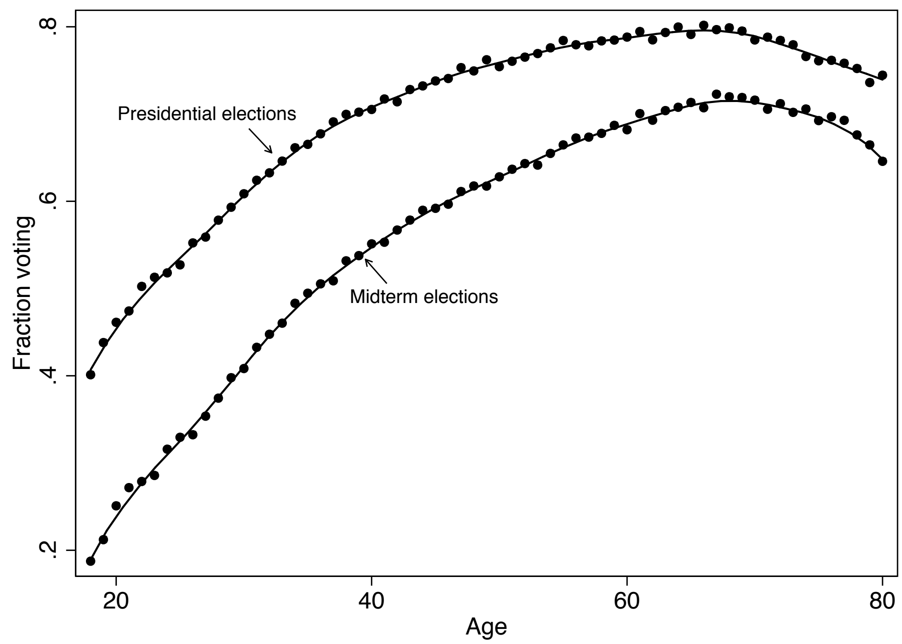

Panel B: Ages 18-27

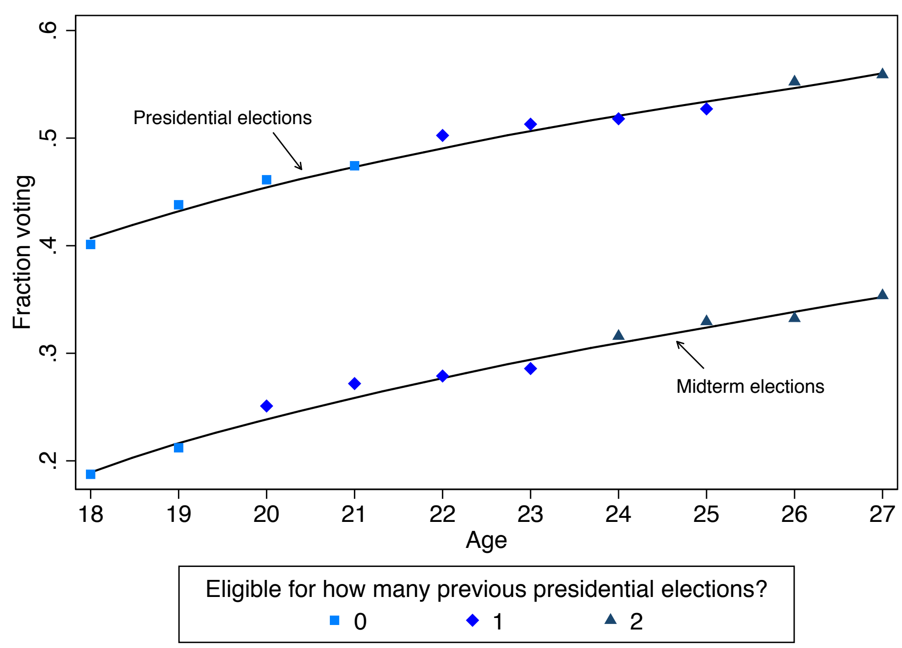

Note: Scatter plots are age-specific rates, while curves are local linear regressions with a bandwidth of 2 years.

Table 1: Descriptive Statistics

			Percentiles				
	Mean	Std. Dev.	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
	(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							
Precipitation on election day (mm)	2.5	6.6	0.0	0.0	0.0	1.4	7.1
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	10.4	10.6	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 49,524 county-year observations, based on presidential elections from 1952-2012 in 3,108 counties.

Table 2: Effect of Contemporaneous Precipitation on Turnout

	(1) Linear model	(2) Bin model	(3) Placebo
Precip. 2 wks before election day, t			0.029
			[0.022]
Precip. 1 wk before election day, t			-0.025
			[0.023]
Precip. on election day, t	-0.069		-0.065
	[0.024]***		[0.023]***
Precip. 1 wk after election day, t			-0.039
			[0.022]*
Precip. 2 wks after election day, t			-0.034
			[0.036]
(4,8] mm precip. on election day, t		-0.654	
		[0.332]*	
(8,12] mm precip. on election day, t		-0.436	
		[0.466]	
(12,16] mm precip. on election day, t		-1.508	
		[0.512]***	
(16,20] mm precip. on election day, t		-1.817	
		[0.667]***	
(20,95] mm precip. on election day, t		-2.003	
		[0.862]**	
Number of county-years	49,524	49,524	49,524
Number of counties	3108	3108	3108
Election years	1952-2012	1952-2012	1952-2012

Note: Brackets contain standard errors clustered at the state level. All regressions include year fixed effects, county fixed effects, and county-specific linear trends. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Effect of Contemporaneous and Lagged Precipitation on Turnout

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Precip. on election day, t+1					-0.028	-0.027		
					[0.027]	[0.035]		
Precip. on election day, t	-0.078	-0.070	-0.077	-0.068	-0.113	-0.110	-0.058	-0.056
	[0.026]***	[0.025]***	[0.025]***	[0.024]***	[0.040]***	[0.037]***	[0.027]**	[0.026]**
Precip. on election day, t-1	-0.070	-0.065	-0.068	-0.063	-0.116	-0.112	-0.058	-0.055
	[0.025]***	[0.024]**	[0.023]***	[0.022]***	[0.045]**	[0.041]***	[0.024]**	[0.023]**
Precip. on election day, t-2					-0.061	-0.057		
					[0.025]**	[0.023]**		
(Precip., t)*(Precip., t-1)							-0.007	-0.007
							[0.002]***	[0.002]***
Precip. 2 wks. after election day, t+1						-0.021		
						[0.026]		
Precip. 2 wks. after election day, t			-0.029	-0.025		-0.008		-0.032
			[0.038]	[0.032]		[0.039]		[0.021]
Precip. 2 wks. after election day, t-1			-0.018	-0.018		-0.008		-0.022
			[0.044]	[0.039]		[0.048]		[0.030]
Precip. 2 wks. after election day, t-2						-0.036		
						[0.034]		
(Precip. 2 wks. after, t)								0.007
*(Precip. 2 wks. after, t-1)								[0.005]
Implied $\boldsymbol{\rho}$	0.89	0.93	0.88	0.92	0.78	0.77		
	[0.28]***	[0.31]***	[0.28]**	[0.32]***	[0.14]***	[0.14]***		
Number of county-years	49,524	49,524	49,524	49,524	43,300	43,300	49,524	49.524
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	1952-2012	1952-2012	1956-2008	1956-2008	1952-2012	1952-2012
County covariates?	No	Yes	No	Yes	No	No	No	No

Note: Brackets contain standard errors clustered at the state level. Regressions include year fixed and county fixed effects and county trends. County covariates are: white pop. share, the over-65 pop. share, log median income, log pop. density. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Heterogeneity in the Effects of Contemporaneous and Lagged Precipitation

	(1)	(2)	(3)	(4)	(5)
Precip. on election day, t	-0.078	-0.195	-0.115	-1.374	-0.152
	[0.026]***	[0.119]	[0.053]**	[0.685]***	[0.041]***
Precip. on election day, t-1	-0.070	-0.195	-0.041	-0.618	-0.159
	[0.025]***	[0.124]	[0.037]	[0.493]	[0.044]***
(% white)*(Precip., t)		0.0015			
		[0.0013]			
(% white)*(Precip., t-1)		0.0015			
		[0.0013]			
(% over 65)*(Precip., t)			0.0034		
			[0.0026]		
(% over 65)*(Precip., t-1)			-0.0024		
			[0.0018]		
(Log median income)*(Precip., t)				0.123	
				[0.064]*	
(Log median income)*(Precip., t-1)				0.052	
				[0.046]	
(Log pop. density)*(Precip., t)					0.017
					[0.006]***
(Log pop. density)*(Precip., t-1)					0.021
					[0.007]***
Number of county-years	49,524	49,524	49,524	49,524	49,524
Number of counties	3,108	3,108	3,108	3,108	3,108

Note: Brackets contain standard errors clustered at the state level. All regressions include year fixed effects, county fixed effects, county-specific linear trends, and the main effects of any demographic variables included in the interaction terms. The sample includes all presidential elections from 1952 to 2012. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Role of County Partisanship and Alignment with the Winner

	Actual partisanship		Predicted partisanship	
	(1)	(2)	(3)	(4)
Precip. on election day, t	-0.071	-0.064	-0.055	-0.058
	[0.023]***	[0.023]***	[0.027]**	[0.027]**
Precip. on election day, t-1	-0.044	-0.045	-0.037	-0.041
	[0.028]	[0.029]	[0.026]	[0.028]
(Aligned, t-1) × (Precip., t)	0.002	0.085	0.173	0.180
	[0.058]	[0.083]	[0.141]	[0.196]
(Aligned, t-1) × (Precip., t-1)	-0.123	-0.131	-0.180	-0.180
	[0.064]*	[0.062]**	[0.069]***	[0.080]**
(Partisan, t-2) × (Precip., t)	0.036		-0.202	
	[0.036]		[0.068]***	
(Partisan, t-2) × (Precip., t-1)	0.021		0.029	
	[0.037]		[0.063]	
(Heavily Dem., t-2) × (Precip., t)		-0.104		-0.194
		[0.047]**		[0.069]***
(Heavily Dem., t-2) × (Precip., t-1)		0.005		0.020
		[0.040]		[0.078]
(Heavily Rep., t-2) × (Precip., t)		-0.070		-0.217
		[0.063]		[0.171]
(Heavily Rep., t-2) × (Precip., t-1)		0.057		0.081
		[0.035]		[0.074]
(Rep. winner, t-1) × (Precip., t)	-0.011	-0.004	0.029	0.029
	[0.038]	[0.035]	[0.038]	[0.040]
(Rep. winner, t-1) × (Precip., t-1)	-0.057	-0.055	-0.068	-0.062
	[0.052]	[0.053]	[0.053]	[0.054]
Aligned, t-1	2.84	2.50	2.55	2.35
	[0.95]***	[1.05]**	[1.33]*	[1.42]
Partisan, t-2	-1.42		0.024	
	[0.37]***		[0.75]	
Heavily Dem., t-2		-1.53		-0.36
		[0.53]***		[0.89]
Heavily Rep., t-2		-0.64		0.40
		[0.66]		[0.90]
Sum of coefs. on (Aligned, t-1)×(Precip, t-1)	-0.102		-0.151	
and (Partisan, t-2) × (Precip., t-1)	[0.043]**		[0.071]**	
Number of county-years	46,329	46,329	46,329	46,329
Number of counties	3,108	3,108	3,108	3,108

Note: Brackets contain standard errors clustered at the state level: asymptotic in cols. (1)-(2), block bootstrapped in cols. (3)-(4). All regressions include year and county fixed effects and county trends. Sample includes all presidential elections from 1956 to 2012. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Precipitation, Turnout, and Vote Shares, Regional Comparison

	Full sample		Non-South		South	
	Turnout	Rep. share	Turnout	Rep. share	Turnout	Rep. share
	(1)	(2)	(3)	(4)	(5)	(6)
Precip. on election day, t	-0.078	-0.022	-0.084	0.078	-0.040	-0.077
	[0.026]***	[0.027]	[0.039]***	[0.028]***	[0.030]	[0.034]**
Precip. on election day, t-1	-0.070	-0.041	-0.039	0.014	-0.066	-0.047
	[0.025]***	[0.034]	[0.011]***	[0.023]	[0.026]**	[0.034]
Number of county-years	49,524	49,441	26,918	26,917	22,606	22,524
Number of counties	3108	3108	1685	1685	1423	1423

Note: Brackets contain standard errors clustered at the state level. All regressions include year fixed effects, county fixed effects, and county-specific linear trends. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Appendix Figure 1: Precipitation and Turnout Residuals, Selected Years

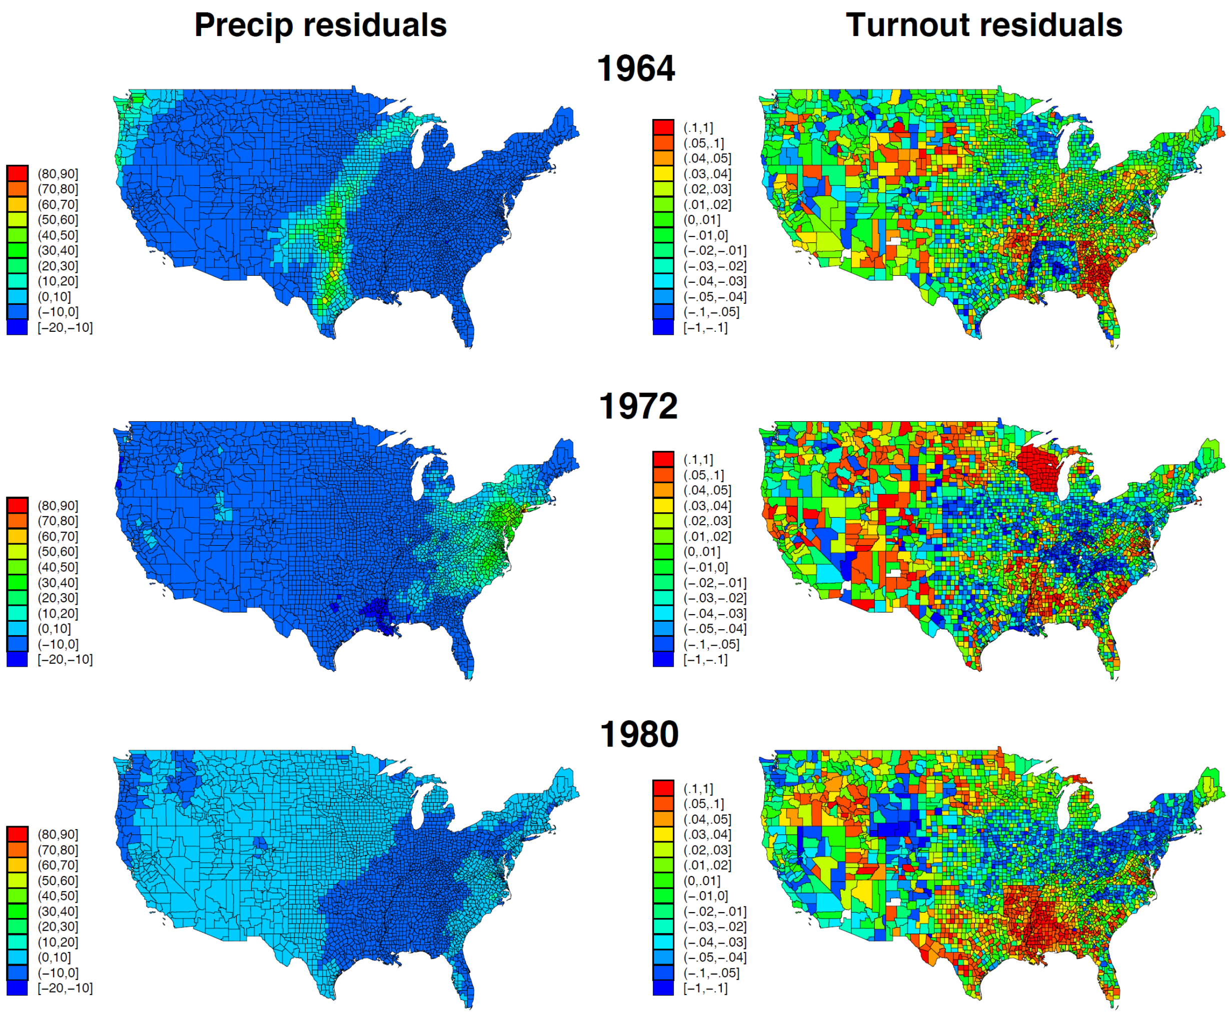

(continued on next page...)

(continued from previous page...)

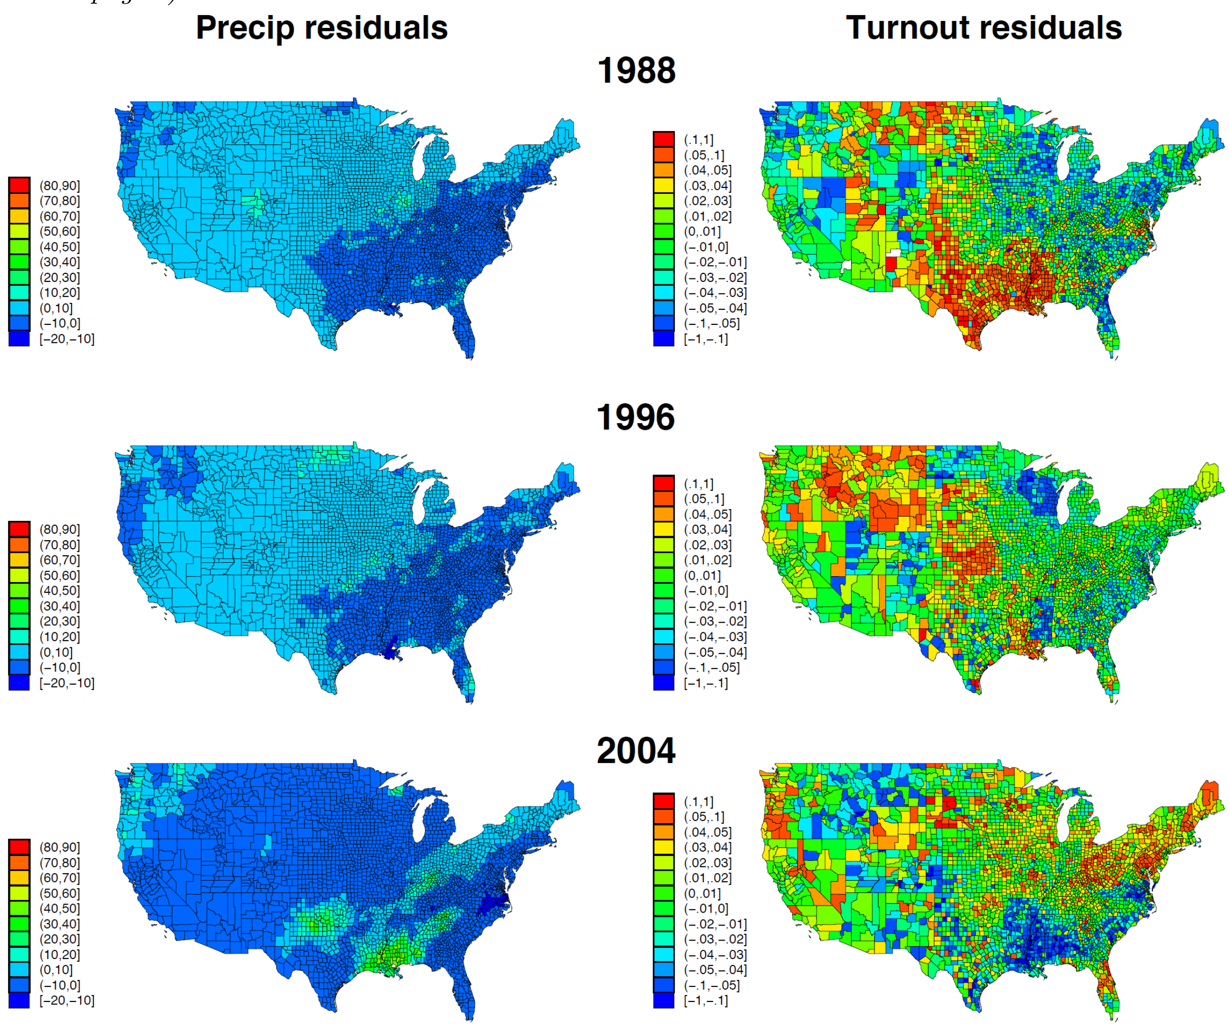

Note: Residuals from regressions of precipitation (mm) and turnout on year and county fixed effects, and county-specific trends.

Appendix Figure 2: Leave-One-Out Checks

Panel A: Leave Out One State

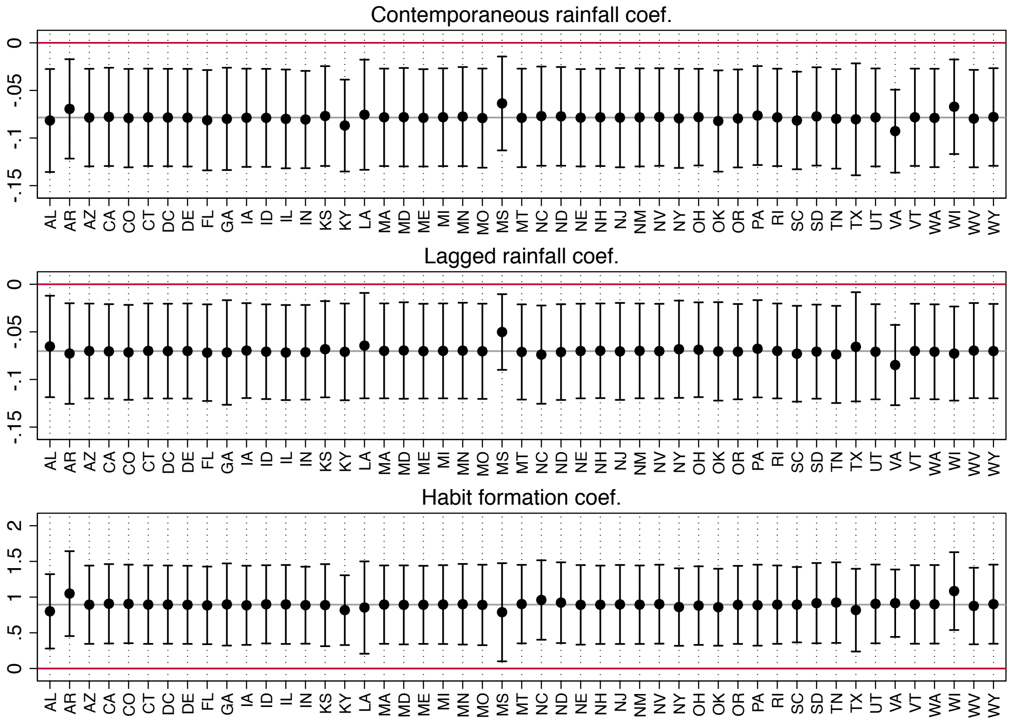

Panel B: Leave Out One Year

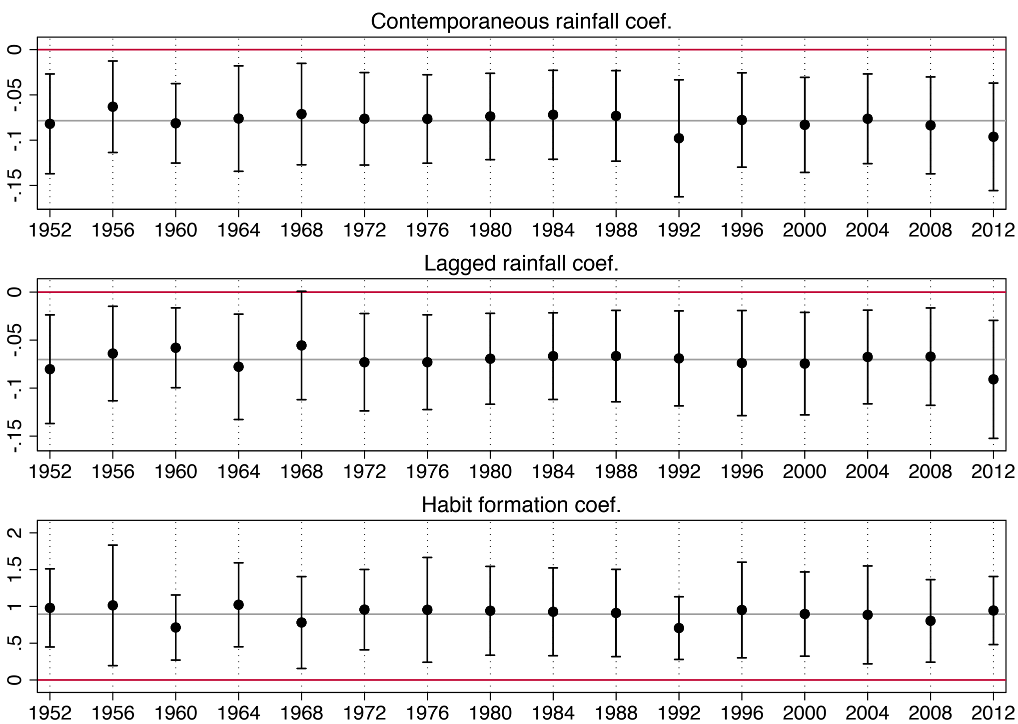

Note: Each estimate is based on a sample that omits the state or year on the *x*-axis. Dots are coefficients; bars are 95% CIs. Light gray horizontal lines represent full-sample estimates.

Appendix Table 1: Higher-Order County-Specific Trends

	(1)	(2)	(3)	(4)
Precip. on election day, t	-0.011	-0.078	-0.062	-0.063
	[0.030]	[0.026]***	[0.023]***	[0.023]***
Precip. on election day, t-1	0.016	-0.070	-0.057	-0.057
	[0.021]	[0.025]***	[0.021]***	[0.021]***
Implied $\boldsymbol{\rho}$	-1.43	0.89	0.92	0.91
	[4.81]	[0.28]***	[0.32]***	[0.32]***
Number of county-years	49,524	49,524	49,524	49,524
Number of counties	3,108	3,108	3,108	3,108
Trend order	0	1	2	3

Note: Brackets contain standard errors clustered at the state level. All regressions include year fixed effects and county fixed effects. The sample includes all presidential elections from 1952 to 2012. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$