Estimating Habit Formation in Voting*

Thomas Fujiwara† Kyle Meng‡ Tom Vogl§

December 2014

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 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. 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 Ethan Kaplan, Ilyana Kuziemko, Doug Miller, Stefano DellaVigna, and seminar participants at Columbia University, 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 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 Ordershook 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 such as 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 role of 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 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. To the best of our knowledge, this latter condition has yet to be satisfied in the literature.

We address this empirical challenge by exploiting unexpected and transitory 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. To ground the analysis conceptually, we present a conceptual framework for studying habit formation based on a simple "calculus of voting" model, in which precipitation 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 precipitation 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 current and lagged election-day precipitation reduce voter turnout. Our main estimates imply that a 1 percentage point decrease in past turnout lowers current turnout by 0.7-0.9 percentage points. All 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 of our interpretation that the effect of lagged precipitation is due to habit formation. First, turnout shows no relation to precipitation on the day of the next presidential election. Second, it shows no relation to precipitation two weeks after both the current and previous election days. These results confirm that only precipitation that fell precisely on the previous election day matters. Third, the effects of both current and lagged precipitation are stronger in rural areas, where the costs of inclement weather are likely greater.

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 between elections. Because precipitation 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. We explain that these existing estimates are subject to identification concerns, but 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 is being driven by an increase in the consumption value of voting. Voting habits thus fall under the purview of 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.[5] 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. Although these prior studies use innovative designs and provide an important benchmark for our findings, our conceptual framework makes clear that their voting persistence results cannot be attributed entirely to habit formation unless their treatments are sure to have no direct effects on the future determinants of turnout. In Gerber et al.'s experiment, the canvassing procedure included messages appealing 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

[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.

[5]Relatedly, Franklin and Hoboldt (2011) show that Europeans whose first eligible election is a (low-turnout) European Parliament election vote less in national elections; 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.

that a transitory and unexpected shock in the cost of voting, such as precipitation, is better suited for estimating habit formation. In addition to this issue related to research design, the paper differs from these previous studies in two other dimensions. First, we estimate county-level habit formation incorporating social multipliers, as discussed above, whereas Gerber et al. and Meredith estimate individual-level effects. Estimates at both levels deserve interest on their own, but estimates that capture social interactions may be more relevant for the analysis of larger-scale policy interventions and other shocks. Second, our sample covers the entire country over 60 years, during which all U.S. counties received precipitation 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.[6]

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 to estimate the effect of participating in Tea Party protests and independence day celebrations on political preferences and behavior.[7] 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.[8] 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.

[6]In both contexts, political competition is low and Democrats dominate federal elections.

[7]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).

[8]More generally, a growing empirical literature examines how an individual's previous experiences affect attitudes such as trust (Aghion et al. 2010; Nunn and Wantchekon 2011) and preferences for equality (Alesina and Fuchs-Schueldeln 2007; Fisman et al. 2009; Voors et al. 2012; Giuliano and Spilimbergo 2013).

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.[9]

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. 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.[10] The natural implication is that the perceived benefits of voting increase with age more 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

[9]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).

[10]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.

highlight this evidence, Panel B of Figure 1 zooms in on ages 18-27, separating the scatter plots by previous presidential election eligibility.[11] 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 source of identifying variation, we propose election-day precipitation. 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 precipitation affects entire communities, rather than individuals.

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

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 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.[12] 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.[13] 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, both our strategy and existing research rely on shocks to D_{it} and C_{it}.

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

[12] 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 $B \equiv \left|V_{it}^{R} - V_{it}^{D}\right|$.

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

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}.[14]

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 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.

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.

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

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 18^{th} 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 Precipitation

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 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 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 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, we note two potential exclusion restriction violations for our research design. First, 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 precipitation 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. Second, precipitation may reduce canvassing activity, which we have argued may have persistent effects on preferences. Unfortunately, data on election-day canvassing are not available for the study period, preventing us from assessing this identification threat.

Supposing that the exclusion restriction holds, election-day precipitation can identify the effect of past turnout on future turnout at the local level. However, because precipitation 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. 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. 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 well-identified estimates may be larger yet, due to a social multiplier. We return to this issue when we describe our econometric method in Section 4.2 and again when we interpret our results in Section 6.

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 precipitation cannot by itself disentangle these mechanisms, we draw on other sources of variation to shed light on this issue in Section 7.

4 Data and Econometric Method

4.1 County-Level Panel Dataset

Mid-latitude precipitation systems, as observed over the United States, can be anywhere between 2 to 1,000 km wide with spatial extents that do not fit naturally onto political boundaries. This implies three data requirements. First, the data pixel resolution of the precipitation data must be fine enough to guarantee that most political units cover at least one pixel so that there is variation in precipitation across neighboring units. Second, turnout data must be at the lowest political unit available so as to reduce measurement error when pixel-level weather data is aggregated up to the political unit. Finally, because daily precipitation is spatially correlated across large areas, the precipitation data must 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.

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.[15] For weather, we acquire data with the highest spatial and temporal resolution available for the continental United States. Daily gridded precipitation data for

[15] 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.

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.[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 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. Similarly, county-level covariates appear symmetrically distributed.

Election-day precipitation exhibits four noteworthy statistical properties. First, election day precipitation is a relatively infrequent event. Table 1 shows that the unconditional 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.[18] Second, when precipitation does occur, it is typically experienced by many counties at once, given the large spatial extent of precipitation systems. Appendix Figure 1 plots the share of U.S. counties that experience any precipitation and precipitation 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 precipitation on election day is rare, all counties experience precipitation at some point in our sample period. Appendix Figure 2 plots the cumulative share of counties that experienced any precipitation over the sample period, indicating that nearly all counties have experienced election day precipitation by 1972, or 20 years into our sample period. This finding implies that our estimations use variation from all counties. Finally, variation in election-day precipitation differs considerably across U.S. counties. The Appendix Figure 3 displays the histogram of the standard deviation in election-day precipitation across counties and shows a fairly large spread in precipitation variability across counties.

[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.

4.2 Econometric Method

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 c in election year t:

$$turnout_{ct} = \beta_0 precip_{ct} + \beta_1 precip_{c,t-1} + \tau_t + \eta_c + \lambda_c t + \varepsilon_{ct} \tag{4}$$

Note that $t-1$ corresponds to the previous election, four years earlier. To assess the robustness of this specification, we perform four checks. First, equation 4 assumes turnout responds linearly to precipitation level. To verify this functional assumption, we also estimate a semi-parametric model with precipitation bin dummies both in levels and in terms of quintiles defined from the within-county precipitation distribution. Second, we gauge the sensitivity of the coefficients on election-day precipitation to the inclusion of the vector of covariates listed in Table 1, and to the inclusion of higher-order county-specific trends. Third, we run placebo tests by including measures of precipitation two weeks after election day, as well as precipitation on future election days. Lastly, we verify the stability of the results to outliers by using subsamples 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.[19] 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, and we map the residuals for an example year, 2004, in Figure 2. Precipitation residuals are clustered over large areas, while turnout residuals tend to cluster within state borders.

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

$$turnout_{ct} = \rho turnout_{c,t-1} + \nu_{ct} \tag{5}$$

[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.

As we discussed in Section 3.1, OLS regression does not identify this model. However, if we define the error term as $\nu_{ct} \equiv \tau_t + \eta_c + c_i t + \varepsilon_{ct}$ (as in equation (4)), 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.[20] This ratio can be seen as 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.

Three 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 identifies 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 6.

Third, as mentioned in Section 3.3, our estimator does 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 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 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 6.

[20] 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.

5 Main 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. To conclude the section, we explore how the effects of contemporaneous and lagged precipitation vary across several demographic characteristics.

5.1 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 implying that a 1 millimeter increase in precipitation decreases voter turnout by 0.07 percentage points. Column (2) shows that linearity is a reasonable funcational form by estimating a semi-parametric model using dummy variables for increasing bins of precipitation. With just one exception, 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. Columns (1) and (2) implicitly assume that the marginal response of voters to precipitation levels are the same regardless of typical local weather. In particular, voters in counties that are historically wetter on election day may respond differently to precipitation levels than voters in historically dryer counties. To address this concern, column (3) uses local precipitation quintiles based on the historical distribution of election-day precipitation for each county. We detect an overall negative response, though our point estimates are noisier and do not decrease monotonically with each quintile. Our preferred specification will be linear in levels because it is the most precisely estimated and because it allows for a cleaner interpretation when recovering our persistence parameter ρ.

As a falsification exercise, we test whether precipitation before and after election day affected election turnout. In column (4), we include linear terms for precipitation on the 7^{th} and 14^{th} days before and after election day. All four placebo coefficients are substantially smaller and less significant than the main effect. Neither of the lagged precipitation terms have statistically significant coefficients, nor does precipitation two weeks after election day. The coefficient on precipitation one week after election day is significant, but only marginally, at the 8 percent level. This result may reflect the fact that counties with large areas experience the same precipitation system within a one week period as the weather system migrates over the county. In unreported analyses, we drop all counties with areas above 1000 sq. km. and find no significant placebo coefficients. Overall, the evidence in Table 2 points to a robust effect of election-day precipitation on contemporaneous turnout. For conciseness, we focus on the 14-days-after placebo check in subsequent analyses.

5.2 Effect of Lagged Precipitation on Turnout

Table 3 shows that the turnout effects of 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. Figure 3 confirms our linear functional forms by plotting coefficients from a semi-parametric specification with dummy variable bins for current and lagged precipitation in levels (Panel A) and local quintiles (Panel B). Returning to Table 3, our finding changes little with the addition of county-level covariates in column (2).[21] In both columns (1) and (2), 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.[22] This estimate of habit formation in voter turnout is substantially larger than existing estimates in the literature, a matter we discuss in Section 6.

The implied habit formation coefficient changes similarly little in columns (3)-(5), which include various combinations of placebos, leads, and lags. As a placebo check, column (3) adds contemporaneous and lagged precipitation two weeks after election day to column (1). Neither placebo is significantly associated with turnout, and the coefficients on election-day precipitation are little changed. In column (4), 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. The magnitudes of the coefficients are unsteady, but the implied estimate of ρ—which we calculate by averaging estimates of $\frac{\beta_1}{\beta_0}$, $\frac{\beta_2}{\beta_1}$, and $\sqrt{\frac{\beta_2}{\beta_0}}$, all of which are consistent estimators for

[21] Our estimates are also unaffected by controlling for temperature (using data from Schlenker and Roberts 2009).
[22] If we weight observations by county population, we obtain nearly identical estimates of ρ.

ρ—falls only slightly, to 0.8.[23] Note also that column (4) 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 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 4 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.

Our main results in Table 3 deal solely with presidential elections, leaving open interesting questions about habit formation in non-presidential elections (as in Green et al. 2003), as well as interactions between non-presidential and presidential elections. Unfortunately, no comprehensive dataset on midterm election turnout exists for our study period, so data constraints prevent a full exploration of this issue.[24] Nevertheless, we can estimate how precipitation on the last midterm election day affects presidential turnout. Because turnout is much lower in midterm than presidential elections (see Figure 1), we conjecture that marginal voters in midterm elections tend to be inframarginal in presidential elections. As such, we anticipate that midterm election precipitation does not affect turnout in the following presidential election. The final column of Table 3 tests this hypothesis, finding no effect of precipitation in the last midterm election.

5.3 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 help shed some light. Table 4 reports regressions that interact these covariates with contempora-

[23] β_0 and β_1 are defined in equation (4), while β_2 is the coefficient on $precip_{c,t-2}$.

[24] Comprehensive county-level midterm results are not available for the 1992-1998 elections. For multiple states (e.g., Arkansas, California, Minnesota, Pennsylvania), county-level results are missing for at least two midterm elections in the main source of our presidential election turnout for the 1950-1990 period (ICPSR Study 8611). Finally, a number of states had, at different points in time, legislation not requiring tabulation of votes for unopposed candidates, creating issues of missing data as well as measurement error in cases of counties with areas that partially overlap with uncontested districts.

neous and lagged precipitation.[25] For comparison, column (1) repeats the main estimate of equation (4) from Table 3 (column 1).

Most of the interaction effects in Table 4 are not statistically significant, although column (5) contains a noteworthy result: the effects of both lagged and contemporaneous precipitation are significantly 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 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. Nevertheless, the interaction effects are not significantly different from each other, so we cannot reject the null hypothesis that rural and urban voters are equally habit-forming. More broadly, although Table 4 provides evidence on heterogeneity that supports the credibility of our precipitation results, the table does not advance our understanding of which voters or communities are most habit-forming.[26] Without finely geocoded individual-level data on turnout, we cannot shed further light on this issue.

6 Assessing Magnitudes

Although we have argued that previous research designs may not fully satisfy exclusion restriction (3), past estimates of habit formation in voting nevertheless serve as an important and interesting basis for comparison. This section discusses the magnitude of our results in the context of existing studies. In interpreting the variation across studies, we pay particular attention to the possible role of spillovers due to social interactions.

[25] 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.

[26] Also of potential interest in Table 4 is the finding that the interactions with the 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. This finding may imply that older citizens are especially habit-forming.

6.1 Comparison with Previous Research

We estimate a habit formation parameter ρ between 0.7 and 0.9. By comparison, Gerber et al. (2003) place their persistence parameter at 0.5 in their get-out-the-vote experiment, while Meredith (2009) estimates persistence to be 0.075 using a regression discontinuity design based on voting age restrictions. In Section 3.1, we described why our strategy to identify habit formation differs from previous efforts to capture persistence in voting, but the biases we discussed there are unlikely to explain why our estimates exceed the others by so much. Specifically, because exclusion restriction violations in the other study designs would bias their estimates upward, one would expect our estimates to be smaller than theirs. Here, we propose other possible explanations for the different magnitudes of our estimates.

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 (the compliers). In contrast, Meredith strategy does not allow for 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 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. We propose four possible explanations. First, as discussed in the next subsection, 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

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

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.

6.2 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 precipitation 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 in Section 3, 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 precipitation 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 $\theta \in [0, 1)$ 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 precipitation 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, the strength of social interactions between election days determines the relationship between the individual- and county-level habit formation parameters.[30]

Unfortunately, little evidence exists on the size of θ. Even if we had individual-level data, we

[29] These derivations use the fact that there is a large number of voters in each county (Glaeser and Scheinkman 2002). 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 precipitation). 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] Glaeser et al. (2003) study the conditions under which the ratio between aggregate and individual estimates pins down the social multiplier parameter.

could not distinguish r from ρ using our estimation strategy because precipitation 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 that their parameter may not map cleanly onto r, it is nonetheless an average of individual-level persistence parameters for a group of marginal voters, making it a useful benchmark. Combined with our baseline estimate of ρ at 0.89, their estimate of r at 0.51 implies $\theta = 0.43$.

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 0.89, this estimate implies $\theta = 0.44$. 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 $\theta = 0.38$. Remarkably, all these exercises yield similar estimates for θ, around 0.4, implying a social multiplier of $\frac{1}{1-\theta} = 1.7$.

7 Mechanisms

Recall from equation (1) that a citizen i votes if and only if $P_{it}B_{it} + D_{it} \geq C_{it}$. Conceivably, any of the framework's terms could depend on past voting experiences, and any of these terms could be influenced by social interactions. In this section, we explore each term's possible role in explaining our results, and we also consider whether partisan politics contributes to the explanation. We conclude, by process of elimination, that the evidence best supports an explanation based on the accumulation of expressive utility, D_{it}. Because D_{it} is consumption utility from the act of voting, this explanation has much in common with the classic models of habit formation in consumption by Pollak (1970) and Becker and Murphy (1988).

7.1 Political Efficacy (P)

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 affecting the election outcome. 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 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. Corroborating evidence is reported Table 8, which Section 7.5 describes in greater detail. Column (3) of the table introduces interaction terms between precipitation and the national margin of victory in the previous election. Neither the effect of current precipitation nor the effect of lagged precipitation vary 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. To do so, Table 6 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 6. 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 following model for turnout in county c at time t:

$$
\begin{aligned}
turnout_{ct} \quad = \quad & \beta_0 precip_{ct} + \beta_1 precip_{c,t-1} + \theta^p partisan_{c,t-2} + \theta^a aligned_{c,t-1} \\
& + \gamma_0^p \left(partisan_{c,t-2} \times precip_{ct} \right) + \gamma_1^p \left(partisan_{c,t-2} \times precip_{c,t-1} \right) \\
& + \gamma_0^a \left(aligned_{c,t-1} \times precip_{ct} \right) + \gamma_1^a \left(aligned_{c,t-1} \times precip_{c,t-1} \right) \\
& + \gamma_0^r \left(Rwinner_{t-1} \times precip_{ct} \right) + \gamma_1^r \left(Rwinner_{t-1} \times precip_{c,t-1} \right) \\
& + \tau_t + \eta_c + \lambda_c t + \varepsilon_{ct}
\end{aligned} \tag{6}
$$

we expect 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 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 (6), we also control for whether a Republican won the previous election ($Rwinner_{t-1}$) because $aligned_{c,t-1}$ is essentially an interaction between $partisan_{c,t-2}$ and $Rwinner_{t-1}$.

Estimates of equation (6), shown in column (1) of Table 6, 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, it does not support a model in which citizens learn about their probability of being pivotal.

A potential concern with regression specification (6) is that the inclusion of $partisan_{c,t-2}$

violates the strict exogeneity assumption required for fixed effects estimation, due to a relationship between turnout and vote shares.[31] To address these concerns, column (3) uses an alternative measure of $partisan_{c,t-2}$, 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.[32] 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 counties exhibit a high degree of habit formation, but dis-aligned 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 aligned counties relative to dis-aligned counties (equivalent to γ_1^a in equation(6)). The coefficient on the interaction of "aligned" with lagged precipitation remains significantly negative.

Notably, the asymmetry of our findings—aligned counties exhibit a higher degree of habit formation, but dis-aligned counties do not significantly differ from non-partisan counties—lends additional support to the importance of social interactions. Without social interactions, even if habit formation were asymmetric at the individual level—so that voting for the winner raised future voting propensity, but voting for the loser had no effect—county-level estimation would still exhibit symmetry. Specifically, dis-aligned counties contain fewer voters who support the winner than non-partisan counties, which in turn contain fewer such voters than aligned counties. As a result,

[31] 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.

[32] The results are not substantively different in specifications that omit the interactions of the South indicator with other covariates.

one would still expect to find greater habit formation in non-partisan counties than in dis-aligned counties. But the data do not confirm this prediction. One potential explanation is that spillovers from social interactions are especially pronounced in areas with fervent political beliefs following a victorious election. More generally, the county-level asymmetry likely reflects some form of social interactions in voting decisions.

7.2 Instrumental Utility (B)

A separate explanation for our results involves the strength of citizens' political preferences. If the act of voting causes an individual i 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 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.

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. Our test thus introduces interactions between precipitation and state-level voter pivotalness, in the following specification:

$$\begin{aligned} turnout_{ct} \;&=\; \beta_0 precip_{ct} + \beta_1 precip_{c,t-1} + \theta pivotal_{ct} \\ &\quad + \gamma_0 \left(pivotal_{ct} \times precip_{ct}\right) + \gamma_1 \left(pivotal_{ct} \times precip_{c,t-1}\right) + \tau_t + \eta_c + \lambda_c t + \varepsilon_{ct} \end{aligned} \tag{7}$$

The 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 $pivotal_{ct}$, 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.[33]

Table 7, which estimates equation (7), shows no evidence that the objective probability of being pivotal plays a role in our results. To ensure that sample selection is not affecting our estimates, 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 (7) using a continuous measure of state-level pivotalness, and none of the estimated parameters of interest are statistically significant.[34] 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 7 suggests that habit formation is not likely to operate through accumulation in the instrumental value of voting.

7.3 Voting Costs (C)

If P_{it} and B_{it} do not drive the results, then the Downsian framework leaves only C_{it} and D_{it}: the intrinsic costs and benefits of voting. Although these terms are theoretically distinct, they are difficult to distinguish empirically because neither is fully measurable. Nevertheless, we organize our discussion of these terms in two separate subsections. We begin with a hypothesis concerning C_{it}: 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

[33] 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.

[34] 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.

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 precipitation 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 have to be particularly salient in partisan counties, *but only when the preferred candidate wins.* We can think of little reason why learning about location, for example, would be stronger in mostly Republican counties after a Republican is elected than when a Democrat wins.

On the institutional side, state and county election offices have at various points implemented laws that purge inactive voters from the registration rolls. After the 1993 National Voter Registration Act (NVRA), automatic purges of non-voters ceased in all states. Nonetheless, while they were in effect, these laws could have produced "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 6, as there is no reason purging would vary by county partisanship with timing that matches the identity of the election winner. Second, we re-estimated our main specification using only the states and elections with no automatic purging of non-voters, finding results very similar to our benchmark findings.[35] These results are consistent with previous findings that the non-voter purges have negligible effects on turnout (Wolfinger and Rosenstone 1980, Mitchell and Wlezien 1995).

Note that we have only discussed costs incurred immediately prior to the act of voting. Longer-

[35] Data on automatic purging are available for 1960-2012 from *The Book of the States* (http://knowledgecenter.csg.org/kc/content/book-states-archive-1935-2009). Half of the observations in this period (and all of the observations since 1993) are from state-elections without automatic purging. 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).

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

7.4 Expressive Utility (D)

Only one Downsian term remains to explain our results, expressive utility D_{it}: the consumption value of voting, stemming from civic duty, ethics, or social pressure. Does accumulation in D_{it} drive habit formation? This hypothesis would conform with traditional interpretations of habit formation (Pollak 1970; Becker and Murphy 1988) in which "habits" are consumption tastes. The "vote for the winner" effects discussed on Section 7.1 does suggest that a psychological mechanism plays a role in habit formation. But because the concepts embedded in D_{it} are so difficult to measure, this hypothesis otherwise has mainly negative evidence in its favor, in the sense that we have ruled out all other explanations in the Downsian framework. To organize the discussion of this hypothesis, we follow the theoretical literature on habit formation (e.g., Chapman 1998) and separate accumulation in D_{it} into two types of processes: intrinsic and extrinsic.

Intrinsic accumulation in D_{it} refers to the individual-level psychological process by which citizens develop attachments to the act of voting, independent of social influences. The idea that pro-social or ethical behavior is self-reinforcing dates at least as far back as Aristotle, who wrote: "men become builders by building and lyreplayers by playing the lyre; so too we become just by doing just acts" (*Nicomachean Ethics* 2.1.32). A possible psychological micro-foundation for this type of accumulation is cognitive dissonance theory (Festinger 1957). In our context, this theory implies that a citizen would adapt her tastes regarding the importance of voting to create a consonance between actions and preferences: the act of voting would cause her to think that voting is a valuable contribution and an important duty.[36] However, other psychological mechanisms may also drive intrinsic accumulation in D_{it}.

Extrinsic accumulation in D_{it} occurs at the social level, with D_{it} responding to the community's voting history, not the individual's. Turnout is a socially visible activity, and people talk to one another about voting, so increases in aggregate turnout may affect a community's information,

[36] Mullainathan and Washington (2009) also study the role of cognitive dissonance in voting, but they focus on the choice of candidate, instead of the turnout choice.

attitudes, and norms about future voting. This class of mechanisms is consistent with the mounting evidence of social influences on the turnout decision (Nickerson 2008; Gerber et al. 2008; Bond et al. 2012; DellaVigna et al. 2013). It also receives support from the social interaction effect calculations in Section 6.2, as well as the vote-for-the-winner estimates in Section 7.1. Given this evidence, we infer that habit formation in voting comes from a mix of intrinsic and extrinsic accumulation in D_{it}.

7.5 What Role for Politics?

Until now, the discussion has treated voters as isolated, rather than as participants in an interactive political process. But actions by political elites may play a role, especially if rain-induced decreases in turnout have a partisan bias. If precipitation shocks affect election outcomes, and if incumbents are especially able to manipulate voter turnout, then the persistent effects of precipitation shocks may have a political explanation. Table 8 explores these issues. Column (1) repeats our main result for reference, while the remaining columns report new results relating to the role of politics.

Two sources of evidence in Table 8 suggest that our results are not primarily driven by politics. First, in column (2), neither current nor lagged precipitation has a significant effect on the Republican vote share in presidential elections. At face value, this result contradicts the finding by Gomez et al. (2007) and Hansford and Gomez (2010) that rainfall benefits Republican candidates. However, Gomez et al. do not include county fixed effects (or trends) in their specification, while Gomez and Hansford include county fixed effects but omit Southern counties from their analysis sample. In unreported results, we confirm in our preferred specification that contemporaneous precipitation raises the Republican vote share in non-Southern counties, but we also find an offsetting effect in the South, where contemporaneous precipitation *decreases* the Republican vote share. Lagged precipitation has no effect on vote shares in either region. In contrast, both regions exhibit the paper's main finding: lagged rainfall raises current turnout. Although these findings elaborate the previous literature, they do not point to a clear political explanation for our main result.

The second source of evidence casting doubt on political explanations comes from interactions between precipitation and national election characteristics. Columns (4)-(5) of Table 8 interact precipitation with the party of the incumbent President or with an indicator for whether one of the presidential candidates is the incumbent. None of the interaction terms are statistically significant, implying that the effect of lagged precipitation does not depend on the party of the incumbent or on whether the incumbent is running for re-election.

Beyond these results, we also note that the majority of precipitation shocks are not large enough to change election outcomes. The 90^{th} percentile of the precipitation distribution is 7.1 mm (Table 1), which given our estimates lowers turnout by approximately 0.5 percentage points. Most elections are won by substantially larger vote margins, especially in local races.[37] By this line of reasoning, the average effect of precipitation on who is elected (even in local races) would likely too small to be a plausible explanation for its sizable lagged impact on turnout.

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 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 these effects, finding that a 1 percentage point decrease in current turnout reduces future turnout by 0.7-0.9 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 occurs 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 finding evidence of habit formation in voting, this paper speaks to the potentially significant 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.

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

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.

Alesina, A., and N. Fuchs-Schundeln. (2007). "Good-Bye Lenin (or Not?): The Effect of Communism on People's Preferences." *American Economic Review* 97(4): 1507-1528.

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.

Aghion, P., Y. Algan, P. Cahuc, and A. Shleifer. (2010). "Regulation and distrust." *Quarterly Journal of Economics* 125(3): 1015-1049.

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. (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.

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.

Chapman, D.A. (1998). "Habit Formation and Aggregate Consumption." *Econometrica* 66(5): 1223-1230.

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

Chiang, C.F., and B.G. Knight, (2011)."Media Bias and Influence: Evidence from Newspaper Endorsements.". *Review of Economic Studies* 78(3): 795-820.

Coate, S., M. Conlin. (2004). “A group Rule-Utilitarian Approach to Voter Turnout: Theory and Evidence.” *American Economic Review* 94(5): 1476-1504.

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.

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.

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.

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

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.

Festinger, L. (1957). *A Theory of Cognitive Dissonance.* Evanston, IL: Row, Peterson.

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

Fisman, R., S. Kariv, and D. Markovits. (2009). “Exposure to ideology and distributional preferences.” Mimeo, Columbia University.

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., and D.P. Green. (2000). "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94(3): 653-663.

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.

Giuiliano, P., and A. Spilimbergo. (Forthcoming). "Growing Up in a Recession" *Review of Economic Studies.*

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.

Glaeser, E.L., B.I. Sacerdote, and J.A. Scheinkman. (2003). "The Social Multiplier." *Journal of the European Economic Association*, 1(2-3): 345-353.

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.

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.

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.

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.

Oberholzer-Gee, F., and J. Waldfoge. (2009). “Media Markets and Localism: Does Local News en Español Boost Hispanic Voter Turnout?” *American Economic Review* 99(5): 2120-28.

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

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): 15594-15598.

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.

Voors, M.J., E.E. Nillesen, P. Verwimp, E.H. Bulte, R. Lensink, and D.P. Van Soest. (2012). “Violent Conflict and Behavior: a Field Experiment in Burundi.” *American Economic Review* 102(2): 941-964.

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

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

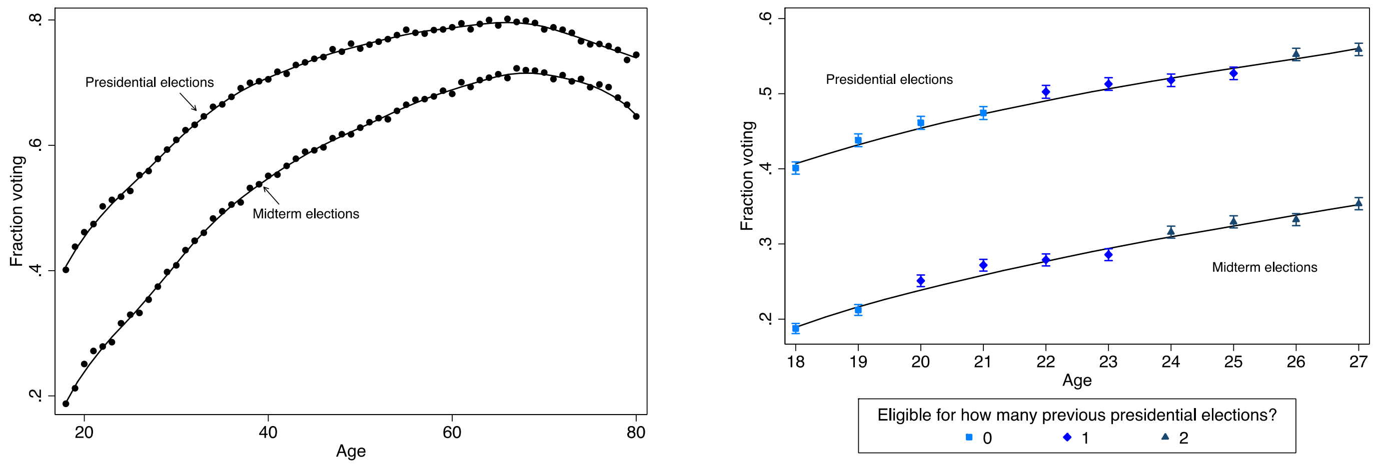

Note: Scatter plots are age-specific rates, while curves are local linear regressions with a bandwidth of 2 years.

Figure 2: Precipitation and Turnout Residuals, 2004

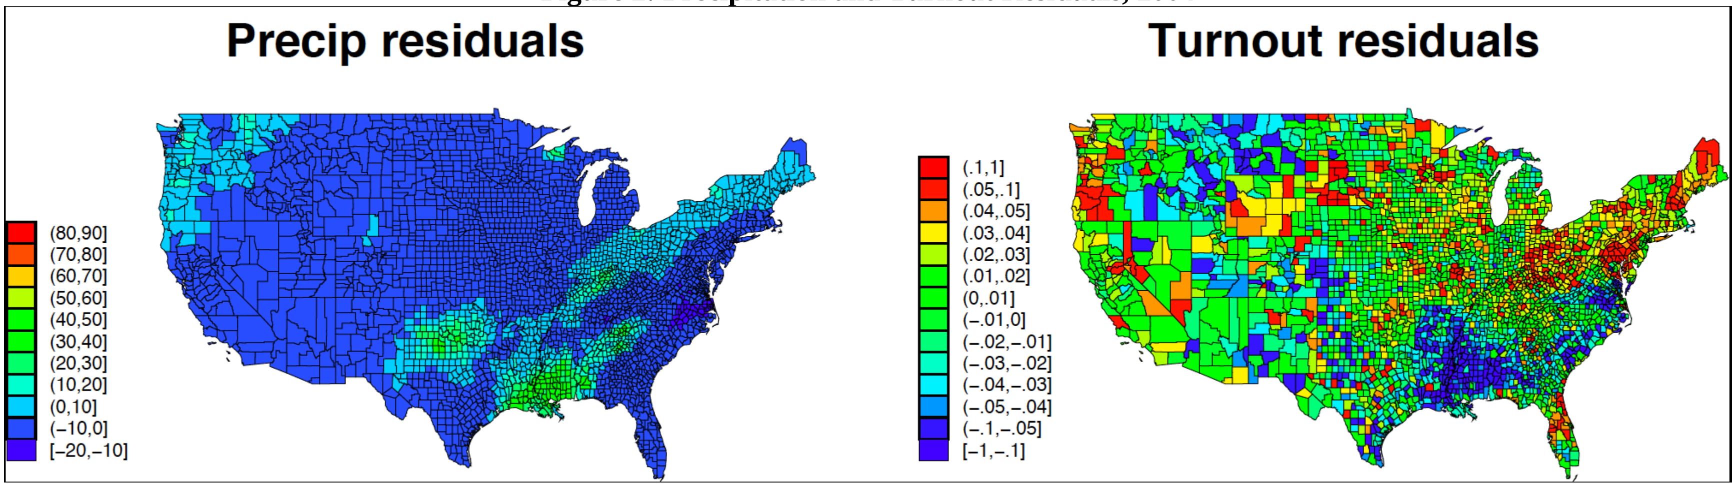

Note: Residuals from regressions of precipitation (mm) and turnout on year and county fixed effects and county trends.

Figure 3: Effect of Contemporaneous and Lagged Precipitation on Turnout: Binned Estimates

A. Precipitation Levels

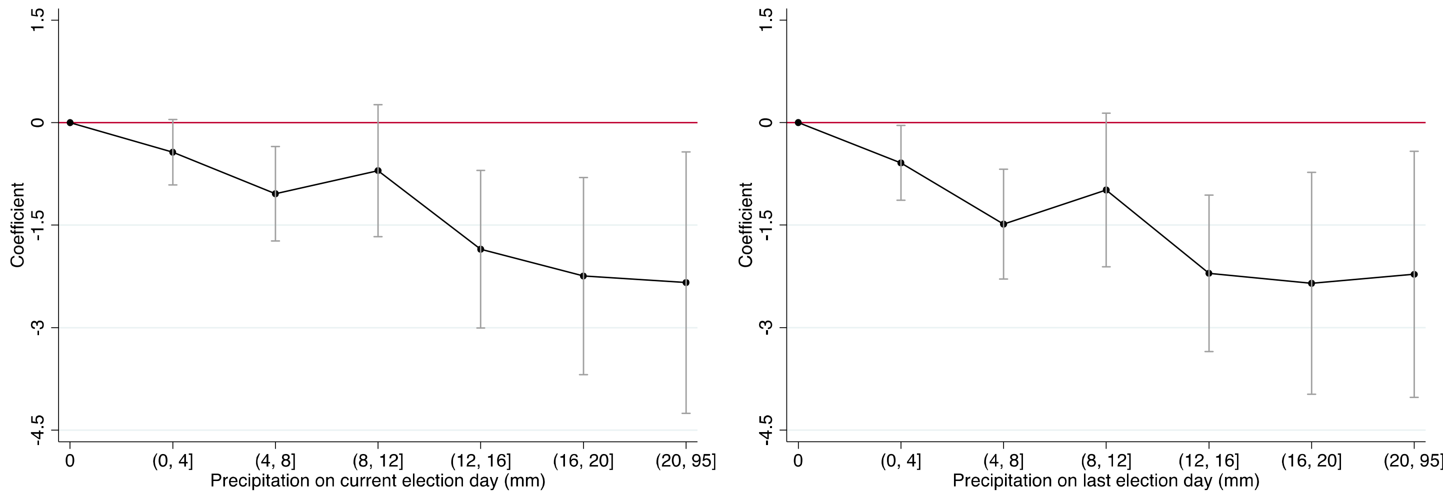

B. Local Precipitation Quintiles

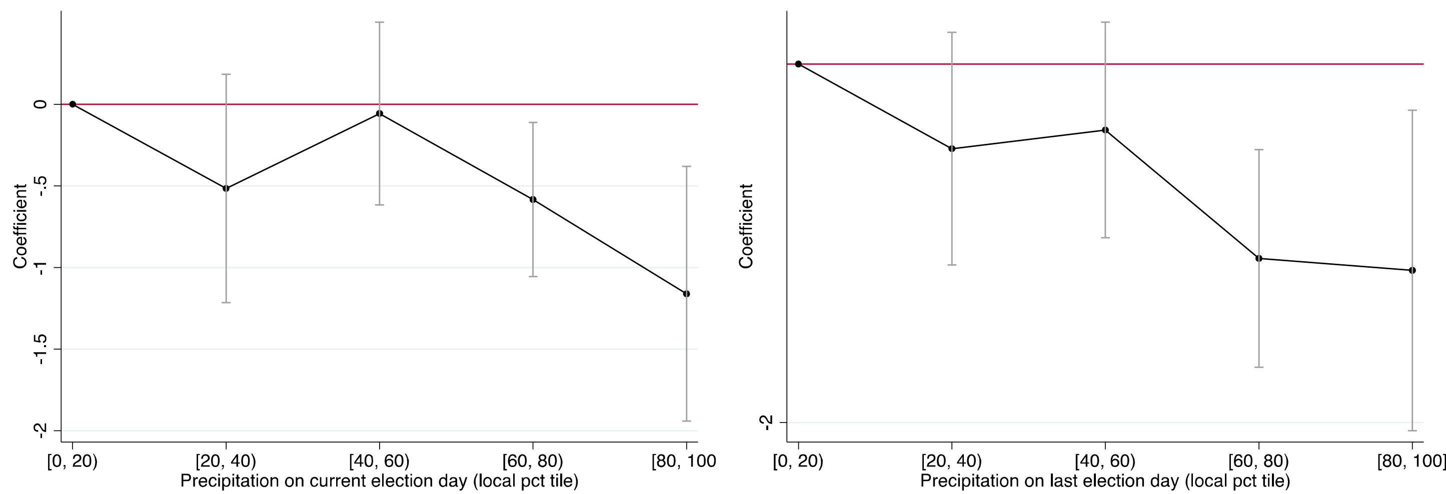

Note: Dep. var. is voter turnout (0-100). Capped spikes represent 95% CIs based on SEs clustered at the state level. Both regressions include year fixed & county FE and county trends.

Table 1: Descriptive Statistics

	Mean	Std. Dev.	Percentiles 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							
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) Bin model	(4) 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]
(0,4] mm precip. on election day, t		-0.350		
		[0.230]		
(4,8] mm precip. on election day, t		-0.860		
		[0.350]**		
(8,12] mm precip. on election day, t		-0.644		
		[0.474]		
(12,16] mm precip. on election day, t		-1.701		
		[0.537]***		
(16,20] mm precip. on election day, t		-2.011		
		[0.684]***		
(20,95] mm precip. on election day, t		-2.183		
		[0.906]**		
[20^{th},40^{th}) local %-tile, t			-0.440	
			[0.238]	
[40^{th},60^{th}) local %-tile, t			0.007	
			[0.267]	
[60^{th},80^{th}) local %-tile, t			-0.457	
			[0.129]**	
$\geq 80^{th}$ local %-tile, t			-1.070	
			[0.383]***	
Number of county-years	49,524	49,524	49,524	49,524
Number of counties	3108	3108	3108	3108
Election years	1952-2012	1952-2012	1952-2012	1952-2012

Note: Dependent variable is voter turnout (0-100). Brackets contain standard errors clustered at the state level. All regressions include year and county fixed effects and county 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)
Precip. on election day, t+1				-0.027	
				[0.035]	
Precip. on election day, t	-0.078	-0.070	-0.077	-0.110	-0.079
	[0.026]***	[0.025]***	[0.025]***	[0.037]***	[0.026]***
Precip. on election day, t-1	-0.070	-0.065	-0.068	-0.112	-0.069
	[0.025]***	[0.024]**	[0.023]***	[0.041]***	[0.025]***
Precip. on election day, t-2				-0.057	
				[0.023]**	
Precip. 2 wks. after election day, t+1				-0.021	
				[0.026]	
Precip. 2 wks. after election day, t			-0.029	-0.008	
			[0.038]	[0.039]	
Precip. 2 wks. after election day, t-1			-0.018	-0.008	
			[0.044]	[0.048]	
Precip. 2 wks. after election day, t-2				-0.036	
				[0.034]	
Precip. on midterm election day, t-½					0.015
					[0.012]
Implied ρ	0.89	0.93	0.88	0.75	0.87
	[0.28]***	[0.31]***	[0.28]**	[0.14]***	[0.28]***
Number of county-years	49,524	49,524	49,524	43,300	49,524
Number of counties	3,108	3,108	3,108	3,108	3,108
Election years	1952-2012	1952-2012	1952-2012	1956-2008	1952-2012
q	No	Yes	No	No	No

Note: Dep. var. is voter turnout (0-100). Brackets contain SEs clustered at the state level. All regressions include year fixed & county FE and county trends. County covariates: white pop. share, over-65 pop. share, log median income, log pop. density. * p<0.1, ** p<0.05, *** p<0.01

Table 4: Interactions with County Characteristics

	(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: 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 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: 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?	No	Yes	Yes	No	Yes	Yes	Yes
Year FE?	No	Yes	Yes	No	Yes	Yes	Yes
State-group/state FE?	No	No	Yes	No	No	Yes	No
County FE?	No	No	No	No	No	No	Yes

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 white pop. share, over-65 pop. share, log median income, and log pop. density.

Table 6: 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: Dep. var. is voter turnout (0-100). Sample includes presidential elections from 1956-2012. 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 FE and county trends. * p<0.1, ** p<0.05, *** p<0.01

Table 7: Role of Predicted Pivotalness

	(1)	(2)	(3)
Precip. on election day, t	-0.109	-0.118	-0.123
	[0.035]***	[0.045]**	[0.040]***
Precip. on election day, t-1	-0.079	-0.077	-0.115
	[0.027]***	[0.023]***	[0.026]***
(Pivotal, t) × (Precip., t)		7,869	
		[21,665]	
(Pivotal, t) × (Precip., t-1)		-1,224	
		[10,342]	
(Pivotal dummy, t) × (Precip., t)			0.038
			[0.040]
(Pivotal dummy, t) × (Precip., t-1)			0.074
			[0.039]*
Pivotal, t		35,486	
		[77,938]	
Pivotal dummy, t			0.069
			[0.431]
Implied ρ (main)	0.724	0.654	0.938
	[0.195]***	[0.216]***	[0.299]***
Implied ρ (dummy interaction)			0.483
			[0.433]
Number of county-years	42875	42875	42875
Number of counties	3,107	3,107	3,107
Election years	1952-2004	1952-2004	1952-2004
Pivotal interaction	None	Continuous	Above median

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 uninteracted effect of pivotalness. The continuous measure of pivotalness is based on Campbell et al. (2006); see footnote 29 for more information. The "pivotal dummy" measures whether the continuous measure is above or below its median. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Precipitation and Politics

	Turnout	Rep. share	Turnout	Turnout	Turnout
	(1)	(2)	(3)	(4)	(5)
Precip. on election day, t	-0.078	-0.022	-0.091	-0.074	-0.059
	[0.026]***	[0.027]	[0.035]**	[0.028]***	[0.048]
Precip. on election day, t-1	-0.070	-0.041	-0.058	-0.037	-0.100
	[0.025]***	[0.034]	[0.029]**	[0.027]	[0.033]***
(Nat'l margin, t-1) × (Precip., t)			0.002		
			[0.003]		
(Nat'l margin, t-1) × (Precip., t-1)			-0.001		
			[0.002]		
(Rep. incumbent, t) × (Precip., t)				-0.004	
				[0.042]	
(Rep. incumbent, t) × (Precip., t-1)				-0.070	
				[0.052]	
(Re-election, t) × (Precip., t)					-0.027
					[0.046]
(Re-election., t) × (Precip., t-1)					0.050
					[0.036]
Number of county-years	49,524	49,441	49,524	49,524	49,524
Number of counties	3108	3108	3108	3108	3108

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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

APPENDIX TABLES AND FIGURES – NOT FOR PUBLICATION

Appendix Figure 1: Share of Counties with Election-Day Precipitation by Year

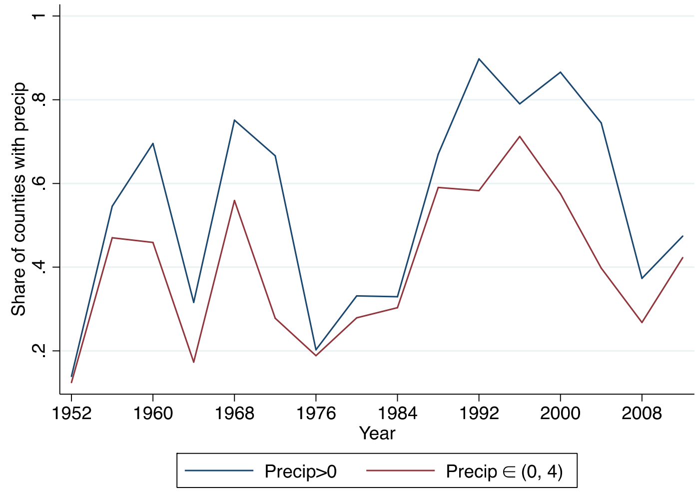

Appendix Figure 2: Cumulative Share of Counties with Election-Day Precipitation

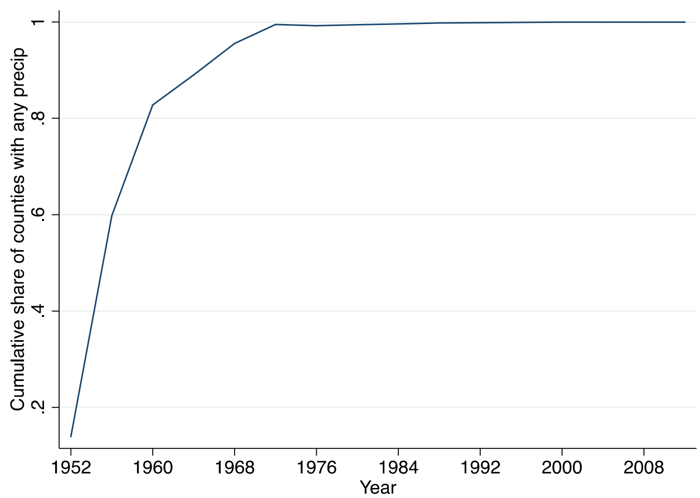

Appendix Figure 3: Histogram of Standard Deviation of Precipitation

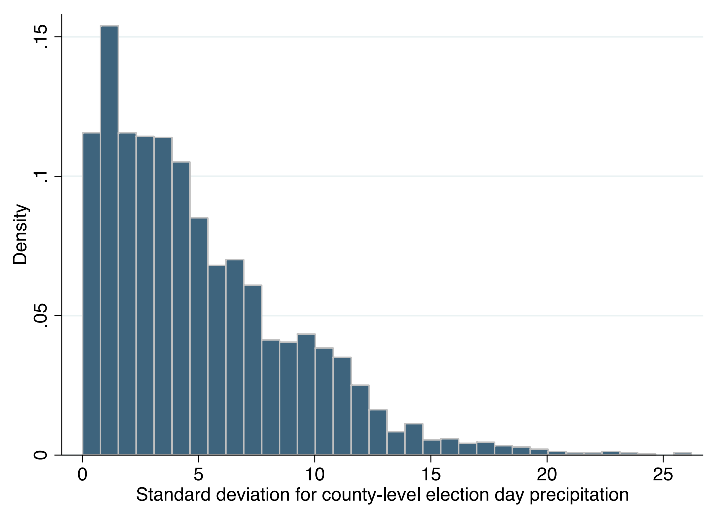

Appendix Figure 4: Leave-One-Out Checks

Panel A: Leave Out One State

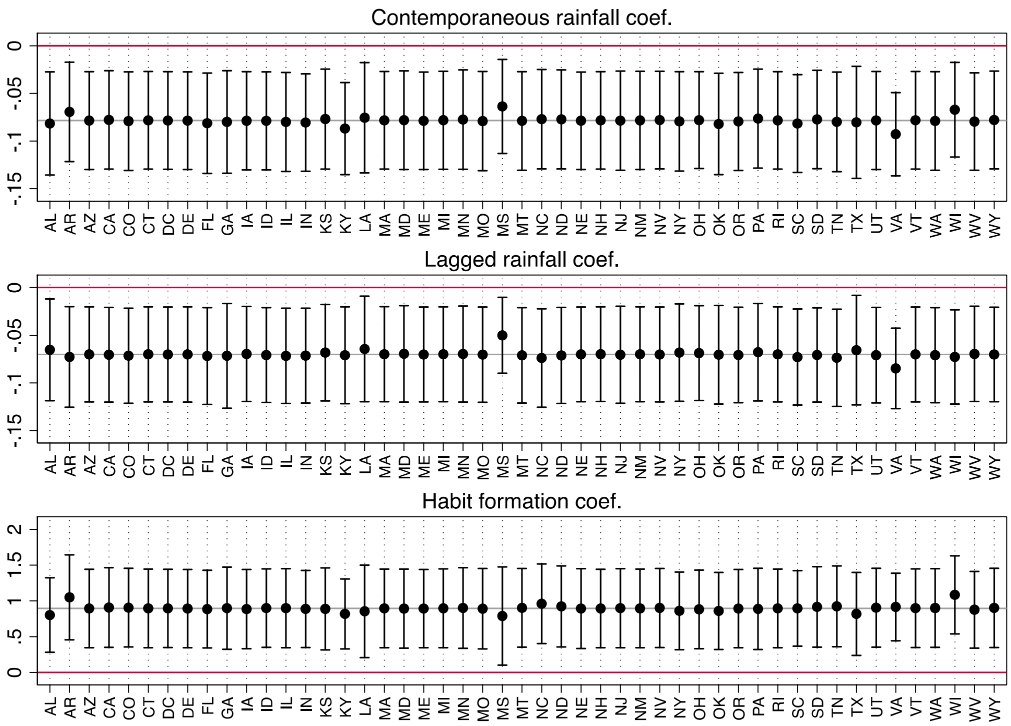

Panel B: Leave Out One Year

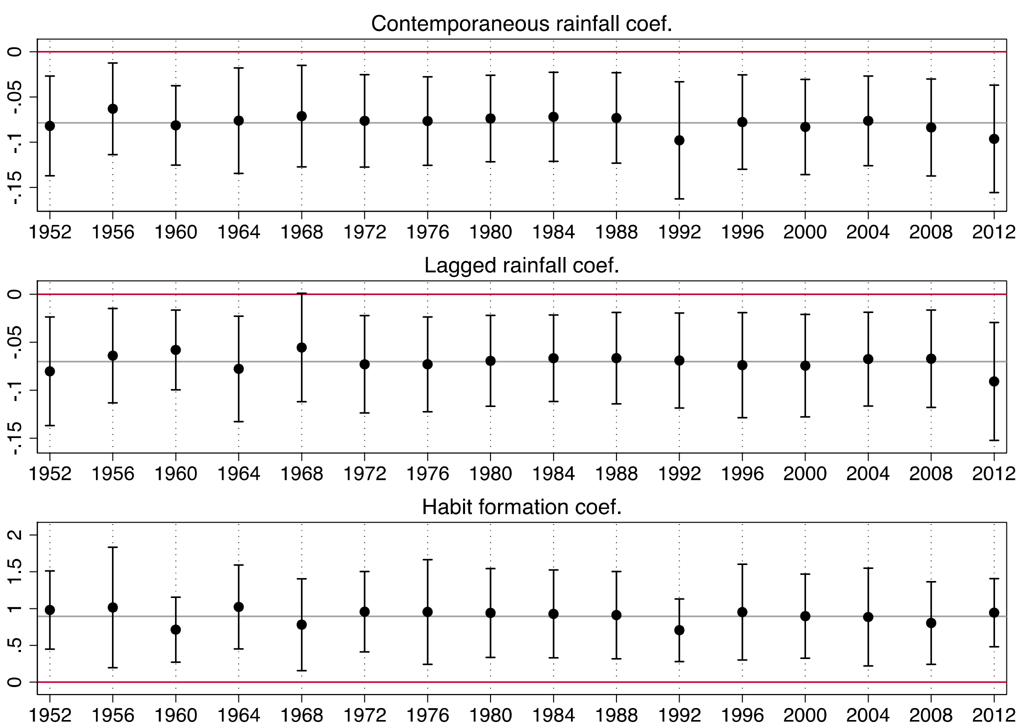

Note: Each estimate is based on a sample that omits the state or year on the x-axis. Dots are coefficients; bars are 95% CIs. Light gray horizontal lines represent full-sample estimates.

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 ρ	-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. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$