

Economic Distress and Voting: Evidence From the Subprime Mortgage Crisis*

Andrew B. Hall[†]
Department of Political Science
Stanford University

Jesse Yoder[‡]
Department of Political Science
Stanford University

Nishant Karandikar[§]
Department of Management Science and Engineering
Stanford University

April 24, 2017

Abstract

Roughly 7 million Americans lost homes to foreclosure during the Great Recession. Despite claims that the subprime mortgage crisis helped fuel recent political turmoil in the U.S., we lack systematic empirical evidence about the effects of this unprecedented spike in home foreclosures on American elections. We combine nationwide deed-level public records data on home foreclosures with election data and administrative voter data to examine the effects of home foreclosures on electoral outcomes and on individual voter turnout. At the aggregate level, county-level difference-in-differences estimates show that counties that suffered larger increases in foreclosures did not punish or reward members of the incumbent president's party more than less affected counties. At the individual level, merging the Ohio voter file with foreclosure data, difference-in-differences estimates reveal that Ohioans whose homes were foreclosed on were somewhat *less* likely to turn out to vote, particularly when foreclosures occurred close to election day. The findings cast doubt on the claim that individual-level economic distress during the Great Recession directly activated angry voters, and raise questions about the posited causal link between economic distress and the electoral punishment of incumbents.

*For comments and suggestions, the authors thank Avi Acharya, Anthony Fowler, Judy Goldstein, Justin Grimmer, Gabe Lenz, John Marshall, Pablo Montagnes, Ken Scheve, Mike Tomz, and participants of the Governing the Global Economy Faculty Workshop at Stanford. For generously sharing data, the authors thank Andy Healy and Gabe Lenz. The protocol for this study has been approved by the Stanford IRB (#39468).

[†]Andrew B. Hall is an Assistant Professor in the Department of Political Science at Stanford University (andrewbenjaminhall@gmail.com, <http://www.andrewbenjaminhall.com>).

[‡]Jesse Yoder is a Ph.D. Student in the Department of Political Science at Stanford University.

[§]Nishant Karandikar is a Masters student in the Department of Management Science and Engineering at Stanford University.

1 Introduction

How do voters hold politicians accountable for their actions? One of the longest-running literatures in political economy links the state of the economy to incumbents' reelection prospects, suggesting that voters reward incumbents when the economy performs well and punish them when it performs badly.¹ American presidential and congressional elections particularly seem to follow this pattern (Fair 1978; Kramer 1971), and so do a number of other electoral contexts across the world (Brender and Drazen 2008; Lewis-Beck 1986). Yet the mechanisms underlying this statistical association remain poorly understood, and there is little causal evidence from real-world elections about this behavior at either the level of the individual voter or in the aggregate.² Does personal economic hardship cause voters to punish their incumbent politicians? Does it drive individuals to become more politically active?

Questions about economic performance and electoral outcomes are at the heart of a large literature in political economy precisely because they matter a great deal for understanding democratic politics. The subprime mortgage crisis, whose political effects we study in this paper, affected the lives of many millions of Americans and, like other large-scale economic crises in the past, has led many to speculate about its role in subsequent political upheaval. A Fox News opinion piece, for example, recently declared that “poor access to housing in America, and the lingering aftereffects of the 2008 housing crash, remains a chief motivator of anger and dissatisfaction amongst an electorate that views the economy and the mortgage market as rigged against ordinary people.”³ Have the people most affected by the housing crash—those whose homes were foreclosed, and those who live in areas where foreclosures were most prevalent—actually changed their voting behavior? Has the economic hardship they experienced during the Great Recession activated them, politically? These are the key questions we study in this paper.

¹This literature goes at least as far back as Gosnell and Colman (1940), in political science, as far back as Pearson and Myers (1948), in economics, and as far back as Tibbitts (1931) in sociology. As of 2007, over 400 articles and books were written about economic voting (Lewis-Beck and Stegmaier 2007), and many more have been written since. For a recent review of some of the literature, see Bagues and Esteve-Volart (2016).

²There is some experimental work using lab participants and/or survey respondents, which we discuss below. Also, a large literature in political science attempts to separate individual-level “pocketbook” voting, in which individuals vote on the basis of their personal economic situation, from “sociotropic” voting, in which individuals vote on the basis of society’s economic situation (see for example Kinder and Kiewiet 1981; Kramer 1983). This literature largely relies on cross-sectional surveys and does not isolate exogenous variation in economic conditions at either the individual or aggregate level.

³<http://www.foxnews.com/opinion/2016/07/16/want-to-understand-voter-anger-in-2016-don-t-overlook-housing.html>

On the one hand, the literature on economic and retrospective voting implies that these individuals should be expected to respond to economic hardship by punishing incumbents. Healy and Lenz (N.d.), a paper we discuss in more detail below, presents evidence for precisely this kind of link at the local level, studying U.S. presidential elections. In so doing, the paper follows a long tradition in political economy and political science that connects economic hardship or prosperity to incumbent punishment or reward. Nordhaus (1989), for example, presents a model that proposes an individual-level retrospective voting rule. Much work in this literature goes further, granting the premise that voters punish incumbents for bad economic performance, but also arguing that voters sometimes misapply this retrospective voting rule, rewarding or punishing incumbents for outcomes that are in fact outside of the incumbents' control. Wolfers (2007) links incumbent governors' electoral performance to oil price shocks outside of gubernatorial control. Bagues and Esteve-Volart (2016) offers experimental evidence that winning the lottery modestly increases support for incumbent politicians in Spain. Huber, Hill, and Lenz (2012) uses a series of lab experiments to suggest several ways in which hypothetical voters might overweight or inappropriately evaluate economic information when evaluating incumbents. In a similar vein, studying observed election returns, Healy, Malhotra, and Mo (2010) suggests that voters reward incumbents for college football wins and punish them for losses, and Healy and Malhotra (2009) makes similar arguments concerning natural disasters (though also see Ashworth, Bueno De Mesquita, and Friedenbergr 2016; Gasper and Reeves 2011). And perhaps most dramatically, Achen and Bartels (2016) suggests that voters punish incumbents for shark attacks and droughts. Some of these findings have been questioned in follow-up work (Fowler and Montagnes 2015; Fowler and Hall 2016), but either way, the literature is fairly unambiguous in predicting that voters should be expected to punish incumbents for personally bad economic outcomes like home foreclosures, whether rationally or not.

On the other hand, we might be skeptical that the individuals most affected by the subprime crisis have viewed political activity as a useful response to their situation. Individuals who experience economic distress receive no increase in the probability their vote is pivotal, and at the same time may experience higher costs from voting, especially in terms of opportunity costs related to the cognitive requirements of voting and to time spent not addressing pressing personal economic issues, transaction costs related to the practical requirements of voting, particularly if they've changed addresses, as well as potential psychological costs. Moreover, the people most

likely to experience economic distress are people with lower income and lower education levels, people who we know are already less likely to be politically engaged (e.g., Verba, Scholzman, and Brady 1995) and less sensitive to interventions encouraging them to become politically engaged (Enos, Fowler, and Vavreck 2014). Economic distress related to individuals' housing situation may be especially discouraging, as previous work suggests that "residential mobility...may strain poor Americans' weak ties to the political system" (Gay 2012: abstract).

In addition, there are doubts over the level of aggregation at which we should observe economic voting, if we observe it. The literature on sociotropic voting argues that it is only visible at the national level, either because voters who vote for economic reasons do it with the entire country's welfare in mind, rather than their own personal circumstances, or because they use the national economy as a forecast of their future wellbeing. Surveying the sociotropic voting literature, for example, Kiewiet and Lewis-Beck (2011: 305) concludes: "[voters'] decisions were strongly influenced by their assessments of how the nation's economy had been performing." As Ansolabehere, Meredith, and Snowberg (2014) argues, however, voters may obtain signals about the macroeconomy at the local level, in which case we should expect to see local variation in incumbent support that is correlated with local variation in economic conditions. Also in contrast to the sociotropic view, the retrospective voting literature cited above also offers many cases where voters seem to respond to events at the local level—like college football games, droughts, locusts, and shark attacks.

We might also be skeptical about the general economic voting literature's one-sided focus on voters. In particular, we might suspect that voter behavior in response to crises depends on the strategic response of politicians, so that predictions about voter behavior that do not consider this strategic interplay will be misleading (e.g., Ashworth and Bueno De Mesquita 2014). Moreover, the differences in economic policy—particularly as it pertains to financial regulation and the housing market—between Democrats and Republicans, and between incumbents and challengers, are not always so clear, and may adapt to changing conditions. Indeed, neither the Democrats nor the Republicans were shy in publicly criticizing the banking system in the aftermath of the subprime crisis, previous legislation to expand access to housing passed with broad partisan support,⁴ and financial support from Wall Street moved fluidly from the Republican party to the Democratic

⁴The Housing and Community Development Act of 1992, for example, passed the House 377-37 and was passed by voice vote in the Senate. <https://www.govtrack.us/congress/bills/102/hr5334>.

party after Barack Obama’s election in 2008.⁵ If the two parties do not manage the financial world or housing markets differently, it is hard to see why voters have any rational reason to punish incumbents or switch parties in the aftermath of home foreclosures or related economic crises. Objections like these are at the heart of Stigler’s (1973) general skepticism towards the theory of economic voting.

To understand the electoral effects of the housing crisis, and to reconcile these countervailing predictions, we use nationwide public records data on individual home foreclosures, combined with county-level election data for the U.S. House, Senate, and President. We also use administrative data on all registered voters in the state of Ohio, which allows us to merge home foreclosures to individual voter records to examine effects on turnout.⁶ To address obvious concerns that the people and places where foreclosures are more likely to occur are unlike the people and places where they are less likely to occur, we employ a series of difference-in-differences designs. Overall, we find precisely estimated null effects on county-level vote share outcomes, and we find small but detectable decreases in the propensity to turn out to vote after a home foreclosure. These decreases grow in magnitude for foreclosures that occur closer to election day.

Our findings raise questions about the aggregate patterns documented in the economic voting literature. We find no evidence that voters punished incumbents because of home foreclosures, or that foreclosures activated individuals and turned them into new voters. We consider several possible ways to explain these patterns.

One possibility is that most voters have strong and polarized ideological preferences, and/or have strong partisan loyalties. These forces could outweigh other incentives to switch to supporting a candidate of the other party. If this were the case, we might expect to find different effects if we zoom in on highly competitive counties, where there may be a larger number of swing voters who are less committed to one party or the other. When we do this, we continue to find precisely estimated null effects on electoral outcomes.

Another possibility is that, for most of the elections we study, the parties were indeed too similar in their positions—the “Stigler hypothesis”—so that there is no reason for affected individuals

⁵For example, Goldman Sachs was the second largest contributor to Obama’s 2008 campaign. See <https://www.opensecrets.org/pres08/contrib.php?cid=N00009638>.

⁶Although both home foreclosure information and the Ohio voter file are matters of public record, we are aware of the sensitive nature of this information. After merging records, we remove all personally identifying information from the dataset. Our procedure was approved by Stanford’s IRB.

to change votes or become more politically active. We find modest support for this hypothesis; although overall we find null results on home foreclosures and votes for the incumbent party, 2016 appears to be an exception to this pattern. Areas suffering from disproportionate amounts of home foreclosures between 2012 and 2016 appear to have increased their vote for Donald Trump and decreased their vote for Hillary Clinton. Because the narrative of the 2016 campaign seems to be that Trump offered a truly novel set of economic policy proposals—while Clinton was perhaps viewed as a conventional candidate—these vote changes may indicate the effect of economic distress on voting when the candidate’s positions are dissimilar. That being said, the effects are modest in size, and are unlikely to represent a meaningful share of the variation in partisan vote shares across elections.

Finally, it may be that individual voters by and large choose not to blame incumbents for personal economic hardship, at least when it comes in the form of home foreclosures. Between the null results on vote choice at the county level, and the negative results on the choice to turn out at the individual level, we can at least say that the relationship between personal economic circumstances and political activity is nuanced and quite muted. Economic upheaval can no doubt bring systemic change to politics, but it seems to have surprisingly little effect on most individuals’ political behavior. Whether the president, senators, and members of Congress deserved blame for the rash of home foreclosures is a matter of debate, and we take no position on whether rational voters affected by foreclosure ought to have punished their incumbents or not. Either way, our findings are at odds with large literatures predicting that individual voters will punish incumbents for these kinds of personal economic difficulties.

Our paper is most closely related to a recent study local economic conditions and voter behavior, as we mentioned earlier. Healy and Lenz (N.d.) studies zipcode-level mortgage delinquencies in California and relates these to changes in support for the Democratic presidential candidate between 2004 and 2008, finding marked increases in support for Obama in places where the housing crisis was more severe.⁷ Synthesizing the data from both studies, we assess the overall evidence for the link between local economic conditions and incumbent electoral fortunes. Using our county-level data on foreclosures, we replicate the California finding from Healy and Lenz (N.d.); if we zoom

⁷Hill, Herron, and Lewis (2010) also estimates some relationships between mortgage delinquencies and support for Obama, finding mixed results, including some evidence for a positive relationship but only conditional on high wage areas.

in on California in the years 2004 and 2008, we, too, find a modest but positive link between foreclosures and changes in support towards Obama. Because of our nationwide deed-level data on foreclosures, our paper is able to extend this analysis of the housing crisis to other states and other time periods. When we do so, we do not generally find the same kind of pattern, instead finding precisely estimated null results for the effect of foreclosures on incumbent electoral performance. We also find null results when we focus only on 2008, like in the Healy and Lenz (N.d.) analysis on mortgage delinquencies, but extend the analysis to include the five states other than California who experienced the largest changes in foreclosures before the 2008 election (see Table A.1 and further discussion in the Appendix). As a result, we conclude that California in 2008 is a special and unusual case.

Healy and Lenz (N.d.) also presents nationwide results linking local economic conditions, measured using changes in county-level wages and employment, to presidential incumbent vote. In the Appendix, we present a detailed analysis of our data and the nationwide data from Healy and Lenz (N.d.). When we use our preferred difference-in-differences design in place of the lagged dependent variable strategy from Healy and Lenz (N.d.), we find precisely estimated null effects in both datasets. We also consider the lagged dependent variable specification that Healy and Lenz (N.d.) employs instead of the difference-in-differences. Using the lagged dependent variable and difference-in-differences estimates on each dataset to bracket the true effect (Angrist and Pischke 2009), we conclude that the link between local economic conditions and incumbent electoral outcomes is, at most, substantively small. See the Appendix for more details.

Although we stress that our design does not evaluate system-wide effects of the housing crisis, writ large—it only estimates county-level and individual-level effects—these findings may nonetheless also help explain why politicians with local reelection concerns have not gone to greater lengths in responding to the subprime crisis. In a recent review of financial reform after the crisis, for example, Duffie (N.d.) discusses the “four core elements” that the Financial Stability Board listed to G20 leaders, and concludes that “only [one] can be scored a clear success, although even here much work remains to be done.” Reelection-minded incumbents in distressed areas may have had little incentive to make major changes when the constituents most affected by the crisis remained politically disengaged after the crisis. Although the crisis may have increased these citizens’ incentives

to become politically engaged, their personal economic challenges during the crisis likely acted to counteract this incentive, leaving them at least as politically disengaged after the crisis as before.

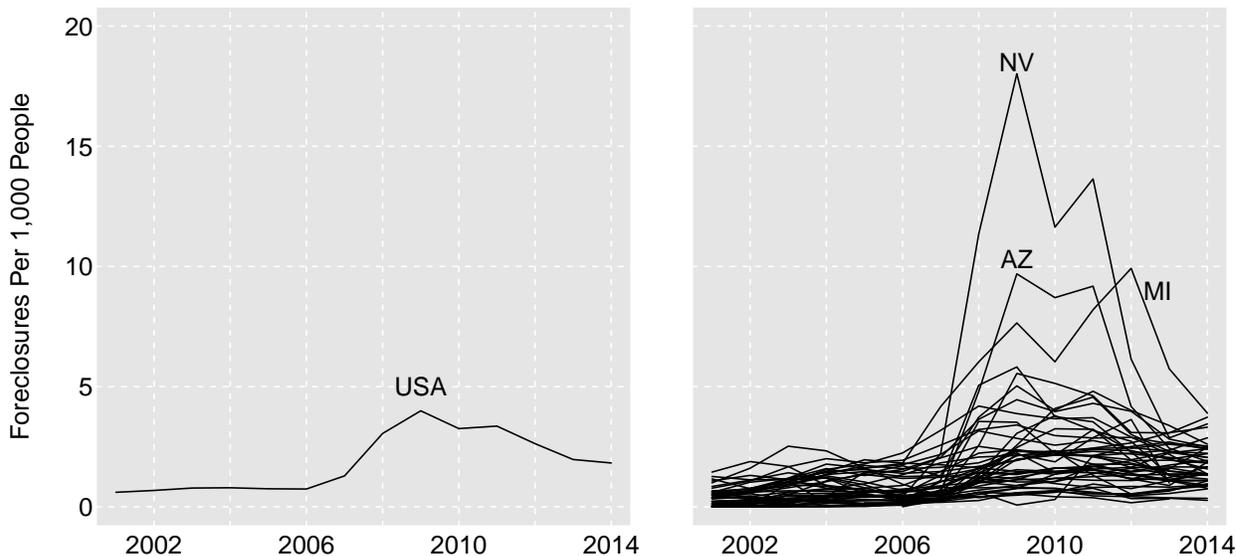
This points to a broader idea concerning the link between economic hardship and political behavior. The individuals most affected by economic downturns are probably unlikely to participate in politics, precisely because they are so affected by economic downturns. People facing default, home foreclosure, and unemployment may have the most reason to be angry with the political system, but they also have personal problems more pressing than casting a non-pivotal vote in an election between candidates they know little about. In the absence of a strong coordinating mechanism that could convince these individuals to yoke their political activities to one another, any one of them has little reason to become politically active. Though politicians and other sophisticated actors may have incentives to claim the mantle of voter anger, those with the most reason to be angry are unlikely to be engaged in post-crisis politics.

2 Studying the Effects of Home Foreclosures

Our data on home foreclosures comes from CoreLogic, a company that compiles information on the housing market culled from public records. The dataset covers foreclosures in the years 2000–2014. After limiting our attention only to deeds held by individuals, we define a foreclosure as any deed that is recorded as entering into a foreclosure, marked in varying cases by a real-estate owned transfer or sale, or by a variety of foreclosure certificates. We do not count a property as foreclosed based on pre-foreclosure events like notices of default or lis pendens. The Appendix provides detailed information on how we process the CoreLogic data. Because CoreLogic data is only available for about half of all U.S. counties, in the Appendix, we also discuss reasons to believe that our results generalize to the full set of U.S. counties. In particular, Table A.3 shows that CoreLogic counties look just like the full set of U.S. counties in terms of important variables like median household income, education, and unemployment.

The events leading up to the Great Recession are well known. Starting in 2007, a large number of home owners defaulted on their mortgages, triggering the failure of several mortgage originators and high-profile financial institutions. During the subprime crisis, foreclosure rates skyrocketed. Figure 1 plots the national and state-by-state average foreclosure rates per 1,000 people over time.

Figure 1 – U.S. Home Foreclosures Over Time, CoreLogic Counties. The left panel shows the national home foreclosure rate over time; the right panel shows the foreclosure rate by state. Both are measured using CoreLogic data, which cover roughly half of U.S. counties, so trends and state levels may vary from nationwide foreclosure data.



The years 2007–2009 saw an unprecedented spike in foreclosures, with especially dramatic increases in Nevada and Arizona, but with increases in many other states as well.

These foreclosures had tremendous consequences. Individuals who experience foreclosure obviously experience significant economic distress, both in the lead-up to foreclosure and in its aftermath. But foreclosures also have substantial spillover effects on local economies. Mian, Sufi, and Trebbi (2015) estimate that foreclosures lead to substantial decreases in home prices, residential investment, and consumer demand at the local level. By studying the effects of foreclosures on political behavior, we are therefore able to study two closely related issues. First, we can study the local effects of foreclosures on aggregated vote choice. Here, foreclosures may exert their own force on electoral outcomes, and they are also likely a proxy for the more general economic downturn that local areas experienced. This is important because, much like other forms of economic distress like unemployment, foreclosures on their own do not affect nearly a large enough number of people to change overall electoral outcomes.⁸ If foreclosures drive aggregate political outcomes, it

⁸Considering the case of unemployment, (Stigler 1973: 162) offers a hypothetical 3 percentage-point change in unemployment, and writes: “The 3 percent of unemployed, largely concentrated in the young, unmarried, and less educated, has substantially less than 3 percent of the vote.”

likely depends on the more widespread economic effects that foreclosures catalyze. Second, within counties, we can study the effects of home foreclosures on individual voter participation.

For the aggregate analysis on vote share, we sum up the total number of individual foreclosures in a given county in each year, and we divide by the county’s population in that year, as estimated from the Census. We linearly interpolate county population for non-Census years. We then take the average foreclosure rate within an electoral cycle—a four-year period, when we study presidential elections, and a two-year period, when we study legislative elections (House and Senate). This information is then merged, by county and election period, with data on county-level vote share for these offices, which comes primarily from Dave Leip’s Atlas of U.S. Presidential Elections.⁹

Table 1 presents descriptive statistics for these county-level datasets. Panel A covers the presidential elections dataset, where each time period is a four-year presidential term; Panel B covers the legislative elections dataset, where each time period is a two-year congressional term. Because we do not have foreclosure data past 2014, we do not use 2016 legislative elections in our analyses, but we keep the 2016 presidential elections in the analyses by using foreclosures from 2012–2014. In addition to raw means, the panels also show population-weighted means, which accounts for the well-known fact that Republicans dominate the low-population rural counties of the U.S. while Democrats dominate the high-population urban counties.

On average, counties experience 1.10–1.11 foreclosures for every 1,000 people (column 1, Panels A and B). As column 2 shows, these averages go up when we weight by population—foreclosures are somewhat more common in more populous, urban areas. As these numbers show, foreclosed individuals are a small fraction of the population, an important fact when considering potential electoral effects at the aggregate level. Foreclosure rates run from a minimum of 0—in certain low population counties—to a maximum of 18.76 over a four-year period (Panel A). This maximum value comes from Archuleta County, Colorado, over the four-year period from 2008 to 2012. When we zoom in on two-year periods in Panel B, we find an even higher maximum (19.25 foreclosures per 1,000 people), which reflects Archuleta County in the period 2010–2012.

Panel A also shows average Democratic presidential vote share. Columns 1 and 2 show the urban-rural divide quite clearly. The raw average Democratic vote share is only 40%; but when

⁹County-level vote shares for the November 2016 election are not yet available through Dave Leip’s Atlas of U.S. Presidential Elections, so 2016 vote shares for presidential races come from Politico’s reported results (see https://github.com/Prooffreader/election_2016_data).

Table 1 – Descriptive Statistics, County Level, 2004–2016.

	Mean (1)	Pop-Weighted Mean (2)	Standard Deviation (3)	Minimum (4)	Maximum (5)	Observations (6)
A. Presidential Elections						
Foreclosure Rate	1.11	1.72	1.37	0.00	18.76	4,411
Dem Vote Pct	39.45	52.41	14.21	4.95	87.17	4,411
County Population	126,759	–	460,813	401	10,139,013	4,411
B. Legislative Elections						
Foreclosure Rate	1.10	1.77	1.48	0.00	19.25	7,344
Dem Senate Vote Pct	28.33	32.92	25.03	0.00	100.00	7,344
Dem House Vote Pct	39.67	49.58	20.95	0.00	100.00	7,344
County Population	130,364	–	467,695	385	10,077,306	7,344

Note: Panel A presents descriptive statistics for the merged dataset on presidential elections. In this panel, every observation is a county within a four-year presidential election period. Panel B presents descriptive statistics for the merged dataset on legislative elections. In this panel, every observation is a county within a two-year congressional election period. The foreclosure rate is calculated as the mean of annual total foreclosures divided by county population in thousands of people, where the mean is computed over the relevant time period (either the four-year presidential term or the two-year congressional term.) Vote percentages are the Democratic percentage of the two-party vote in a given county. County population is linearly interpolated for non-Census years. The second column provides means that are weighted by county population.

we weight by county population, this average increases to over 50%. The same phenomenon is seen for the U.S. House in Panel B, though it is more muted for U.S. Senate elections. House and Senate elections, as Panel B shows, include cases where the Democrats received 0 or 100% of the vote—these are uncontested races.

For the individual analysis on voter turnout, we focus on the state of Ohio. Ohio is a good test case because it is a closely contested state, electorally, which also happens to offer administrative data on individual voter turnout at no cost. We obtained the Ohio voter file from the Ohio Secretary of State website, and we joined it to CoreLogic’s individual-level foreclosure data using the full name and county of each individual. The voter file provides four important pieces of information. First, anyone who appears in the voter file is, by definition, registered to vote. Second, the voter file shows who, among registered voters, actually turns out to vote in each election. Third, the voter file records the party with which each voter has chosen to register, if the voter has elected to register with a party. And fourth, the voter file shows the date on which each individual registered to vote.

In studying the voter file, we face a difficult post-treatment problem. Individuals only appear in the voter file when and if they register to vote, but the choice to register could itself be influenced by home foreclosure. If we study the effect of foreclosures only for the set of people who choose to

Table 2 – Descriptive Statistics, Individual Level, Ohio, 2004–2014.

	Mean (1)	Standard Deviation (2)	Observations (3)
A. Full Voter File			
Prob of Foreclosure	0.004	0.066	56,067,424
Turnout Rate	0.44	0.50	56,067,424
Dem Registration %	0.17	0.37	55,016,038
Rep Registration %	0.26	0.44	55,016,038
Ind Registration %	0.57	0.49	55,016,038
B. Registered Before 2004			
Prob of Foreclosure	0.002	0.039	20,719,377
Turnout Rate	0.70	0.46	20,719,377
Dem Registration %	0.23	0.42	20,719,377
Rep Registration %	0.41	0.49	20,719,377
Ind Registration %	0.36	0.48	20,719,377

Note: Panel A presents summary statistics for all individuals who appear either in Ohio’s voter file or in the Ohio foreclosure records. Panel B presents these same summary statistics only for individuals who registered to vote prior to the timeframe of our study. The unit of analysis for all rows is an individual-year pair. An individual is marked as foreclosed on if she experienced at least one foreclosure at any time in the year up to the election date in early November. Turnout rate is measured as the percentage of all individuals in the dataset who are marked as turning out in the general election. Dem, Rep, and Ind registration percentages are calculated using party registration as provided in the Ohio voter file, and reflect the proportion of registered voters who register with each party. A negligible number of Ohio voters register for third parties. The number of observations in Panel A is slightly larger for foreclosures and turnout rate than for the party registration variables because we include foreclosed individuals who do not register to vote. These individuals are not included in Panel B, so all sample sizes in Panel B are equal.

register, we risk missing many people who either (a) experienced foreclosure, but did not register to vote, and so did not turnout or (b) did not experience foreclosure and did not register to vote, and so did not turnout. We address this issue in two ways. First, we include anyone from the CoreLogic data who experiences foreclosure, whether or not they appear in the voter file; anyone who does not appear in the voter file is recorded as not turning out to vote, as are any individuals who register to vote but do not turn out. This ensures that we do not eliminate people who were foreclosed on but did not vote. Second, we also perform analyses where we include only people who were registered to vote before the period of our study—thus avoiding post-treatment issues. In particular, for these analyses, we study only people who registered to vote in 2003 or earlier.

Table 2 presents summary statistics for this individual-level analysis. Panel A provides these statistics for the full voter file + CoreLogic dataset, while Panel B provides them for the set of people registered to vote before 2004 (we do not provide minimums or maximums in this table because all variables run from 0 to 1). In both panels, the unit of analysis is the individual-year. The probability that any individual is foreclosed on in any given year is 0.004, in the full dataset, and 0.002 among people registered to vote before 2004—this makes sense since we have reasons to suspect that long-time registered voters are relatively more affluent than non-voters or non-habitual voters (Verba, Schlozman, and Brady 1995).

The second row of each panel shows the turnout rate. Among the full dataset, 44% of person-year observations show that the individual turns out to vote (Panel A). This rate is substantially higher among long-time registered voters (Panel B), as we might expect. The final three rows of each panel show the breakdown of party registration, among registered voters. The majority of registered voters do not affiliate with a party (Panel A), but the majority of long-time registered voters do (Panel B). Among both sets of people, Republican registration is significantly more common than Democratic registration. In the top panel, we drop roughly 1 million person-year observations when we calculate the party registration percentages—these are individuals in the CoreLogic data who do not register to vote, and thus do not have party registrations. This drop does not occur in Panel B because here we have subset to only people appearing in the voter file.

3 Electoral Non-Effects of Home Foreclosures

We first estimate the effects of home foreclosures on incumbent vote share at the county level. Specifically, we estimate equations of the form

$$Dem\ Vote\ Pct_{it} = \alpha Foreclosures_{it} + \beta Foreclosures_{it} \cdot Dem\ Inc_{it} + \gamma_i + \delta_t + \epsilon_{it}, \quad (1)$$

where $Dem\ Vote\ Pct_{it}$ measures the Democratic candidate’s percentage of the two-party vote in county i at time t , running from 0 to 100. The variable $Foreclosures_{it}$ measures total foreclosures per capita in county i during the period from the last election to the election at time t . This variable is also interacted with $Dem\ Inc_{it}$, which takes the value 1 when the Democratic party

held the office in the previous cycle and -1 when the Republican party held it. Finally, γ_i and δ_t stand in for county and year fixed effects, respectively. In many specifications, δ_t is made more flexible, either using state-by-year fixed effects or population decile-by-year fixed effects. For many cases, including when studying the presidency, the main effect on $Dem\ Inc_{it}$ is omitted because it is absorbed by the time fixed effects.

As the above specification makes clear, the analysis is a difference-in-differences design in which we compare within-county changes in incumbent and non-incumbent party vote shares over time across within-county changes in home foreclosures. For the resulting estimate to be causal, the trends in vote share in counties with smaller increases (or decreases) in foreclosures must provide valid counterfactuals for the trends we would have observed in counties with larger increases (or decreases) in foreclosures, had these counties instead had smaller increases (or decreases). The main advantage of the difference-in-differences design is that this assumption can be tested. One way to test it, which we do throughout the results below, is to relax the assumption of parallel trends in a variety of ways and see if the results change. Because the results do not change meaningfully, we conclude that the parallel trends assumption appears to be valid.

Given that we want to study how foreclosures affect incumbent performance, it might seem more logical to use incumbent party vote share, rather than Democratic party vote share, as our dependent variable. This would allow us to forego the interaction term between foreclosures and Democratic incumbency. However, it seems unlikely that counties trend in terms of their general support for incumbents, and far more likely that they might trend in terms of their partisanship. As such, it makes more sense to use the interactive specification with Democratic vote share as the dependent variable, so that we can account for these trends directly. The main quantity of interest is therefore the interaction between foreclosures and the incumbent party variable.

Table 3 presents the estimates for four possible specifications, with and without county population weights. The first column is the most simple, in which we use county fixed effects and state-by-year fixed effects—implicitly doing a separate difference-in-differences for each state and averaging the estimates together. Here we find a statistically significant but substantively small *benefit* to the incumbent party in presidential elections when a county suffers more foreclosures. Specifically, an increase of one foreclosure for every 1,000 people in the county is estimated to increase the incumbent party’s presidential vote share by roughly a quarter of a percentage-point.

Table 3 – Effects of Housing Foreclosures on Presidential Elections, County Level, 2004–2016.

	Dem Presidential Vote Pct (0-100)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Foreclosures Per 1,000 People	-0.38 (0.11)	-0.05 (0.11)	-0.01 (0.11)	-0.28 (0.15)	-0.17 (0.21)	0.12 (0.16)	0.37 (0.16)	0.20 (0.19)
Foreclosures × Inc Party	0.26 (0.08)	-0.06 (0.08)	-0.07 (0.07)	0.10 (0.11)	0.10 (0.14)	-0.26 (0.12)	-0.15 (0.10)	-0.06 (0.12)
N	4411	4411	4411	4411	4411	4411	4411	4411
# Counties	1328	1328	1328	1328	1328	1328	1328	1328
County Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes	No	No	Yes	Yes
Population Weights	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors clustered by county in parentheses. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

Put another way, it is estimated to take an increase of 4 foreclosures per 1,000 people to move the incumbent party’s vote share by a full percentage-point—roughly 3 standard deviations in the foreclosure rate variable (see Table 1). Moreover, this substantively small effect is in the opposite direction predicted by the economic voting literature.

We have reasons to be skeptical of this specification. The state-year fixed effects mean that we compare the changes over time in counties with more foreclosures to the changes over time in counties with fewer foreclosures in the same state. These counties may not give us the best counterfactual trend. Counties that experienced big spikes in foreclosures tend to be urban or suburban counties with high populations—the same types of counties that have become increasingly Democratic over this same time period. As such, it might be better to use similarly populous counties as counterfactuals, even if they are in different states. We do this in column 2, where we create fixed effects for every population decile-year. That is, we chunk cities into deciles based on their population as of 2003 (before the treatment window), and we create year fixed effects within each population decile, so that our difference-in-differences counterfactual trends come from counties in the same population decile who vary in their foreclosure rates over time. When we do this, the estimates shrink even more. As column 2 shows, we now estimate that an increase of 1 foreclosure per 1,000 people in a county decreases incumbent vote share by a tiny 0.06 percentage

Table 4 – Effects of Housing Foreclosures on Legislative Elections, County Level, 2002–2014.

Dem Senate Vote Percentage (0-100)								
Foreclosures Per 1,000 People	0.11 (0.10)	-0.16 (0.33)	0.12 (0.13)	-0.56 (0.43)	0.15 (0.17)	-0.79 (0.45)	0.18 (0.22)	-0.76 (0.50)
Foreclosures × Inc Party	0.10 (0.10)	0.25 (0.29)	0.04 (0.15)	0.33 (0.45)	-0.42 (0.30)	1.18 (0.71)	-0.66 (0.41)	1.37 (1.01)
N	7344	7344	7344	7344	7344	7344	7344	7344
# Counties	1338	1338	1338	1338	1338	1338	1338	1338
Dem House Vote Percentage (0-100)								
Foreclosures Per 1,000 People	-0.19 (0.23)	-0.20 (0.21)	-0.10 (0.28)	-0.33 (0.25)	-0.03 (0.20)	0.11 (0.22)	0.14 (0.21)	0.12 (0.23)
Foreclosures × Inc Party	0.21 (0.19)	0.04 (0.18)	0.19 (0.23)	0.26 (0.20)	0.17 (0.17)	0.05 (0.14)	0.07 (0.18)	0.16 (0.14)
N	7344	7344	7344	7344	7344	7344	7344	7344
# Counties	1338	1338	1338	1338	1338	1338	1338	1338
County Fixed Effects	Yes							
State-Year Fixed Effects	Yes	No	Yes	No	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes	No	No	Yes	Yes
Population Weights	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors clustered by county in parentheses. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

points, or 6 basis points. In this specification, it would take an increase of almost 17 foreclosures per 1,000 people to swing incumbent vote share by a full percentage point—an increase that spans almost the entire range of the data, where foreclosures per 1,000 people range from 0 to roughly 19.

Besides being substantively tiny, this null result is relatively precisely estimated. The 95% confidence interval for this estimate ranges from roughly -0.22 to +0.10, which mean we can reject any substantively meaningful relationship between changes in foreclosure rates and changes in incumbent electoral support.

The remaining specifications explore the robustness of this null result. We include county-specific linear time trends, and we also explore differences in the results when we weight by county population. In all cases, we continue to find precisely estimated, substantively small effects.

Next, we turn to possible effects on legislative elections. Table 4 presents the results for House and Senate elections, respectively, using the same specifications from Table 3. We continue to define party incumbency based on the presidency (in the Appendix, we present results using party incumbency of senators, also.) In the interest of brevity, we do not walk through each individual estimate; but they are largely null. The only cases in which the estimates are arguably large are for effects on Senate races in some columns. However, these estimates are primarily positive, not negative. The only estimate over 1 is positive, indicating that incumbent party Senatorial candidates may benefit, slightly, from an increase in foreclosures. However, these estimates are imprecise, and the bulk are substantively small.

In summary, looking across presidential, U.S. Senate, and U.S. House races, we find no link between home foreclosures and the punishment of incumbents. Counties that have suffered disproportionate home foreclosures do not appear to have voted against incumbents at a different rate than counties experiencing fewer foreclosures. These null results are not the result of noise; in many cases we have very precise estimates which can rule out any substantively meaningful effects. Despite the salience of home foreclosures, they do not seem to alter vote choices across localities. To understand this overall result, we now turn to an analysis of individual turnout behavior. After that, we also analyze several possible mechanisms for these null results.

4 Home Foreclosures Reduce Individual Turnout

We now estimate effects of foreclosure on individual voter turnout. Specifically, we estimate equations of the form

$$\textit{Turns Out}_{it} = \beta \textit{Foreclosed}_{it} + \gamma_i + \delta_t + \epsilon_{it}, \quad (2)$$

where $\textit{Turns Out}_{it}$ is a binary variable that takes the value 1 if individual i turned out to vote in the election at time t , and 0 otherwise. The variable $\textit{Foreclosed}_{it}$ is an indicator for whether individual i experienced a foreclosure during election period t ; γ_i stands in for individual fixed effects, while δ_t stands in for county-by-year fixed effects. In some specifications, we also add interactions of the $\textit{Foreclosed}$ variable with indicators for how close the foreclosure was, in time, to election day.

Table 5 presents the results. The first two columns use the full sample, while the second two use only the individuals registered to vote before 2004, as discussed along with Table 2. As columns 1

Table 5 – Difference-in-Differences Effects of Housing Foreclosures on Individual Turnout in Ohio, 2004 to 2016.

	Turn Out in General Election (Yes/No)			
	Full Voter File (1)	(2)	Registered Before 2004 (3)	(4)
Home Foreclosed	-0.049 (0.001)	-0.046 (0.001)	-0.010 (0.002)	-0.009 (0.002)
Foreclosed w/in 90 Days of Elec		-0.022 (0.004)		-0.010 (0.009)
Foreclosed w/in 30 Days of Elec		-0.012 (0.006)		-0.021 (0.013)
Foreclosed w/in 7 Days of Elec		-0.002 (0.009)		-0.010 (0.020)
N	56,067,655	56,067,655	20,719,377	20,719,377
Individual Fixed Effects	Yes	Yes	Yes	Yes
County-Year Fixed Effects	Yes	Yes	Yes	Yes

Robust standard errors clustered by county in parentheses.

and 3 show, we find modest but highly precise negative estimates. Individuals who suffer foreclosure are less likely to turn out to vote—as much as, perhaps, 5 percentage-point (column 1). Effects are more modest in column 3, likely because the set of people who are in the voter file and registered to vote before 2004 are habitual voters who are correspondingly less sensitive to any stimuli that might reduce their propensity to vote.

Columns 2 and 4 show how these effects vary as the foreclosure occurs closer to election day. We interact the foreclosures variable with indicators for whether the foreclosure occurred within 90 days of the election, within 30 days, and within 7. Because these are nested indicators, effects are cumulative—e.g., to get the overall estimated effect of a foreclosure within 7 days of the election, we add the main effect and all three of the interaction coefficients. Estimated effects grow—become more negative—when foreclosures occur closer to the election, as we might expect. Consider the results in column 2. A foreclosure that happens any time outside of 90 days of an election reduces turnout by 4.6 percentage-points; this effect grows to 6.8 percentage-points if the foreclosure occurs within 90 days, to 8.0 percentage points if the foreclosure occurs within 30 days, and to 8.2 percentage points if the foreclosure occurs within 7 days of the election (this final increase is insubstantial compared to the 30-day increase, and not significantly different).

Interestingly, these estimates are quite similar to those in a previous study of residential mobility. Gay (2012) studies individuals who randomly received the opportunity to move out of public housing and into private apartments during the 1994 Moving to Opportunity for Fair Housing Demonstration Program. Experimental estimates of the effect of this opportunity on turnout suggest that it decreased turnout by roughly 2–4 percentage points. The similarity in our estimates further supports the idea that home foreclosures have not galvanized political activity; instead, foreclosures seem to affect individuals in much the same way other forms of residential mobility—even those that are the result of an intentional opt-in program and that are not associated with a political crisis—do.

In sum, experiencing a home foreclosure makes an individual in Ohio less likely to turn out and vote, on average. This effect is not massive but is very precisely estimated, which means we can rule out positive effects. It does not appear to be true that home foreclosures galvanize individuals to participate in politics, in Ohio at least.

5 Why Don't Foreclosures Affect Voter Behavior, Generally?

Why don't foreclosures galvanize people, the way some popular accounts have suggested they might? At the individual level, there are probably a number of reasons, but the results across temporal proximity to elections suggest one in particular. Foreclosures no doubt upset people, and this may encourage them to become politically active, but they also cause stress and many new obligations. Foreclosed individuals have complicated financial situations to deal with, not to mention that they must find a new residence. More prosaically, simply being forced to change addresses may reduce their probability of registering to vote (e.g., Squire, Wolfinger, and Glass 1987). Strains like these raise the opportunity cost of participating in politics, and likely reduce participation as a result. We should probably be skeptical of posited links between other forms of economic distress and individual political behavior for the same reasons. In general, the people who experience economic distress may have the greatest incentives to pay new attention to the political process, but they are also the people experiencing life events that reduce their ability or desire to spend time involved in politics.

Table 6 – Effects of Housing Foreclosures on Presidential Elections, County Level, 2004–2016: Testing for Trump-Clinton Effects.

	Dem Presidential Vote Pct (0-100)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Foreclosures Per 1,000 People	-0.15 (0.07)	-0.09 (0.07)	-0.08 (0.09)	-0.22 (0.11)	-0.08 (0.11)	-0.04 (0.10)	0.12 (0.14)	0.08 (0.13)
Foreclosures × 2016	0.04 (0.21)	-0.15 (0.19)	-0.15 (0.25)	-0.18 (0.27)	-1.41 (0.39)	-1.83 (0.31)	-1.16 (0.42)	-0.83 (0.42)
N	4411	4411	4411	4411	4411	4411	4411	4411
# Counties	1328	1328	1328	1328	1328	1328	1328	1328
County Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes	No	No	Yes	Yes
Population Weights	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors clustered by county in parentheses. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

The turnout results do not speak to aggregated vote choices, however. Even if they do not galvanize the individuals who experience foreclosure, why don't foreclosures alter localities' aggregated vote choices? One possibility is that voters don't react to foreclosures in their area by altering their support for one party or the other because the parties' positions on policies related to the housing market and to financial regulation are relatively similar, most of the time. One way to investigate this possibility, though it is crude, is to see whether effects seem to be different in the 2016 election. Donald Trump deviated sharply from policy positions traditionally held by both parties, particularly related to free trade, infrastructure spending, student loans, and other economic issues that voters might care about. Whether any of these issues linked to home foreclosures is unclear, but as we have mentioned earlier, upticks in home foreclosures in an area are likely signs of other forms of economic distress. Trump's overall separation on economic policy might potentially make the 2016 election different from previous ones.

In Table 6, we reestimate the effects adding an interaction of the foreclosures variable with a dummy variable for the 2016 election cycle. Results are mixed but suggest some evidence that Trump benefited from foreclosures, electorally. The results in the first four columns show no such effects; however, the final four columns, when we weight by county population, show noticeable effects. Across specifications, an increase of 1 foreclosure per 1,000 people is estimated to decrease

Table 7 – Effects of Housing Foreclosures on Presidential Elections, County Level, 2004–2016: Including Only Competitive Counties.

	Dem Presidential Vote Pct (0-100)							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Foreclosures Per 1,000 People	-0.63 (0.16)	-0.04 (0.14)	-0.10 (0.16)	-0.32 (0.20)	-0.10 (0.29)	0.11 (0.22)	0.61 (0.20)	0.29 (0.21)
Foreclosures × Inc Party	0.40 (0.11)	-0.07 (0.10)	-0.12 (0.10)	0.12 (0.14)	0.03 (0.14)	-0.29 (0.15)	-0.24 (0.12)	-0.15 (0.14)
N	2695	2695	2695	2695	2695	2695	2695	2695
# Counties	817	817	817	817	817	817	817	817
County Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes	No	No	Yes	Yes
Population Weights	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors clustered by county in parentheses. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

the Democratic vote share (the Clinton vote share) by 0.83 to 1.83 percentage points more in 2016 than in past years (this the range of the interaction coefficient estimates, not their confidence intervals). These larger effects in the weighted specifications suggest that, if foreclosures led voters to punish Clinton (and the incumbent Democratic party), and to reward Trump, this behavior was concentrated in more populous counties. It is difficult to be confident in the identification of these effects, relying as they do on changes within only one four-year period, but they are at least suggestive. Voters may rarely switch parties or punish incumbents in response to personal economic distress not only because they do not care about politics, but also because most of the time the two parties offer relatively similar economic policies. To the extent we find evidence in favor of this claim in 2016, however, the magnitudes are not large, and most people do not seem to have altered their voting behavior. The majority of voters seem unaffected.

It is also possible that voters might not typically react to foreclosures because of strong ideological views and/or strong partisan loyalties. Even if historically voters did punish incumbents for a host of economic outcomes, polarization may have dulled this response in more recent elections. Although there is no perfect way to test this claim, we try to get at it indirectly by examining if the effects of foreclosures are different in counties that are more politically competitive—counties where voters, as a whole, have not consistently supported one party in the past. This is not a perfect test

since a county might be competitive only because it contains partisan or ideological voters of each party in equal proportion, rather than containing more swing voters, but it is at least a possibility, and one which would be consistent with empirical research that suggests incumbents offer more moderate positions in more competitive places (Ansolabehere, Snyder, and Stewart 2001). Specifically, we compute each county’s “normal vote” as the average Democratic vote share for president in the county in all elections from 1972–2000 (thus avoiding any post-treatment issues). We then re-estimate the regressions from Table 3 focusing only on competitive districts, which we define to be districts where the normal vote is between 40% and 60%. As Table 7 shows, we continue to find null results. If even more politically competitive districts do not appear to punish incumbent party candidates for foreclosures, it seems less likely that the null results are driven by ideological voters or partisan loyalty.

6 Conclusion

The political aftermath of the Great Recession has raised a number of questions about the role of economic conditions in electoral politics. While aggregate correlations between economic conditions and incumbent party success are apparent, the individual-level mechanisms underlying these patterns are not well understood, and little evidence exists for whether any economic links, whether individual or aggregate, actually *cause* changes in political outcomes. In this paper we have investigated the individual link between economic distress and political behavior in the context of home foreclosures, surely one of the most salient individual experiences with economic distress, and an especially relevant one in the shadow of the Great Recession. Our findings raise questions about the supposed causal links between economic distress and individual political behavior; we find no evidence that areas that experienced more foreclosures punish incumbent candidates for the presidency or for legislative offices. At the individual level, our findings may help explain one of the reasons why we observe no such link. Individuals who suffer economic distress in the form of home foreclosures are somewhat less likely to turn out and vote—not more. Rather than being galvanized, the increased personal costs of voting after a home foreclosure deter individuals from participating in politics. These turnout effects alone do not rule out changes in vote choice among those who do vote, but they suggest to us a lack of political activism in the aftermath of these crises,

at least among those directly affected by the crises. In general, we suspect, those with the most reasons to be angry in the aftermath of economic crises are not very likely to become politically engaged.

We have also sought to test some of the possible mechanisms for the lack of a link between economic conditions and incumbent vote share at the county level. Null results in competitive counties suggest to us that the non-effect is not purely the result of partisan or ideological rigidity among voters. On the other hand, the positive link we find between foreclosures and vote share for Trump at the county level in 2016 suggests to us that the similarity of the parties' historic positions on policies related to the economy may mute voter responses. Although these results are tentative, and should be explored further in future work, the 2016 results make us suspect that links between local economic conditions and voter behavior may be stronger for cases where the two parties' candidates have clear policy differences associated with current economic conditions. Hypotheses like this one, which explore the strategic interplay between candidate and voter behavior, are largely absent from the literature on economic voting, which has mainly focused on one-sided responses of voters to economic conditions.

In addition to its substantive findings, our paper also adds to a growing literature concerned that publication bias has systematically skewed empirical estimates of many important phenomena across many disciplines (De Long and Lang 1992; Gelman and Loken 2014; Gerber and Malhotra 2008; Simonsohn, Nelson, and Simmons 2014). Historically, it has been difficult for null results to appear in published academic journals. Most published estimates are non-null as a result, even in cases when underlying relationships actually are null. The large literature on economic and retrospective voting is just as likely to face these issues as any other literature. While each individual published study may be perfectly executed, the selection of positive results for publication has most likely led our aggregate view of economic and retrospective voting to overstate the effect of economic stimuli on incumbent electoral fortunes. Our results are consistent with this possibility, and they add a contrasting finding to the empirical literature on retrospective voting. Losing one's home to foreclosure, or living in an area where foreclosures are prevalent, constitutes a clearly visible, highly salient, and emotional experience, yet we find no major response in voter behavior.

Understanding how voters evaluate incumbents in real elections is at the core of empirically evaluating democratic accountability. We take no position here on how voters *ought* to evaluate

incumbents, or on whether the patterns we have observed are “rational” or not. Either way, a bevy of recent work, which we reviewed in the Introduction, argues that voters respond, and even over-respond, to local and personal economic conditions, arguing that these over-responses erode democratic accountability. Our results suggest that the way voters respond to economic conditions is more complicated. Perhaps surprisingly, local economic conditions related to the housing market do not appear to play an important role in U.S. election results in our sample, and the individuals most directly affected by the housing crisis have, if anything, been demobilized rather than galvanized in its aftermath.

References

- Achen, Christopher H., and Larry M. Bartels. 2016. *Democracy for Realists: Why Elections Do Not Produce Responsive Government*. Princeton University Press.
- Angrist, Josuha D., and Jorn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Ansolabehere, Stephen, James M. Snyder, Jr., and Charles Stewart, III. 2001. "Candidate Positioning in US House Elections." *American Journal of Political Science* 45(1): 136–159.
- Ansolabehere, Stephen, Marc Meredith, and Erik Snowberg. 2014. "Macro-Economic Voting: Local Information and Micro-Perceptions of the Macro-Economy." *Economics & Politics* 26(3): 380–410.
- Ashworth, Scott, and Ethan Bueno De Mesquita. 2014. "Is Voter Competence Good for Voters? Information, Rationality, and Democratic Performance." *American Political Science Review* 108(3): 565–587.
- Ashworth, Scott, Ethan Bueno De Mesquita, and Amanda Friedenberg. 2016. "Learning About Voter Rationality." Working Paper.
- Bagues, Manuel, and Berta Esteve-Volart. 2016. "Politicians' Luck of the Draw: Evidence from the Spanish Christmas Lottery." *Journal of Political Economy* 124(5): 1269–1294.
- Brender, Adi, and Allan Drazen. 2008. "How Do Budget Deficits and Economic Growth Affect Reelection Prospects? Evidence From a Large Panel of Countries." *The American Economic Review* 98(5): 2203–2220.
- De Long, J. Bradford, and Kevin Lang. 1992. "Are All Economic Hypotheses False?" *Journal of Political Economy* 100(6): 1257–1272.
- Duffie, Darrell. N.d. "Financial Regulatory Reform After the Crisis: An Assessment." *Management Science*. Forthcoming.
- Enos, Ryan D., Anthony Fowler, and Lynn Vavreck. 2014. "Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate." *The Journal of Politics* 76(1): 273–288.
- Fair, Ray C. 1978. "The Effect of Economic Events on Votes for President." *The Review of Economics and Statistics* 60(2): 159–173.
- Fowler, Anthony, and Andrew B. Hall. 2016. "Do Shark Attacks Influence Presidential Elections? Reassessing a Prominent Finding on Voter Competence." Working Paper.
- Fowler, Anthony, and B. Pablo Montagnes. 2015. "College Football, Elections, and False-Positive Results in Observational Research." *Proceedings of the National Academy of Sciences* 112(45): 13800–13804.
- Gaspar, John T, and Andrew Reeves. 2011. "Make it Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters." *American Journal of Political Science* 55(2): 340–355.

- Gay, Claudine. 2012. "Moving to Opportunity: The Political Effects of a Housing Mobility Experiment." *Urban Affairs Review* 48(2): 147–179.
- Gelman, Andrew, and Eric Loken. 2014. "The Statistical Crisis in Science Data-dependent Analysis—A "Garden of Forking Paths"—Explains Why Many Statistically Significant Comparisons Don't Hold Up." *Am Sci* 102(6): 460.
- Gerber, Alan, and Neil Malhotra. 2008. "Do Statistical Reporting Standards Affect What Is Published? Publication Bias in Two Leading Political Science Journals." *Quarterly Journal of Political Science* 3(3): 313–326.
- Gosnell, Harold F., and William G. Colman. 1940. "Political Trends in Industrial America: Pennsylvania An Example." *Public Opinion Quarterly* 4(3): 473–486.
- Healy, Andrew, and Gabriel S. Lenz. N.d. "Presidential Voting and the Local Economy: Evidence from Two Population-based Datasets." *Journal of Politics*. Forthcoming.
- Healy, Andrew, and Neil Malhotra. 2009. "Myopic Voters and Natural Disaster Policy." *American Political Science Review* 103(3): 387–406.
- Healy, Andrew J, Neil Malhotra, and Cecilia Hyunjung Mo. 2010. "Irrelevant Events Affect Voters' Evaluations of Government Performance." *Proceedings of the National Academy of Sciences* 107(29): 12804–12809.
- Hill, Seth J., Michael C. Herron, and Jeffrey B. Lewis. 2010. "Economic Crisis, Iraq, and Race: A Study of the 2008 Presidential Election." *Election Law Journal* 9(1): 41–62.
- Huber, Gregory A., Seth J. Hill, and Gabriel S Lenz. 2012. "Sources of Bias in Retrospective Decision Making: Experimental Evidence on Voters? Limitations in Controlling Incumbents." *American Political Science Review* 106(4): 720–741.
- Kiewiet, D Roderick, and Michael S Lewis-Beck. 2011. "No man is an island: self-interest, the public interest, and sociotropic voting." *Critical Review* 23(3): 303–319.
- Kinder, Donald R., and D. Roderick Kiewiet. 1981. "Sociotropic Politics: the American Case." *British Journal of Political Science* 11(2): 129–161.
- Kramer, Gerald H. 1971. "Short-term Fluctuations in US Voting Behavior, 1896-1964." *American Political Science Review* 65(1): 131–143.
- Kramer, Gerald H. 1983. "The Ecological Fallacy Revisited: Aggregate-versus Individual-level Findings on Economics and Elections, and Sociotropic Voting." *American Political Science Review* 77(1): 92–111.
- Lewis-Beck, Michael S. 1986. "Comparative Economic Voting: Britain, France, Germany, Italy." *American Journal of Political Science* 30(2): 315–346.
- Lewis-Beck, Michael S., and Mary Stegmaier. 2007. "Economic Models of Voting." In *The Oxford Handbook of Political Behavior*, ed. Russell Dalton, and Hans-Dieter Klingemann. Oxford University Press pp. 518–537.
- Mian, Atif, Amir Sufi, and Francesco Trebbi. 2015. "Foreclosures, House Prices, and the Real Economy." *The Journal of Finance* 70(6): 2587–2634.

- Nordhaus, William D. 1989. "Alternative Approaches to the Political Business Cycle." *Brookings Papers Econ. Activity* (2): 1–68.
- Pearson, F.A., and W.I. Myers. 1948. "Prices and Presidents." *Farm Economics* 163: 4210–4218.
- Simonsohn, Uri, Leif D. Nelson, and Joseph P. Simmons. 2014. "P-curve: A Key to the File-drawer." *Journal of Experimental Psychology: General* 143(2): 534.
- Squire, Peverill, Raymond E. Wolfinger, and David P. Glass. 1987. "Residential Mobility and Voter Turnout." *American Political Science Review* 81(1): 45–65.
- Stigler, George J. 1973. "General Economic Conditions and National Elections." *The American Economic Review* 63(2): 160–167.
- Tibbitts, Clark. 1931. "Majority Votes and the Business Cycle." *American Journal of Sociology* 36(4): 596–606.
- Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics*. Harvard University Press.
- Wolfers, Justin. 2007. "Are Voters Rational? Evidence from Gubernatorial Elections." Unpublished Manuscript.

Online Appendix

Intended for online publication only.

Contents

A.1	Unifying Our Results with Existing Research	28
A.1.1	Evaluating General Effects of the Local Economy	28
A.1.2	Evaluating Effects of the Housing Crisis in California	30
A.2	Senate Results Using Senate Incumbency	31
A.3	Processing the Deed-Level Data	32
A.4	Assessing the Generalizability of the CoreLogic Data	34

A.1 Unifying Our Results with Existing Research

Because Healy and Lenz (N.d.) studied the effects of local economic conditions with different data, and because the authors of the study have made their replication data public, we have a somewhat unusual opportunity to consider evidence across the two studies and aggregate them together to obtain more information about local economic conditions and incumbent electoral fortunes.

In this section, we compare our results to the two main findings from Healy and Lenz (N.d.), and we use these comparisons to offer an aggregated view of the estimated links, considering both datasets in context. Specifically, the two main findings from Healy and Lenz (N.d.) are: (1) that an index of local wage and employment growth predicts incumbent electoral performance at the county level between 1990 and 2012; and (2) that mortgage delinquencies at the zip-code level in California predicted Obama vote gains in 2008.

We draw two main conclusions from the analyses below. First, using our preferred specification, which is the difference-in-differences design, we estimate precise null results of local economic conditions on incumbent electoral fortunes using both our dataset on home foreclosures and the Healy and Lenz (N.d.) dataset on local economic conditions. Second, using the estimates from the difference-in-differences design and the lagged dependent variable model preferred by Healy and Lenz (N.d.), applying both strategies to both datasets after rescaling the explanatory variables, we bracket the true effect and argue that it is quite modest in either case. In sum, while there may be unusual cases where effects exist—like California in 2008, a case that Healy and Lenz (N.d.) identifies and for which we, too, find an effect (see later discussion below)—in general, there does not appear to be a substantial link between local economic conditions and incumbent electoral performance.

A.1.1 Evaluating General Effects of the Local Economy

To start, we compare the nationwide analysis from Healy and Lenz (N.d.) to our nationwide results. The main explanatory variable that Healy and Lenz (N.d.) uses in the nationwide analysis is an index of local economic growth—the simple average of wage and employment growth over the past 6 months before the election—that has been de-meanned by county, and then rescaled so that it is 0 for the .001 percentile and 1 for the .999 percentile. Healy and Lenz (N.d.) uses this variable in a lagged dependent-variable framework to estimate its effect on incumbent vote share.

A key difference between the main results presented in this paper and those from Healy and Lenz (N.d.) is that they come from different empirical designs.¹⁰ While Healy and Lenz (N.d.) uses the lagged dependent variable model, the main results in our paper come from the difference-in-differences framework. We prefer the difference-in-differences framework to the lagged dependent variable setup because its key identifying assumption, parallel trends, is clear and testable. In the body of our paper, we test the validity of the difference-in-differences design by showing that our conclusions do not change when we add county linear time trends, or when we add several different types of time fixed effects. Both of these strategies are ways to relax the parallel trends assumption of the difference-in-differences design, and the fact that the estimates are similar across them suggests that the assumption is sound. These kinds of robustness checks are not possible in the lagged dependent variable framework, which is one reason we prefer the difference-in-differences.

¹⁰The reader might also note a second difference, which is that the Healy and Lenz (N.d.) regressions are weighted by total votes while our results are weighted by county population. We have examined these two approaches and found that it makes only negligible differences in the results. That said, the Healy and Lenz (N.d.) results do attenuate considerably when the counties are not weighted, as do our 2016/Trump estimates in Table 6. We do not focus too much on this issue, because it seems reasonable to down-weight low population counties.

Moreover, we know that the lagged dependent variable approach is biased when parallel trends, the difference-in-differences assumption, is met (Angrist and Pischke 2009)—and, as we have just discussed, we have reasons to believe the parallel trends assumption is met in this case.

All that said, both strategies require assumptions and neither is a silver bullet. In cases where the parallel trends assumption is violated, like in cases where there is substantial year-to-year variability in treated and control potential outcomes, the lagged dependent-variable model can perform better. Accordingly, it is instructive to look at estimates for both designs on both datasets. To try to make estimates comparable, we rescale the main explanatory variables, foreclosures per 1,000 people and local wage and employment growth, so that each takes the value 0 at the median and 1 at the .99 percentile.¹¹ Finally, to make the estimates go in the same direction, we take the negative of the point estimates from our foreclosure data (since an increase in foreclosures is the opposite of an increase in economic growth). Figure A.1 plots the results along with 95% confidence intervals from robust standard errors clustered by county.

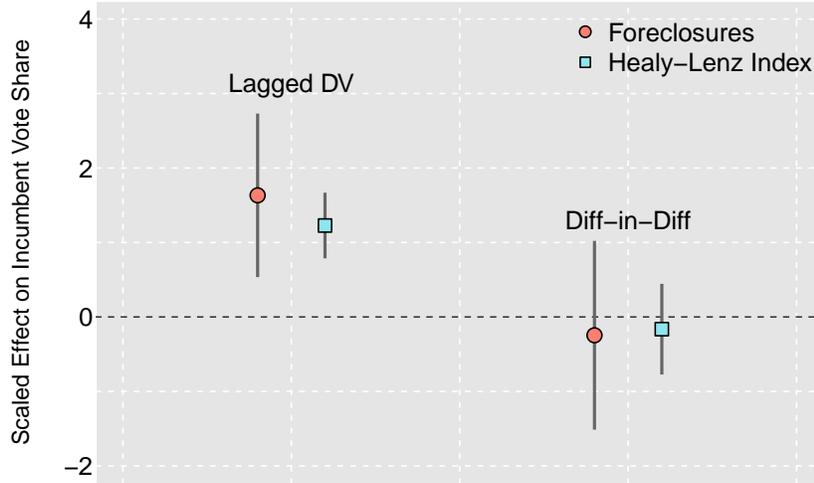
As the figure shows, the estimates from the two datasets are surprisingly similar, once variables are rescaled to be comparable and the same empirical design is used. This is despite differences in the time period and the economic measures used, which makes the similarity all the more remarkable. How do we interpret these similar effects? As the plot shows, when we use the lagged dependent variable setup, we find positive effects, but these effects are quite modest in size. A move from the median to the 99th percentile is an extreme shift, and even this shift is estimated to increase the incumbent’s vote share by fewer than 2 percentage points in both datasets.

These effects become very close to zero (and in fact slightly negative), when we use the difference-in-differences design on either dataset. These null results are precisely estimated; as the figure shows, the 95% confidence intervals rule out effects even as large as +1 percentage point, for the foreclosures data, and 0.42 percentage points (42 basis points) for the local economic growth data.

Determining whether we think there are positive effects thus requires choosing between the two empirical strategies. As Angrist and Pischke (2009) explains, we can use the two estimators to bracket the true effect, if we assume that one of the two identifying assumptions is correct. The two foreclosures estimates give us the most extreme estimates in either direction, positive and negative; as such, using the bracketing principle, we can offer an estimated effect range from -0.25 to +1.63, where again these are effects on incumbent vote share from increasing the foreclosure rate from the median rate to the .99 percentile rate. If we use the bracketing principle with the 95% confidence intervals, we bound the effect of good local economic changes, defined using the shift from the median to the .99 percentile, between -1.50 and +2.71 percentage points. Again, these are estimated effects scaled to what is among the most extreme plausible swings in economic conditions that we could see in the data. That such large swings are only estimated to move incumbent vote by at most a few percentage points—and that what we think are the most plausible estimates even

¹¹This is not the one-unit shift that Healy and Lenz (N.d.) uses; instead, the paper reports the effect for a shift from the .001 percentile to the .999 percentile. This is a very large shift, because both the .001 and the .999 percentile are extreme outliers in the distribution of the economic growth index variable. And because the index variable is already a growth variable, i.e., a measure of changes in the data, this resulting scaled variable reflects the difference between what is one of the most negative changes observed in the data (the .001 percentile of growth) and one of the most positive observed changes in the data (the .999 percentile of growth). As a result, the coefficient on this scaled variable corresponds to the change in incumbent vote share for a very extreme and unusual change in local economic conditions. Specifically, the distance between the .001 percentile and the .999 percentile of the growth variable, after de-meaning by county but not rescaling, is roughly 0.4. The .01 and .99 percentiles of the growth variable are -0.08 and +0.08, respectively. This means that changes of 0.08, like that from the median to the .99 percentile, which are more extreme than 98% of all the data, are smaller than the hypothetical shift from the .001 percentile to the .999 percentile by a factor of 5.

Figure A.1 – Comparing Estimates to Healy and Lenz. We compare estimates from our foreclosures dataset to estimates from the Healy and Lenz data by estimating effects in each dataset using the same two designs after rescaling the explanatory variables so that the regression coefficient of interest represents a shift from the median to the 0.99 percentile of each variable. Foreclosures estimates are multiplied by -1 so that they are comparable to the Healy and Lenz results (since a decrease in foreclosures is positively correlated with economic growth).



go in the opposite direction—suggests to us that the link between local economic conditions and incumbent electoral fortunes is in general very modest, if it exists.

A.1.2 Evaluating Effects of the Housing Crisis in California

Having considered overall effects, we now zoom in on estimating the effect of foreclosures on support for Obama, studying California in the period 2004-2008, mirroring the Healy and Lenz (N.d.) analysis on mortgage delinquencies in California. Given the overall results, the question now is whether we have evidence for a salient case producing a larger link between local economic conditions, specifically measured via the housing collapse, and incumbent electoral fortunes.

We start by replicating the Healy and Lenz (N.d.) finding using our foreclosure data. We use our preferred difference-in-differences design, like in the body of this paper.¹² Following Healy and Lenz (N.d.), we relate changes from 2006 to 2008 in foreclosures per 1,000 people at the county level to changes in Democratic vote from 2004 to 2008. The results are presented in the first two columns of Table A.1; in the first column, we do not weight by population county, in the second, we do. Consistent with Healy and Lenz (N.d.), we find a positive and statistically significant link between foreclosures and support for Obama in California in 2008. To make the estimates comparable to those from the previous section, we again rescale the foreclosures per 1,000 people variable so that it is 0 at the median foreclosure rate and 1 for the 0.99 percentile. The increase from the median to the 0.99 percentile is estimated to increase Obama vote in 2008 by 4.42 or 3.57 percentage points, depending on the specification. These effects are much larger than those

¹²Because we only have two time periods for this replication, we difference the y and x variables rather than using county and year fixed effects.

Table A.1 – Effects of Housing Foreclosures on Presidential Elections, 2004–2008

	Δ Dem Pres Vote (0-100)			
	California	Top 5		
	(1)	(2)	(3)	(4)
Δ Foreclosures Per 1,000 People	4.42 (0.58)	3.57 (0.67)	0.86 (1.29)	0.18 (1.03)
N	35	35	43	43
State Fixed Effects	–	–	Yes	Yes
Population Weights	No	Yes	No	Yes

Robust standard errors clustered by county in parentheses.

estimated overall—indeed, the difference-in-differences estimates on the full sample sometimes even have the opposite sign—and, while they are not huge, they are clearly discernible.¹³

Our main results (Table 3) suggest that this relationship does not generalize to other times and places. But might we find it in other highly salient contexts, like we seem to in California during the heyday of the subprime crisis? To investigate this, in the last two columns of Table A.1, we apply the design to the 5 other states who experienced the largest changes in foreclosures per 1,000 people between 2006 and 2008 in the CoreLogic data. Specifically, these states are: Arizona, Florida, Michigan, Nevada, and Rhode Island. As the final two columns show, we find null results for these cases. Although the estimates here are not precise, the magnitude is much smaller than the California case, and we cannot reject the null of no effect. Combined with the results in the body of the paper, it appears that home foreclosures do not affect incumbent electoral fortunes in general. We cannot rule out that there are special cases where they do—perhaps including California in 2008 and certain areas of the U.S. in 2016 (see Table 6 in the body of our paper)—but there does not seem to be a general phenomenon underlying these effects.¹⁴

A.2 Senate Results Using Senate Incumbency

Here, we re-estimate the results on legislative elections, defining the incumbent party for the election based on the incumbent senator rather than president. Because the data is at the county level, we do not repeat this exercise for the U.S. House, where there is not a one-to-one correspondence of incumbent members of congress to counties. As the table shows, we continue to find null results with this alternate specification.

¹³We mean opposite sign in the sense that here we find incumbents doing worse when the economy does worse. Mechanically the sign is positive in both cases, but that’s because here we are just using Democratic vote share without an indicator for incumbency, because the Democrats are the non-incumbent party in 2008.

¹⁴In considering these more specific effects, we should keep in mind that zooming in on cases with only two time periods does raise the probability of finding false positives. With only one difference to compute over time, resulting standard errors are highly suspect, so many patterns that result from noise may look statistically significant. This could be an issue for the California analysis as well as for the 2016/Trump analysis in Table 6.

Table A.2 – Effects of Housing Foreclosures on Legislative Elections, County Level, 2002–2014.

	Dem Senate Vote Percentage (0-100)							
Foreclosures Per 1,000 People	0.19 (0.10)	0.05 (0.25)	0.16 (0.16)	-0.16 (0.43)	-0.26 (0.22)	0.09 (0.48)	-0.26 (0.32)	0.14 (0.64)
Foreclosures \times Inc Party	-0.04 (0.10)	-0.05 (0.22)	-0.06 (0.16)	-0.76 (0.43)	0.26 (0.21)	-0.14 (0.42)	0.14 (0.30)	-0.40 (0.60)
N	7344	7344	7344	7344	7344	7344	7344	7344
# Counties	1338	1338	1338	1338	1338	1338	1338	1338
County Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes	No	No	Yes	Yes
Population Weights	No	No	No	No	Yes	Yes	Yes	Yes

Robust standard errors clustered by county in parentheses. Inc Party is 1 for two Dem Senators, -1 for two Rep Senators, and 0 for split party. Main effect for Inc Party is absorbed by fixed effects.

A.3 Processing the Deed-Level Data

Each row in the CoreLogic Deed data represents a deed-level event for a property. Deed-level events include typical transactions like sales and foreclosures, as well as lesser-known transactions like the submission of a Notice of Default. Our research concerns itself with foreclosure events (e.g., sale of property by a real estate organization) and pre-foreclosure events (e.g., Notice of Default, lis pendens); thus, we seek to narrow the dataset of all transactions into a dataset of only foreclosure records.

We begin by sorting every row in the Deed data into one of three categories: foreclosure event, pre-foreclosure event, or no indication of this being a foreclosure-related event. Four fields in each row provide us the necessary information to divide the rows into categories.

- First, we inspect the Secondary Deed Cat Codes in each row. The contents of these character positions are “detailed category codes providing additional deed information” per the CoreLogic documentation. If the field contains the value “O” (meaning Real-Estate Owned transfer) or “P” (meaning Real-Estate Owned sale), then we classify the row as a foreclosure event, since the property is under real estate control.
- If the Secondary Deed Cat Codes are inconclusive, we continue to the character positions in the row corresponding to the Mortgage Document Type, which is the “Type of Deed Used for Recording” per CoreLogic documentation. Various states signify foreclosures in different ways, so, if any of the following flags appears, we classify the transaction as a foreclosure event: “CO” (Commissioner’s Deed (foreclosure)), “FD” (Foreclosure Deed), “MF” (Mortgage Foreclosure Deed), “NT” (Notice of Trustee’s Sale), “SC” (Sheriff’s Certificate of Foreclosure), “TE” (Trustee’s Deed Upon Sale (Foreclosure)), “U” (Foreclosure Deed).
- If the Mortgage Document Type is inconclusive, we inspect the character positions in the row corresponding to Document Type, which is the “Type of Transfer Document Recorded” per CoreLogic documentation. If this field takes the value “U” (a Foreclosure event), then

we classify the row as a foreclosure event. Else, if this field takes the value “N” (Notice of Default (NOD)) or “L” (Lis Pendens), we classify the row as a pre-foreclosure event.

- If the Document Type is inconclusive, we examine the character positions in the row corresponding to Foreclosure Field, an “Indicator Showing the Transaction is a Foreclosure” per CoreLogic documentation. If the value of this indicator is “O” or “Y” (REO-Nominal transaction or a Transfer Between Bank and FNMA, FHA, etc.) or “P” (REO Sale - Sale from Government to Third Party), then we classify the event as a foreclosure event.
- If none of the above described criteria are met, we classify the row as having no indication of being a foreclosure event and omit it from the dataset.

Following this narrowing to just the foreclosure and pre-foreclosure transactions, we employ a three-step filter to further narrow these rows into the most relevant data points.

First, we filter out transactions that correspond to businesses or trusts and not to individual homeowners. Because we seek to observe relationships between individuals’ voting behavior and their home foreclosures, information about the foreclosures of enterprise properties does not provide us useful information. One character in the Corelogic Deed rows corresponds to an indicator of whether the property in question is owned by a corporation or trust. If the character assumes the value of “C” (for “Corporation”), “T” (for “Trust”), or “Y” (for “Yes-Corporation”), then we omit the transaction from further consideration. If the character position is vacant (i.e. occupied by whitespace), we assume the property does not belong to a corporation or trust and retain the record.

Next, we filter out transactions for which no home address numbers are available. We do so because, in order to link foreclosure information to voter files, we require a name and address of residence. Simply knowing that an anonymous foreclosure occurred in a county does not provide us with information we can use to model the relationship between foreclosures and voting behavior. To filter these anonymous transactions, we check the character positions in the Corelogic Deed data row corresponding to the house number; if they are occupied by whitespace, we omit the record.

Finally, we filter out transactions for which no date information was available. Our goal is to tally the foreclosures that occurred between elections, so as to analyze their possible impact on vote behavior in the election cycle containing the foreclosures. For example, if a foreclosure occurred on January 3rd, 2001, we would tally it as a foreclosure relevant to the 2002 election. If no date information is present, we cannot arbitrarily assign the foreclosure record to a time period out of concern for biasing the results of our inquiry. To prevent mis-allocation of undated foreclosure records, we omit the Deed events that contain no date information. If the character positions in the Deed data row corresponding to the date are vacant, we disregard the record.

We now possess a list of transactions classified as either a foreclosure event or as a pre-foreclosure event. However, we still need to transform these to voter-level records. One voter owning multiple properties could appear multiple times within our current document. To overcome this, we build a dictionary; we group all transactions that share a unique owner; a unique owner is one with a distinct four-tuple comprising FIPS (county) code, city, owner’s last name, and owner’s first name. Using the transactions within each group, we tally the number of foreclosure-related transactions that occurred prior to the elections (November 7th) in 2004, 2006, 2008, 2010, 2012, 2014, and 2016 for each unique owner. For each of these election years, we record the number of foreclosure-related events for that unique owner since the last election year, the date of the earliest foreclosure-related event that occurred in that time interval, and the severity of the most severe foreclosure-related event that took place. Severity can assume three values: either an “I” (incomplete) foreclosure if only pre-foreclosure events (Notice of Default or lis pendens) occurred in that election cycle, or “C”

(complete) if even one foreclosure event occurred. If neither is true, it assumes the value of “N” (no foreclosure).

Following this consolidation of transaction-level data into unique-owner-level data, each row in our data contains a FIPS, a city, the owner’s last name, the owner’s first name, and the three aforementioned items per election year. Equipped with this information, we seek to now associate as many homeowners in this file as possible to a voter file containing vote behavior for each election of interest. Each row in our voter file corresponds to an individual voter; for each entry, it contains a FIPS, a city, a first name, and a last name. We use these fields to perform a left-join between our foreclosure data and the voter file, retaining all rows in our foreclosure data with appended voter information where available. Where unavailable, the voter information is merely populated as “NA”. We then proceed with further analysis.

A.4 Assessing the Generalizability of the CoreLogic Data

The CoreLogic data only covers roughly half the counties in the United States. Although this does not induce any obvious bias into our statistical analyses, we would like to know how specific the resulting estimates are to the CoreLogic context. Accordingly, Table A.3 compares the CoreLogic sample to the full set of U.S. counties. Specifically, we compare means of a large number of covariates available in the ACS. As we see, CoreLogic counties are a bit more populous, with an average population of about 130,000 vs. about 100,000 for all counties in the U.S. However, on essentially every other covariate they look extremely similar. Perhaps most importantly, CoreLogic counties do not appear to have higher or lower unemployment rates, education rates, or household median incomes. As a result, we suspect our results generalize to non-CoreLogic counties.

Table A.3 – County Covariates (ACS 2009-2014)

Covariate	Mean (All Counties)	Mean (CoreLogic Counties)
Total Population	100125	130809
Percent Age 18 to 29	0.168	0.169
Percent Age 30 to 44	0.196	0.198
Percent Age 45 to 64	0.264	0.262
Percent Age 65 and up	0.137	0.134
Percent Female	0.508	0.507
Percent White	0.628	0.604
Percent Black	0.122	0.114
Percent Hispanic	0.169	0.190
Percent Asian	0.049	0.060
Percent Less than High School	0.138	0.140
Percent High School Degree	0.279	0.268
Percent Some College	0.291	0.293
Percent 4-year College Degree	0.183	0.186
Percent Post-Graduate	0.110	0.112
Unemployment Rate	0.093	0.094
Household Median Income	55790	56826
# Counties	3137	1567