Online Appendix

Online Appendix for "Economic Distress and Voting: Evidence from the Subprime Mortgage Crisis."

Intended for online publication only.

Contents

A.1 Unifying Our Results with Existing Research

Because Healy and Lenz (2017) 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 (2017), 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 (2017) 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 (2017) 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 (2017), 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 (2017) 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 (2017) to our nationwide results. The main explanatory variable that Healy and Lenz (2017) 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-meaned by county, and then rescaled so that it is 0 for the .001 percentile and 1 for the .999 percentile. Healy and Lenz (2017) 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 (2017) is that they come from different empirical designs.[10] While Healy and Lenz (2017) 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 we can provide a series of tests to interrogate

[10] The reader might also note a second difference, which is that the Healy and Lenz (2017) 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 (2017) results do attenuate considerably when the counties are not weighted, as do our 2016/Trump estimates in Table 4. We do not focus too much on this issue, because it seems reasonable to down-weight low population counties.

its key identifying assumption, parallel trends. 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 might be satisfied. These kinds of robustness checks are not possible in the lagged dependent variable framework, which is one reason we prefer the difference-in-differences. 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 per-

[11]This is not the one-unit shift that Healy and Lenz (2017) 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.

centage point, for the foreclosures data, and 0.42 percentage points (42 basis points) for the local economic growth data.

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).

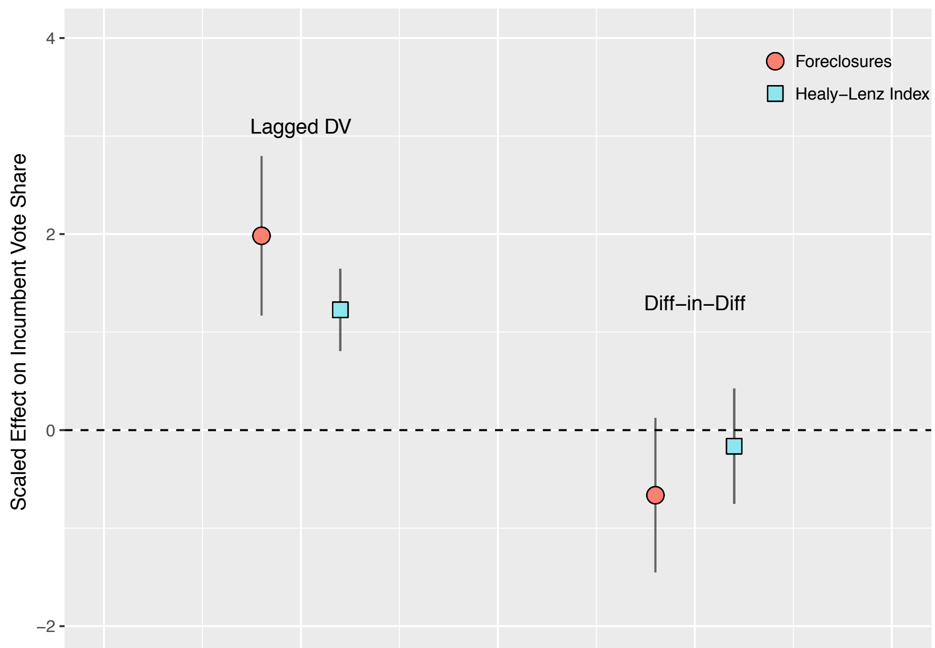

Determining whether we think there are positive effects thus requires choosing between the two empirical strategies. As Angrist and Pischke (2009) explains, we can use the two estimators to bracket the true effect, if we assume that one of the two identifying assumptions is correct. The two foreclosures 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 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 (2017) 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 (2017) 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 (2017), 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 (2017), 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 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 4 in the body of our paper)—but there does not seem to be a general phenomenon underlying these effects.[14]

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

[13] We 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.

[14] In 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 4.

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	4.19	3.57	2.50	0.23
	(0.53)	(0.64)	(1.10)	(0.88)
N	58	58	127	127
State Fixed Effects	–	–	Yes	Yes
Population Weights	No	Yes	No	Yes

Robust standard errors clustered by county in parentheses.

A.2 Using a Lagged Dependent Variable Approach

In this section, we present versions of our main county-level results using a lagged dependent variable approach. To do so, we estimate equations of the form

$$Dem\ Vote\ Pct_{it} = \beta_0 Dem\ Vote\ Pct_{i,t-1} + \beta_1 Foreclosures_{it} + \beta_2 Foreclosures_{it} \cdot Dem\ Inc_{it} + \delta_t + \epsilon_{it}, \tag{3}$$

where $Dem\ Vote\ Pct_{i,t-1}$ represents county i's Democratic vote share in the previous election. To interpret the coefficient on the interaction term, β_2, as causal, we need to make a conditional independence assumption, where the expected value of the potential outcomes under control are the same for counties with both high and low foreclosure rates, after conditioning on the lagged outcome and other covariates. As discussed earlier, we use a difference-in-differences approach, which is consistent if the parallel trends assumption is satisfied. If we assume one of the two designs' identifying assumption is correct, we can use the estimtes from the two designs to bracket the true effect (Angrist and Pischke 2009; Ding and Li 2019).

Table A.2 – Effects of Housing Foreclosures on Presidential Elections, Lagged Dependent Variable Approach, 2004–2016.

	Dem Pres Vote Pct (0-100)	
	(1)	(2)
Dem Pres Vote Pct (t-1)	1.07	0.99
	(0.01)	(0.01)
Foreclosures Per 1,000 People	0.02	0.02
	(0.08)	(0.07)
Foreclosures × Inc Party	-0.33	-0.30
	(0.08)	(0.07)
N	9536	9540
# Counties	2839	2840
State-Year Fixed Effects	Yes	No
Pop Decile-Year Fixed Effects	No	Yes
Population Weights	Yes	Yes

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

We re-estimate the first two columns of Table 3 in the main body of the paper, but instead use the lagged dependent variable model, and we show those results in Table A.2. As we can see, the estimates are negative, but substantively small. In column 1, an increase of 1 foreclosure per 1,000 people leads to about a 0.33 percentage point decrease in the incumbent party's performance. While this decrease is statistically significant, it is substantively small. It would require more than a 3 standard deviation increase in the county-demeaned foreclosure rate to affect the incumbent vote share by a full percentage point. We interpret these effects as evidence that the link between local economic conditions and incumbent party performance is, in general, very modest.

Next, we again re-create the first two columns of our other county-level results in the main body of the paper, but instead use the lagged dependent variable approach. In Table A.3 we re-estimate Table 4, and in Table A.4 we re-estimate Table 5. In both cases, the estimates from the difference-in-differences approach and the lagged dependent variable approach are quite similar – counties with large increases in foreclosures increased their vote share for Trump in 2016, and in 2016 these effects were especially large in places that experienced large increases in foreclosures in the six months leading up to the election compared to previous years.

Table A.3 – Effects of Housing Foreclosures on Presidential Elections, Lagged Dependent Variable Approach, 2004–2016: Testing for Trump-Clinton Effects.

	Dem Pres Vote Pct (0-100)	
	(1)	(2)
Dem Pres Vote Pct (t-1)	1.07	0.99
	(0.01)	(0.01)
Foreclosures Per 1,000 People	0.11	-0.00
	(0.06)	(0.05)
Foreclosures × 2016	-1.70	-1.16
	(0.23)	(0.15)
N	9536	9540
# Counties	2839	2840
State-Year Fixed Effects	Yes	No
Pop Decile-Year Fixed Effects	No	Yes
Population Weights	Yes	Yes

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Main effect for 2016 is absorbed by fixed effects.

Table A.4 – Effects of Recent Housing Foreclosures on Presidential Elections, Lagged Dependent Variable Approach, 2004–2016: Testing for Trump-Clinton Effects.

	Dem Pres Vote Pct (0-100)	
	(1)	(2)
Dem Pres Vote Pct (t-1)	1.07	0.99
	(0.01)	(0.01)
Foreclosures Per 1,000 People	0.39	0.24
	(0.10)	(0.09)
Foreclosures × 2016	-4.35	-3.92
	(0.43)	(0.35)
N	9181	9185
# Counties	2818	2819
State-Year Fixed Effects	Yes	No
Pop Decile-Year Fixed Effects	No	Yes
Population Weights	Yes	Yes

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Main effect for 2016 is absorbed by fixed effects.

A.3 Checking for Pre-Trending

In this section we carry out a placebo test to examine, albeit indirectly, the parallel trends assumption. We construct leads of the foreclosure rate variable to see if future foreclosures affect current presidential vote share. Finding a large coefficient on this lead would suggest there might be pre-trending.

To carry out this test, we code a foreclosure rate lead variable, which takes the value the county's foreclosure rate in time $t + 1$. We show these results in Table A.5, which do not show evidence of pre-trending. The coefficients on these leads are statistically insignificant, but more importantly, substantively small and reasonably precisely estimated. Again, we interpret this as suggestive evidence that the parallel trends assumption might hold.

Table A.5 – Effects of Future Foreclosures on Current Presidential Elections, Parallel Trends Test, 2004–2012.

	Dem Pres Vote Pct (0-100)	
	(1)	(2)
Foreclosures Per 1,000 People	0.01	0.07
	(0.10)	(0.09)
Foreclosures × Inc. Party	0.13	-0.07
	(0.06)	(0.06)
Foreclosures ($t+1$)	-0.04	-0.09
	(0.09)	(0.09)
N	6275	6278
# Counties	2247	2248
State-Year Fixed Effects	Yes	No
Pop Decile-Year Fixed Effects	No	Yes
Population Weights	Yes	Yes

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

A.4 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 occured 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, afirst 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.5 Assessing the Generalizability of the CoreLogic Data

The CoreLogic data covers about 90% of the counties in the United States. Although this lack of complete coverage does not induce any obvious bias into our statistical analyses, we would like to know how specific the resulting estimates are to the CoreLogic counties versus all U.S. counties. Accordingly, Table A.6 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 American Community Survey (ACS). As we see, CoreLogic counties are very similar in population, with an average population of 108,391 compared to 100,099 for all counties. Moreover, the set of CoreLogic counties and the overall set of counties are similar on essentially every covariate in the ACS. 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.6 – County Covariates (ACS 2009-2014)

Covariate	Mean (All Counties)	Mean (CoreLogic Counties)
Total Population	100099	108391
Percent Age 18 to 29	0.168	0.168
Percent Age 30 to 44	0.196	0.196
Percent Age 45 to 64	0.264	0.264
Percent Age 65 and up	0.137	0.137
Percent Female	0.508	0.508
Percent White	0.628	0.626
Percent Black	0.122	0.123
Percent Hispanic	0.169	0.171
Percent Asian	0.049	0.050
Percent Less than High School	0.138	0.137
Percent High School Degree	0.279	0.277
Percent Some College	0.291	0.291
Percent 4-year College Degree	0.183	0.184
Percent Post-Graduate	0.110	0.111
Unemployment Rate	0.093	0.093
Household Median Income	55793	55985
# Counties	3136	2850

Table A.7 – Effects of Housing Foreclosures on Presidential Elections for Small Counties, 2004–2016.

	Dem Pres Vote Pct (0-100)			
	(1)	(2)	(3)	(4)
Foreclosures Per 1,000 People	-0.27	-0.31	-0.19	-0.67
	(0.11)	(0.14)	(0.10)	(0.17)
Foreclosures × Inc Party	-0.10	0.12	-0.09	0.22
	(0.09)	(0.10)	(0.07)	(0.11)
N	4625	4632	4625	4632
# Counties	1412	1413	1412	1413
County Fixed Effects	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes
Population Weights	Yes	Yes	Yes	Yes

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

A.6 Effects in Small Counties

In this section, we examine the effect of foreclosures in small counties. We do not find effects of foreclosures on aggregated vote choice generally, but because foreclosures are felt by a small number of people, it could be difficult to detect effects of foreclosures in large counties. If foreclosures are a better measure of the typical experiences in small counties, we might be most likely to pick up an effect when we subset the analysis to small counties. In Table A.7, we estimate the effect of foreclosures just for small counties, defined as those with a 2003 population at or below the median. Much like the results using the full sample of counties, we find null results across specifications. In Table A.8, we estimate the effect of foreclosures in House and Senate races. Again, we find no evidence that voters in small counties reward or punish House or Senate incumbents based on housing foreclosures. The largest coefficient is in the fourth column for Senate incumbents, but the coefficient is in the opposite direction that the economic voting literature would predict.

Table A.8 – Effects of Housing Foreclosures on Legislative Elections for Small Counties, 2002–2016.

	Dem Senate Vote Pct (0-100)			
Foreclosures Per 1,000 People	-0.05	-0.54	-0.04	-0.91
	(0.08)	(0.27)	(0.09)	(0.40)
Foreclosures × Inc Party	-0.11	0.46	-0.08	1.46
	(0.07)	(0.20)	(0.06)	(0.26)
N	6080	6092	6080	6092
# Counties	1473	1474	1473	1474
	Dem House Vote Pct (0-100)			
Foreclosures Per 1,000 People	-0.10	0.29	-0.05	0.19
	(0.24)	(0.28)	(0.29)	(0.26)
Foreclosures × Inc Party	0.04	-0.52	0.08	-0.30
	(0.25)	(0.26)	(0.28)	(0.23)
N	9059	9075	9059	9075
# Counties	1560	1561	1560	1561
County Fixed Effects	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes
Population Weights	Yes	Yes	Yes	Yes

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

Table A.9 – Effects of Housing Foreclosures on Presidential Elections, County Level, 2004–2016.

	Dem Pres Vote Pct (0-100)			
	(1)	(2)	(3)	(4)
Foreclosures Per 1,000 People	-0.50	-0.16	-0.11	-0.30
	(0.06)	(0.06)	(0.05)	(0.06)
Foreclosures × Inc Party	0.16	0.00	-0.13	0.08
	(0.05)	(0.04)	(0.03)	(0.04)
N	9369	9373	9369	9373
# Counties	2671	2672	2671	2672
County Fixed Effects	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes
Population Weights	No	No	No	No

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Inc Party is 1 for Dem, -1 for Rep. Main effect for Inc Party is absorbed by fixed effects.

A.7 Effects Without County Population Weights

In this section, we report estimates without weighting by county population. In the main text, we weight all of our estimates by county population. As Table A.9 shows, the null effects of foreclosures on presidential elections (from Table 3) hold when we do not use county population weights.

Next, we show that the Trump effects in 2016, which we show in Tables 4 and 5 for all foreclosures and recent foreclosures, respectively, are somewhat sensitive to the decision to weight by county population. In Table A.10, we test for Trump effects using all foreclosures in the electoral cycle and without weighting by county population. The estimates are all still negative, but some are null. Compared to 4, where the point estimates are more negative, we can conclude that if foreclosures led voters to punish Clinton and reward Trump, this behavior was concentrated in more populous counties.

There are also larger effects in the weighted specifications for recent housing foreclosures. In Table A.11, we test for a Trump effect using only foreclosures in the last six months before the election, but we do not weight by county population. The estimates, while still all negative and statistically significant, are much smaller in magnitude than in the corresponding Table 5, where we use population weights and find massive effects of foreclosures on punishing Clinton and rewarding Trump. Again, this suggests that, if foreclosures in the last six months before the election led to voters to support Trump, these effects were particularly concentrated in more populous counties.

Table A.10 – Effects of Housing Foreclosures on Presidential Elections, County Level, 2004–2016: Testing for Trump-Clinton Effects.

	Dem Pres Vote Pct (0-100)			
	(1)	(2)	(3)	(4)
Foreclosures Per 1,000 People	-0.33	-0.13	-0.27	-0.27
	(0.04)	(0.04)	(0.04)	(0.05)
Foreclosures × 2016	-0.32	-0.16	-0.74	-0.23
	(0.10)	(0.09)	(0.10)	(0.10)
N	9369	9373	9369	9373
# Counties	2671	2672	2671	2672
County Fixed Effects	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes
Population Weights	No	No	No	No

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Main effect for 2016 is absorbed by fixed effects.

Table A.11 – Effects of Recent Housing Foreclosures on Presidential Elections, County Level, 2004–2016: Testing for Trump-Clinton Effects.

	Dem Pres Vote Pct (0-100)			
	(1)	(2)	(3)	(4)
Foreclosures Per 1,000 People	-0.52	-0.22	-0.22	-0.25
	(0.07)	(0.06)	(0.05)	(0.07)
Foreclosures × 2016	-0.53	-0.26	-0.78	-0.25
	(0.11)	(0.09)	(0.11)	(0.11)
N	8973	8977	8973	8977
# Counties	2609	2610	2609	2610
County Fixed Effects	Yes	Yes	Yes	Yes
State-Year Fixed Effects	Yes	No	Yes	No
Pop Decile-Year Fixed Effects	No	Yes	No	Yes
County Linear Trends	No	No	Yes	Yes
Population Weights	No	No	No	No

Standard errors generated from 1,000 iterations of a county-level block bootstrap procedure. Main effect for 2016 is absorbed by fixed effects.