Relieving Financial Distress Increases Voter Turnout: Evidence from the Mortgage Market*

Haoyang Liu, W. Ben McCartney, Rodney Ramcharan, Calvin Zhang, and Xiaohan Zhang†

This Version: April 12, 2025

Abstract

Borrowers who refinanced mortgages between 2009 and 2012, a period marked by mortgage distress and dislocated housing markets, but also falling interest rates, were more likely to vote in the 2012 general election than similar borrowers who did not refinance. We exploit an eligibility cutoff in the Home Affordable Refinance Program (HARP) to identify a causal relationship. Consistent with the resource model of voting, the effect of refinancing on turnout is strongest among borrowers with lower incomes and larger debt service reductions. Our findings shed new light on an important channel linking economic conditions and political outcomes.

JEL Classification: D12, D14, D72, E43, H31, R20

Keywords: Household Finance, Mortgages, Interest Rates, Political Participation, Voter Turnout

*We thank Manuel Adelino, Nathaniel Baum-Snow, Simcha Barkai (discussant), John Cochrane, Olivier Coibion, Jim Conklin, Andreas Fuster, Bob Hunt, Stephanie Johnson, Lauren Lambie-Hanson, Jeanna Kenney (discussant), Ben Keys, Samuel Kruger, David Low, Waldo Ojeda, Tess C. Scharlemann, Antoinette Schoar, Tien Foo Sing, Toni Whited, and Tony Zeyer; and seminar participants at the Junior Household Finance Seminar, the 2023 Northeastern University Finance Conference, the AREUEA virtual seminar, and the 2024 ASSA-AREUEA Conference for helpful comments and suggestions. Claire Nelson provided excellent research assistance.

†Liu: Federal Reserve Bank of Dallas, haoyang.liu@dal.frb.org; McCartney: McIntire School of Commerce, University of Virginia, ben.mccartney@virginia.edu; Ramcharan: Marshall School of Business, University of Southern California, rodney.ramcharan@marshall.usc.edu; Zhang, Calvin: University of Oregon, calvinz@uoregon.edu; Zhang, Xiaohan: Federal Reserve Bank of Dallas, xiaohan.zhang@dal.frb.org. Disclaimer: This paper represents preliminary research that is being circulated for discussion purposes. The views expressed are solely those of the authors and do not necessarily reflect the views of the the Federal Reserve Bank of Dallas or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

1 Introduction

Does household debt distress influence political participation? Financial crises and other sources of economic dislocation increase household debt distress, and both historical narratives and more recent statistical evidence suggest that greater debt distress among households can shape political outcomes across a range of dimensions.[1] One important dimension centers on households' participation in elections, as disparities in political participation can shape the overall size of government and the focus of public programs.[2] The impact of debt distress on participation is however ambiguous. Some theories note that greater financial distress can depress political participation, for heavily indebted households might lack the resources to engage politically relative to those with less debt distress. But other theories hypothesize that debt distress might either increase or leave participation unchanged.[3] This lack of consensus reflects the complex motivations behind political engagement and the severe endogeneity problems that arise when simply correlating changes in economic status and civic engagement.

In this paper, we turn to the U.S. mortgage market in the years between 2009 and 2012, and examine the impact of mortgage refinancing on political participation. During our sample period, sharply falling mortgage rates created significant opportunities for homeowners to reduce their debt burdens and relieve financial distress. Yet pervasive frictions in credit markets limited access to refinancing for many households, inducing substantial variation in debt distress and welfare among homeowners. Housing costs are a major component of household spending and changes in the mortgage debt service burden have dramatic implications for most households' disposable incomes and overall financial health.[4] For those able to refinance during this period, monthly debt service pay-

[1]Some historical narratives suggest that debt distress can fuel populist movements on both the left, as in the case of agrarian movements in the U.S. in the early 20th century that sought to expand bank credit for debt distressed farmers (Rajan and Ramcharan, 2011), and on the right, as in the case of Nazi Germany (Evans, 2005). Recent empirical studies examine political polarization and right wing movements after the 2008-2009 financial crisis (Gyöngyösi and Verner, 2022) and more generally (Funke et al., 2016).

[2]See, for example, Avery (2015); Bartels (2016); Franko et al. (2016); Gilens (2012); Leighley and Nagler (2013); Lijphart (1997); Stigler (2021).

[3]Theoretically, the causal nature of the relationship between an individual's financial health and their turnout decision is not obvious (Rosenstone, 1982). Improved financial health might lead to greater participation by increasing individuals' social capital (Wolfinger and Rosenstone, 1980), making the costs associated with voting more affordable (Brady et al., 1995), or boosting individuals' sense of political efficacy (Beaumont, 2011). Participation may decrease if increases in economic well-being cause individuals to feel less connected to or reliant on government programs, thus reducing the perceived stakes of voting (Brunner et al., 2011). Or, finally, if people vote as a response to social pressures or out of a sense of duty, then their personal financial situations might be irrelevant (Gerber et al., 2003; Plutzer, 2002). The empirical evidence does little to settle the debate (Margalit, 2019; Smets and Van Ham, 2013).

[4]The literature on refinancing and consumption is large. A selected list of references include: Agarwal et al. (2023);

ments typically fell by over $300, generating present value interest savings in the tens of thousands of dollars while boosting consumption, voluntary deleveraging, and credit scores (Abel and Fuster, 2021; Fuster and Willen, 2017).

To understand the relationship between reduced debt burdens and political participation, we begin by creating an individual-level dataset that jointly describes mortgage and voting decisions over time. Our final sample includes borrowers who originated fixed-rate mortgages between 2005 and 2008 and were eligible to vote in both the 2008 and 2012 general elections. We then ask if those who refinanced their mortgages between 2009 and 2012, decreasing their rates by an average of 161 basis points, were more likely to vote in the 2012 election than those who did not. We find that they were. Specifically, using our final sample of more than ten million people, we document that borrowers who refinanced were 2.3 percentage points more likely to vote in the 2012 election than otherwise similar borrowers who did not refinance.

Of course, because voting and rate refinancing are both active decisions, selection bias can plague attempts to measure the impact of lower housing costs on voting. Poorer, less educated citizens are on average less politically engaged. At the same time, the refinancing literature observes that these agents are more likely to refinance sub-optimally or not at all, incurring substantial welfare costs (Andersen et al., 2020). Poorer less educated citizens are also more likely to attrite out of voter rolls, compounding endogeneity concerns. Conversely, higher income, better educated individuals are both more likely to vote and optimally refinance in order to reap the large economic gains on offer during our sample period. Thus, a positive association between refinancing post-2008 and voting in 2012 *could* causally reflect the variation in welfare gains from refinancing. But latent income, education, or cognitive factors that allowed some homeowners to refinance optimally could also explain the positive association we observe.

To help address this endogeneity concern, we use our detailed individual- and mortgage-level data to control for many factors hypothesized to influence voting and refinancing. Specifically, the basic model estimates the effect of refinancing on voting by comparing people of the same sex, similar age, and affiliated with the same party; living in homes of similar size and age in the same county; and who originated similar mortgages of similar amounts at the same time. Furthermore, since we see complete voting histories, we can also control for each individual's participation in the 2008, "pre-

Andersen et al. (2020); Campbell (2006); Di Maggio et al. (2017); Ganong and Noel (2020); Keys et al. (2016).

period," general election. This is helpful because it allows us to absorb the effects of even unobserved characteristics that drive both participation and refinancing decisions, at least to the extent that these characteristics and their effects are time invariant. Finally, for roughly 60% of our sample, we observe each households' income at the time their mortgage was originated. Using this sample, we show that controlling for income leaves the main result unchanged.

We next use the heterogeneity in the impact of refinancing on voting to make progress on the potential mechanisms at work. The logic of this approach builds on the argument that because voting requires time, money, and effort, resource theories of voting predict that citizens that gain access to additional resources will become more likely to vote (Brady et al., 1995; Schlozman et al., 2018, 2012; Verba and Nie, 1987). In our setting, then, the resource theory of voting predicts that any extra income from refinancing that relaxes resource constrains should have the largest marginal impact on voting among poorer citizens. When we allow the treatment effect to vary as a function of income, we see that the effect of refinancing has a roughly 30% smaller effect size on those households making more than \$125,000 a year relative to those making less than \$50,000 a year. We next allow treatment effects to vary with debt service savings and again document results consistent with the resource theory of voting. Borrowers who decreased their mortgage rate by less than 100 basis points were 1.7 percentage points more likely to participate than otherwise comparable borrowers who did not refinance. But borrowers who saved the most – those who decreased their mortgage rate by 200 basis points or more – were a full 2.5 percentage points more likely to participate.

These results not only shed light on potential channels, but also help rule out an alternative explanation of our main finding: That relative to those who fail to refinance, those who do are simply more literate or more engaged with all institutions and thus more likely to vote. This narrative predicts that neither borrowers' incomes nor debt service savings should mediate the effects of refinancing on voting. Our heterogeneity results are thus inconsistent with this alternative explanation. That said, it is hard to rule out all possible alternative narratives based on unobservables. Each election could be *sui generis* for example, so that controlling for past participation may not fully absorb the current unobservable determinants of participation.

To further isolate the causal effects of refinancing on political participation, we build on the work of Abel and Fuster (2021) and Karamon et al. (2017) and use the plausibly exogenous variation in refinancing eligibility from the Home Affordable Refinance Program (HARP). HARP was announced by the Obama Administration in March 2009 and aimed to help highly leveraged homeowners refinance their mortgages into lower-rate fixed-term mortgages. HARP's benefits were large. Debt service savings were on average nearly $200 per month and the program significantly decreased default risk. Moreover, a feature of the program's design helps us address endogeneity concerns: Only those loans purchased by Fannie Mae or Freddie Mac (the "government-sponsored enterprises" or GSEs) on or before May 31, 2009, regardless of when borrowers took out the loan, were HARP eligible. Since the date that the lending institution sold the loan to the GSE is completely outside the borrower's control, whether an individual's loan was sold before the cutoff date, and is thus eligible for HARP, is a potentially useful instrument for whether or not they refinanced their mortgage.

To implement the HARP identification strategy, we build a sample of GSE-eligible loans originated in the first half of 2009 and conduct a series of balance tests which show, consistent with the exclusion restriction, that borrowers whose loans were purchased before the cutoff date are indistinguishable along a number of dimensions, including participation in the 2008 general election, from those whose loans missed the cutoff. Furthermore, those eligible for HARP refinancing were dramatically more likely to refinance their mortgage than those who were not, satisfying the relevance condition. Using the fact that HARP eligibility – determined by the GSE cutoff date – is likely exogenous in our sample and increases the odds of refinancing among eligible homeowners, we return to the question of whether refinancing increases voter participation using an intention to treat identification strategy.

Our central finding is that after correcting for endogeneity by instrumenting for refinancing with HARP eligibility, refinancing increases the likelihood of voting in the 2012 election by about 4.4 percentage points among highly levered borrowers. This higher magnitude echoes the heterogeneity results in the main tests and reflects the characteristics of the compliers in this "intention to treat" setting. Borrowers that "complied" with HARP – those who both required HARP to refinance and would not have refinanced without it – were precisely those borrowers that experienced the largest increases in unplanned leverage during the housing bust, faced the largest frictions to refinancing, and were the most liquidity constrained during the sample period. It is for these borrowers that HARP eligibility likely improved financial health the most. That is, the larger IV point estimates for this sub-sample of borrowers is consistent with the predictions from the resource model of voting. There are of course many other theories of voting, some of which overlap, and our results do not suggest that other motives are unimportant.[5]

Our paper adds to the growing literature in financial economics that investigates how changes to individuals' financial health affect their turnout in elections. As in this paper, McCartney (2021) merges deeds data and voter rolls, though just in the state of North Carolina, and finds that registered voters exposed to falls in house prices are less likely to vote. In a very different political context, Gyöngyösi and Verner (2022) use zip code level data from Hungary and find that greater debtor distress among voters increases support for the populist far right while modestly reducing turnout.[6] Our paper contributes in two ways to this nascent literature. First, we use millions of individual-level voter records data, including a rich set of individual- and mortgage-level characteristics, to generalize the results and address better some of the thorny identification challenges that can affect inference in this literature. For example, evidence from house price falls in North Carolina may not apply to other states and has little say about the impact of refinancing and debt relief on voting, while zip code level correlations can be subject to alternative interpretations. Second, we use the plausibly exogenous HARP experiment to make further progress on these identification challenges and to highlight the sizable heterogeneity in the impact of changes in household debt distress on voting.

We also contribute to the literature on how mortgage refinancing affects households (see, e.g., Agarwal et al. (2023); Andersen et al. (2020); Campbell (2006); Di Maggio et al. (2017); Ganong and Noel (2020); Keys et al. (2016)). Especially related to our paper are Di Maggio et al. (2017) and Abel and Fuster (2021). Di Maggio et al. (2017) take advantage of variation in the timing of when households with adjustable rate mortgages saw their rates reset to find that reductions in the size of mortgage payments increases car purchases and voluntary deleveraging. Abel and Fuster (2021),

[5]Some of these theories range from the the purely rational — where individuals weigh the benefits of their preferred candidate winning and the likelihood that their vote is the decisive one with the cost of voting — to the largely behavioral — where voting is a social activity subject to peer influences, norms, and habits (Blais, 2000; Cancela and Geys, 2016; Downs, 1957; Smets and Van Ham, 2013).

[6]These papers, of course, belong to the much broader political science and economics literature on the resource model of voting (Blais, 2000; Margalit, 2019; Smets and Van Ham, 2013). Looking first to labor markets, Burden and Wichowsky (2014) and Cebula (2017) both find that unemployment drives people to the polls. In contrast, Emmenegger et al. (2015) show that unemployment depresses voter participation. The effects of unconditional cash transfers yield similarly mixed results. Hirvonen et al. (2024) and Loeffler (2023) document that the random recipients of a universal basic income (UBI) in Finland and Alaska, respectively, were more likely to vote in subsequent elections. Geys and Sørensen (2025) take advantage of the exogeneity of lottery wins in Norway to find that positive wealth shocks increase turnout. On the other hand, Akee et al. (2020), Brännlund et al. (2024), Broockman et al. (2024), and Pop-Eleches et al. (2012) find turnout unresponsive to unexpected unconditional cash transfers among adult Native Americans in North Carolina, lottery winners in Sweden, unconditional cash transfers among low-income adults in Illinois and Texas, and quasi-randomly assigned voucher awardees in Romania, respectively.

off of whose HARP research design we build, find that refinancing reduces borrower default rates while also inducing borrowers to access other types of debt, like auto loans, consistent with increases in consumption. Our work adds to this important literature by showing that the bundled treatment received following a decline in mortgage payments has impacts not just on households' pocketbooks, but on households' participation in the political process and thus society more broadly.

Finally, a large literature has examined the political origins of housing policy and mortgage credit in the U.S., especially in the wake of the 2008-2009 financial crisis (see, for example, Calomiris and Haber (2014), Fieldhouse et al. (2018), Green and Wachter (2005), Mian et al. (2010), and Rajan (2011)). To our knowledge, ours is the first paper to examine the possible implications of improved mortgage credit access on voters' political participation, showing that housing and mortgage market policies that improve the welfare of homeowners may indeed increase political participation. In short, our work further demonstrates how consumer financial markets and politics are deeply intertwined.

In what follows, we first describe the data and sample creation in Section 2. Section 3 presents the main results and Section 4 discusses how the main effect varies as a function of voter and borrowers attributes. In Section 5, we present the results of the HARP research design. We conclude in Section 6 with a discussion of potential channels and the paper's primary contributions.

2 Data

2.1 Data Sources

2.1.1 Voter Data

We obtain voter registration and historical voter participation data from the non-partisan data provider Labels & Lists, Inc. (hereafter "L2").[7] The voter registration data is a snapshot of all people registered to vote and includes each voter's full name, complete mailing address, party affiliation, date of registration, date of birth, and sex.[8] The historical voter participation data lists each election in which the voter cast a ballot.

One important limitation of the voter registration data is that it cannot be used as a census of

[7] Our access to this data comes via a licensing agreement between L2 and the University of Virginia. For more information on L2, visit https://www.l2-data.com/datamapping/.

[8] Party affiliation is state reported except in the following states where it is modeled by L2: AL, GA, HI, IL, MI, MN, MO, MT, ND, SC, TX, VA, VT, WA, WI.

current residents. There are two reasons for this. First, not all people are registered to vote. They may not be legally allowed to vote or may have never opted to register. Second, most voter rolls are purged rarely which means that many people who have long since passed away or moved to different states are still registered to vote. If L2 were our only data source, then this would be a crippling limitation since we would not know whether a voter does not participate because they abstained, because they died, or because they were living, registered, and participating somewhere else. But, by looking to deeds that tracks home sales, we can make relatively safe assumptions about who is truly registered and eligible to participate in any election.

A second limitation of the voter data is what we observe regarding participation. Specifically, we know only whether or not a voter submitted a ballot for the election that took place on a given day. We see, for example, that a voter submitted a ballot on November 4, 2008. But, since the ballots themselves are confidential, neither we nor anyone else can legally know which specific elections – presidential, congressional, local – each voter voted in nor for whom they voted. In this paper, therefore, we focus on participation decisions in the 2008 and 2012 general elections. Another benefit of focusing on these two elections is that, across all jurisdictions, the presidential candidates are the same on each election date, and the elections take place at the same time. This helps ensure a level of standardization across the sample.

Despite these limitations, the voter data have useful attributes for answering the questions we ask in this paper. First, with few exceptions, the coverage of registered voters and their complete participation histories is universal, allowing us to examine participation decisions of the entire electorate in the United States over time. Second, since the data is at the individual-level, we can reach much more precise conclusions about potential channels than allowed by aggregate data. And, finally, the public availability of the data, which includes voters' full names and addresses, means we can merge in information from other public data sources, like county deeds and assessors' offices, to paint detailed pictures of households' financial situations and civic engagement.

2.1.2 Mortgage Data

We use information on property transactions and mortgage originations from county deeds and assessors' records. These publicly available data are collected and processed by CoreLogic Solutions, LLC (hereafter "CoreLogic") and come in three main parts.[9] First, we observe the complete history of ownership changes. We see the date of the transaction, the nature of the transaction, the sale price (if applicable), and the names of the new and previous owners. Second, we observe every time a mortgage that uses the property as collateral is originated. Here we see the amount of the mortgage, the date the mortgage was originated, and the names of the borrowers. Third, we see parcel-level characteristics used by local tax assessors' offices including full address, square footage, and year built.

These data are ideally suited to understanding the effects of mortgage refinancing on civic engagement for two reasons. First, we observe not only original purchase transactions and details of any purchase mortgage, but also all subsequent refinancing activity. Second, the ownership transfer data includes not just arms length sales, but also quitclaims, affidavits of death, and inter-family deed transfers among many other transaction types. This wealth of information allows us to create a very accurate and detailed picture of who owns a given parcel at any point in time.[10] We can therefore overcome the central limitation of the voter data – the often stale nature of registration data that does not reflect when individuals pass or move away.

2.1.3 Loan Ownership Data

Implementing our HARP eligibility research design requires that we know whether the mortgage is held by either of the GSEs as of the end of May 2009 – which we use as our instrumental variable for refinancing. Since there could be unobservable differences in loan characteristics that affected a bank's decision of whether to sell a loan or keep it on balance sheet, especially during this period of falling home prices, we restrict the sample we use for this research design to just those loans that were owned by the GSEs as of the end of December 2009. Restricting the sample to loans sold to the GSEs guarantees that, within our sample, HARP eligibility was simply a function of timing rather than loan characteristics. To create these two dummy variables – held by a GSE as of May 31st, 2009 and held by a GSE as of December 31st, 2009 – we use loan performance data from ICE, McDash.

[9] Our access to this data comes via a licensing agreement between CoreLogic and the University of Virginia's White Ruffin Byron Center for Real Estate. For more information on CoreLogic, visit https://www.corelogic.com/360-property-data/.

[10] A limitation of the deeds data is that we observe only the names of each property's owners, not the names of its residents. We therefore focus our paper on owner-occupants, defined as property owners who are also registered to vote at that address, leaving an analysis of the political behaviors of landlords and renters to future work.

We present more details on how we incorporate this data in Section 5.

2.1.4 Other Data

We do not observe the rate on each loan, so we assign each loan a rate equal to the prevailing mortgage rate during the week it was originated. We use publicly available current mortgage rates data from Freddie Mac.[11] To measure borrower's income, we match loans in our sample to mortgage applications in the publicly available Home Mortgage Disclosure Act (HMDA) data.[12] We use the crosswalk described in Bayer et al. (2022) and achieve an overall match rate of roughly 60%. Finally, we determine if a loan is under the conforming loan limit by comparing it to the publicly available conforming loan limit values published by the Federal Housing Finance Agency (FHFA).[13]

2.2 The Final Sample

Most simply, our empirical strategy compares the participation decisions of people who refinanced to those who did not. Because our focus is on the impact of refinancing on political participation, selecting a sample from the population of homeowners with fixed rate mortgages is economically appropriate. To this end, we use the deeds data to identify all individuals that meet the following three criteria. First, we limit the sample to those households who have a mortgage to refinance. Specifically, we include only those borrowers whose most recent first lien mortgage as of January 1, 2009 was originated some time after January 1, 2005. We include both purchase mortgages and refinances of previous mortgages. Second, we require that this mortgage have a fixed interest rate. Fixed-rate mortgages represent the majority of the mortgage market and provide a more consistent basis for analyzing the effects of refinancing than adjustable-rate mortgages due to their predictability and standardization.

Our third criteria is to limit the sample to just those borrowers who had not sold their house or otherwise transferred ownership as of the 2012 election. This third requirement is needed in order to match homeowners to state voting records. Moving can be for random reasons, such as divorce or death of a family member. But foreclosure, job loss, or other economic factors can also precipitate moving and possibly cause non-random attrition out of the sample. In Section 5, we discuss how

[11] https://www.corelogic.com/360-property-data/, accessed September 17, 2024.
[12] https://www.consumerfinance.gov/data-research/hmda/, accessed April 23, 2024.
[13] https://www.fhfa.gov/data/conforming-loan-limit-cll-values, accessed November 17, 2021.

non-random attrition leads to sample selection bias and what we do to address these concerns. These sample restrictions leave us with a final sample of just over 14 million mortgages whose borrowers we attempt to find in L2.

We merge in information from the L2 voter data using the borrower/registered voter's complete mailing address, first name, and last name.[14] In most cases, we require a match on complete address and full name, but we also allow some flexibility to pick up spouses who might not match perfectly. We do so in two ways. First, we include in the sample voters whose first name is on the deed/mortgage at that address if that address also has another person with a full name match. This captures cases where, for example, John Doe is on the deed and registered to vote and his wife is on the deed as Jane Doe and registered to vote as Jane Smith. Second, in some cases one spouse is left off the mortgage. We therefore include voters whose last name is the same as the voter with the full name match to the deed/mortgage and whose age is within fifteen years of the voter with a full name match. By including this age restriction, we aim to capture spouses, but not children or parents.[15]

Our merged data set includes nearly 11 million distinct registered voters representing 7 million distinct mortgages (half of the 14 million qualifying mortgages described above). There are three explanations that likely explain most of the unmatched mortgages. First, these mortgages might not be for owner-occupied housing, as in the case of vacation homes or rental properties. Note that this lack of match is actually a strength of our data, not a limitation, since it correctly assigns the mortgage to the voter who holds it. Second, while the homes might be owner-occupied, the owners themselves might be either unable to register, as is the case for non-U.S. citizens, or not interested in registering. Third, we require voters to have been registered and eligible to vote in both the 2008 and 2012 elections which drops voters that either registered for the first time after the 2008 election or unregistered before 2012.[16] Our final dataset is conservatively built along a number of dimensions but nevertheless consists of well over ten million borrowers for whom we observe a very complete picture of mortgage and election participation decisions.[17] We describe this sample in detail in Table

[14]For merging purposes, we define a parcel by its house number, unit number (if applicable), street direction, street name, street mode, street quadrant, and ZIP code.

[15]Our final sample consists of 93.42% full name, or type 1, matches, 1.24% type 2 matches, and 5.34% type 3 matches. Our conclusion are unchanged if we use only type 1 matches.

[16]Analysis of registration decisions is challenging because the reasons that people might not register to vote – no interest, not affordable, non-citizen, registered elsewhere – are often very difficult to disentangle and we leave an investigation of the drivers of registration decisions to future work.

[17]We visualize the creation of this dataset in Appendix B.

1.

TABLE 1 HERE

The first two sections of Table 1 summarize our dependent and independent variables of interest. We see first that among this population of homeowners with recently originated mortgages who then did not move for at least four years, participation in general elections is very high. This is consistent with prior work; (DeSilver, 2022) and (Hall and Yoder, 2022) both document participation rates in presidential elections of nearly 95% among homeowners registered to vote. Also consistent with prior work, less than half of these borrowers refinanced, but those that did decreased their rate by an average of 161 basis points (Keys et al., 2016). The next three sections of Table 1 describe voter's demographic information, the characteristics of their properties, and the characteristics of the mortgages they had as of 2009.

3 The Effect of Refinancing on Voter Participation

3.1 Empirical Framework

We first establish the basic relationship between mortgage refinancing and civic engagement using a simple linear regression with a comprehensive suite of controls. The wealth of information we have at the individual-level that describes voters' demographics, properties, and outstanding mortgages allows us to control for many of the known factors that influence both civic engagement and financial literacy.

Crucially, a key advantage of our panel dataset is that we can observe the participation decisions made by borrowers both after *and before* the period of low-interest rates during which some borrowers refinanced and others did not. This allows us to absorb the effects of those unobserved individual characteristics that influence both refinancing and voting, like wealth or education. At least to the extent that wealthier or better educated individuals are both more likely to refinance their mortgages and to vote, this should be reflected in their 2008 participation decisions. Specifically, we estimate the following specification:

$$
\begin{aligned}
\text{Voted in the 2012 General Election}_i = \beta \times \text{Refinanced Between 2009 and 2012}_i \\
+ \gamma \times \text{Voted in the 2008 General Election}_i + \text{Controls}_i \times \Phi + \text{Party-by-County}_{p,c} \quad (1)
\end{aligned}
$$

where our individual-level control variables include personal demographics (sex, birth year, registration year, and income), property characteristics (year built and property size), and characteristics of the outstanding mortgage (quarter of origination, purpose, amount, whether or not conventional, and whether or not under the conforming loan limit).

3.2 2009-2012 Refinancing is Associated with Increased 2012 Voter Participation

TABLE 2 HERE

We begin by estimating a model with no control variables and just a party affiliation-by-county fixed effect. This result, presented in column (1) of Table 2, is essentially the correlation between having refinanced between 2009 and 2012 and voting in the 2012 election. As expected, the relationship is economically large and statistically significant. Those that refinanced are nearly four percentage points more likely to vote than those who did not. Of course, a number of individual and environmental factors are likely correlated with both refinancing decisions and civic engagement.

In column (2), we control for whether or not the individual voted in the general election in 2008. Persistent civic participation is correlated with education, social capital, parental behavior and a host of other time-invariant factors that might also shape financial literacy and refinancing behavior. Not surprisingly, 2008 participation loads very strongly, with those who participated in 2008 being 37.3 percentage points more likely to vote in 2012 than those who did not. The inclusion of just this control, by controlling for many of the time invariant characteristics that predict both voting and refinancing, decreases the magnitude of the estimate of interest from .0377 to .0287.

In column (3), we present the results of a model that removes the prior participation control and replaces it with our full suite of other personal, property, and mortgage attributes. Intriguingly, while this does reduce our estimated treatment effect, it does so by significantly less than the inclusion of just the 2008 participation control. So, to the extent that our estimated effect in column (1) is biased upwards, the inclusion of the 2008 participation control absorbs more of this than all the rest of the control variables combined. Column (4) includes all of these control variables and produces an estimated treatment effect of .0234. This positive association between refinancing and voting is consistent with the resource constraints theory of voting. To wit, by increasing disposable income, relaxing financial constraints, and relieving debt distress, refinancing between 2009 and 2012 can increase the odds of participation in 2012.[18]

There are many different and overlapping models of voter turnout. While some of these models are in conflict, resource models generally coexist with them and simply predict that relaxing resource constraints have particularly large effects on those citizens who are otherwise near some participation threshold (Arceneaux and Nickerson, 2009; Enos et al., 2014; Hirvonen et al., 2024). These "sometimes voters" stand in contrast to those who always participate in the political process — either out of civic obligation, habit, trust in political institutions, or other factors – the so called "always voters." For these voters, resource models of voting predict that additional income from refinancing should matter less for their participation in the 2012 election.

To test this prediction, we examine the mediating effects of political affiliation. Voters that identify with a particular party – registered Republicans or Democrats – are much more likely to be those "always" voters who vote because of ideology, trust in political institutions, or the myriad other factors identified in the literature on political participation. In contrast, "sometimes" voters whose participation decisions are much less certain are more likely to be unaffiliated with any party. Resource models of voting thus predict that refinancing should have the largest effect on participation in the 2012 election among unaffiliated voters.

TABLE 3 HERE

To test this, columns (1), (2), and (3) of Table 3 estimate the baseline specification separately on sub-samples of Republican, Democrat, and unaffiliated voters. The impact of refinancing between 2009 and 2012 on voting is similar among voters affiliated with a party, but significantly smaller than the estimate based on the sample of unaffiliated voters in column (3). Column (4) nests this test and shows that the effects of refinancing on voting among unaffiliated voters is about 1.9 percentage points higher relative to registered Republicans that refinanced.

[18] Instead of conducting a cross sectional test and controlling for 2008 participation, another possibility is to estimate a difference-in-difference model. We present the results of this analysis in Table A1.

While we acknowledge that political participation is complex and multiple overlapping theories help explain political participation, our results to this point are especially consistent with the predictions of resource models of voting. Of course, the evidence thus far is suggestive and mechanical selection bias can also produce a spurious positive association between refinancing between 2009 and 2012 and voting. We next attempt to parse better the underlying mechanism and address issues of bias.

4 Heterogeneity Tests

In this section, we present a number of heterogeneity results that aim to both validate our main findings and shed more light on a particularly plausible mechanism linking households' financial situations to their civic engagement: the resource model of voting.

4.1 Are Those Who Refinanced Different?

While our models of voter participation have included various controls, the fundamental concern that refinancers are different on salient unobservables — education, awareness of current events, cognitive abilities — remain. One way to gauge the extent of any such bias is to use information on past refinancing behavior. Homeowners that refinanced in response to falling rates before 2009 are likely to be more aware or educated, for example. By including information on past refinancing behavior, we can therefore proxy for these potentially salient homeowner unobservables that drive refinancing behavior, and possibly voting.

TABLE 4 HERE

To implement this test, column (1) of Table 4 restricts the sample to those homeowners we have never seen refinance, while column (2) runs our same regression using the sample of homeowners that refinanced at least once before 2009.[19] The sample in column (2) thus arguably includes the sample of more attentive and financially literate homeowners. Strikingly, the impact of refinancing

[19] To be clear, we do not observe the mortgage decisions that people made at prior residences. Every borrower in our sample has exactly one outstanding, first-lien mortgage at their current residence as of January 1, 2009. This mortgage is a purchase mortgage if the borrowers bought their home with a mortgage between 2005 and 2008 and have not since refinanced. Alternatively, the mortgage is a refinance if the borrowers bought their home at some point prior to 2009 and then refinanced between 2005 and 2008. Note: borrowers who bought a home prior to 2005 and have not since refinanced are not included in our sample. Borrowers whose last refinance was before 2005 are similarly excluded.

between 2009 and 2012 on voting in the 2012 presidential election is about the same across the two groups. Column (3) combines the two samples and formally tests for differences in behavior across the two groups. The differences are tiny, and if anything, the impact of refinancing between 2009 and 2012 on voting in 2012 is slightly smaller among past refinanciers. The evidence in Table 4 thus weighs against the idea that those who refinance and those who do not differ in ways that bias upwards our coefficient of interest.

4.2 Heterogeneity with Respect to Household Income and Savings

Resource based theories of voting make two additional predictions in our setting regarding heterogeneity with respect to income and dollars saved by refinancing. First, regarding income, we should expect to see a larger effect of refinancing among poorer citizens. Because voting requires time, money, and effort, all citizens that gain access to additional resources will become more likely to vote at the margin. But the extra income from refinancing should relax financial constraints, and thus voting constraints, the most among poorer, more financially constrained, citizens. This logic motivates the following two tests. We begin by focusing on the subset of our sample for whom we observe income.

TABLE 5 HERE

To set the stage for these tests, Column (1) of Table 5 re-estimates our main model described by Equation 1 on the sample for whom we observe income. The estimated treatment effect of .0236 is nearly identical to the estimate of .0234 from Table 2, suggesting this sample is representative of the full sample. In column (2), we control for income with four income bins that correspond roughly to quartiles of income. There are two noteworthy findings. First, the estimated average treatment effect is remarkably similar to the estimate in column (1). In other words, controlling for income, surely one of the most obvious omitted variables in Table 2, does not alter our estimate of interest. This should reassure us as to the validity of our main result. Second, the effect of income behaves exactly as predicted by the resource model of voting, with high income voters being markedly more likely to vote than low income voters.

In column (3), we get at the heart of the matter and allow the effect of refinancing to vary with income. We find that the effect of refinancing is much higher, and significantly so, for low income borrowers. This is, as with the main effect estimates in column (2), consistent both with our understanding of household finance and with resource models of voting. Perhaps curiously, the effect of refinancing is positive even on high earners. It seems unlikely that for borrowers making more than $125,000, a smaller monthly mortgage payment will matter that much, on the margin, for purely resource reasons. This suggests that other potential mechanisms are in play.

We thus return to the resource model of voting and investigate another of its predictions: that those borrowers who experienced the biggest welfare gains should also evince the greatest increase in political participation. To evaluate this hypothesis, we compute for each refinancing borrower the difference between the prevailing rate at the time they originated their initial mortgage and the prevailing interest rate on the date that they refinanced. We use this rate gap to proxy for the size of the financial gains from refinancing and then test whether borrowers with a larger rate gap are indeed more likely to vote.

TABLE 6 HERE

Specifically, instead of comparing those who did not refinance to those who did, we split the refinancers into three mutually exclusive groups as a function of how much their rate decreased when they refinanced. In Table 6, we compare those who did not refinance (61% of the sample) to those who saved less than 99 basis points (5% of the sample), those who saved between 100 and 199 basis points (22% of the sample), and those who saved more than 200 basis points (12% of the sample).[20] In column (1) of Table 6, we document that the greater the savings due to refinancing, the larger the effect refinancing had on voting. And these differences are not small. Borrowers who saved fewer than 100 basis points were 1.72 percentage points more likely to vote than those who did not refinance. This treatment effect is 45% larger, at 2.51 percentage points, for those borrowers who saved 200 basis points or more. In column (2), we show, even conditional on income, an increasing relationship between dollars saved and participation likelihood.

This result is both consistent with the resource model of voting and also pushes back against the notion that unobservable differences between refinancers and non-refinancers explains the main result. That is, if it were simply the case that our main results are being driven by differences between borrowers who ever refinance and those who do not, perhaps because those who refinance

[20] We choose round number cutoffs for ease of exposition, but the results are robust to other cutoffs.

are savvier or have deeper ties to their communities, then we would *not* expect to see the effect on voting vary by the dollars saved. But we do. The alternative story thus has to say, for example, that borrowers who timed their refinancing to save 225 basis points instead of 150 basis points are especially savvy and thus especially more likely to vote in 2012. Few theories of voting posit such a refined participation strategy.

We conclude this section by discussing an alternative explanation for the findings in Table 6. A determined skeptic can argue that since rates were mostly falling between 2009 and 2012, the results of that analysis reflect the *salience* or recency of the savings rather than the *amount* of the savings. That is, it might be that voters who refinanced more recently have their lower monthly payment front and center in their mind. And it is this sort of recency effect that makes them more likely to vote, and not actual dollars saved.

The ideal way to address this would be to control for both the time the initial mortgage was originated and the time that the mortgage was refinanced. But, if we control for both of these times, then we have no variation left in the rate drop itself. Instead, we leverage an econometrically useful feature of the 2010-2011 period: Rates on 30-year fixed-rate mortgages were *not* uniformly falling. Indeed, during one spell, prevailing rates increased nearly a full percentage point in just three months starting at the end of 2010. These large swings mean that the precise savings each household got when they refinanced was somewhat outside of their control, shutting down an endogenous timing channel. And since rates were not uniformly decreasing, we can horse-race the dollars saved mechanism with a salience story by comparing people who saved more money further in time with those who saved less money more recently.

FIGURE 1 HERE

To do so, we estimate a model that allows the refinancing treatment effect to vary monthly by adjusting our main model slightly. Instead of comparing those who did not refinance to those who did, we compare those who did not refinance to those who refinanced in the first month of 2009, to those who refinanced in the second month of 2009, and so on. We then plot these monthly effects. At the same time, we plot the prevailing rate each month. We present our findings in Figure 1 and document a striking negative correlation between the prevailing rate in a given month and the effect of refinancing in that month on 2012 participation. That is, in periods when rates were relatively high, and the benefits from refinancing relatively low, the treatment effect of refinancing on future participation was also low. Our finding of a time varying treatment effect that coincides with the path of mortgage rates fits much better with a resource model of voting explanation for our main finding than either endogenous timing or recency narratives.

5 Evidence from a Policy Experiment

In Section 3, we presented evidence that refinancing is associated with an increased likelihood of voting, even after controlling for an extensive set of observable borrower characteristics, including prior (2008) electoral participation. Results from heterogeneity tests further support our hypothesis that debt service reductions following refinancing drives this relationship, consistent with the resource theory of political participation. According to this theory, reductions in debt distress that alleviate financial constraints increase individuals' capacity to participate in political activities such as voting.

However, two threats to the validity of our conclusion remain. First, within the sample, refinancing is still an active choice. Thus, unobserved borrower characteristics correlated with both refinancing decisions and voter participation could bias our estimates upward. For instance, financial sophistication could simultaneously increase refinancing likelihood and political engagement, thereby inflating our estimated effect. Second, our sample itself may be endogenously constructed. To refinance a mortgage, an individual must have a mortgage, and so we pick a sample of individuals with mortgages originated between 2005 and 2008. This is a well-defined sample given the focus of our question. However, by further requiring that these individuals remain in the same home until at least the 2012 election, we introduce the potential for non-random attrition that can bias our estimates. For example, our study focuses on a time period marked by high foreclosure rates (Zhang, 2019), resulting in a sample that excludes individuals who lost their homes or were otherwise forced to move. If the borrowers who remain in our sample are relatively more educated or financially literate, our estimates would likely be biased downward, since our comparison group (non-refinancers) would disproportionately consist of financially stable households more likely to vote regardless of refinancing status.

Although our heterogeneity analyses in Section 4.2 suggest that these biases alone do not fully explain our findings – since effects are clearly stronger among low-income borrowers and those who experience greater financial relief – we acknowledge these limitations. Thus, to better isolate a causal relationship, we exploit quasi-random variation in refinancing opportunities generated by eligibility for the Home Affordable Refinance Program (HARP). The specific institutional details of HARP provide an ideal setting to address these endogeneity and selection concerns directly, as eligibility for refinancing under HARP was plausibly exogenous with respect to voter characteristics. The following analysis leverages this natural experiment to further strengthen our causal interpretation of the refinancing-voting relationship.

5.1 The Home Affordable Refinance Program

The Home Affordable Refinance Program (HARP) was introduced by the Obama Administration in March 2009 with the explicit goal of enabling highly leveraged homeowners to refinance into lower-rate, fixed-term mortgages (Abel and Fuster (2021), Ehrlich and Perry (2015), and Karamon et al. (2017)). Previous research documents substantial economic benefits: HARP refinancing generated average monthly debt service savings of approximately $174 and reduced subsequent mortgage default probabilities by around 40 percent. A crucial institutional feature of HARP was that eligibility was restricted to loans securitized by Fannie Mae or Freddie Mac (the government-sponsored enterprises, or GSEs) on or before May 31, 2009. This clearly defined eligibility cutoff provides a natural experiment, enabling us to construct plausible counterfactuals to identify the causal impact of refinancing on political participation.

Specifically, lenders typically sell conforming loans to the GSEs as part of the mortgage securitization process. Although borrowers actively choose when to originate their mortgages, the timing of any sale to the GSEs – how quickly the loan exits the lender's balance sheet after the loan is originated – is outside the borrower's control and largely unknown to them. Consequently, for mortgages originated in the months preceding the May 31, 2009 cutoff, assignment to HARP eligibility can be considered quasi-random from the borrower's perspective. Leveraging this quasi-random assignment, we use HARP eligibility as an instrument for refinancing decisions between 2009 and 2012. Under this identification assumption, the resulting intention-to-treat estimates allow us to measure the local average treatment effect of refinancing on voter participation.

5.2 Empirical Strategy and Sample Construction

Our empirical strategy closely follows the approach in Abel and Fuster (2021). We focus specifically on fixed-rate purchase mortgages originated during the first six months of 2009. Given the importance of accurate measurement of borrowers' loan-to-value ratios (LTV) for our identification strategy, we restrict our analysis to purchase mortgages, for which the LTV at origination is precisely defined. This ensures clarity about borrowers' initial equity positions and thus their likely eligibility for refinancing through HARP. Specifically, we restrict our sample to loans with initial loan-to-value (LTV) ratios of at least 75%.[21] This sample restriction leaves us with those borrowers most likely targeted by HARP: Homeowners with significant leverage at origination who subsequently faced challenges refinancing through conventional means. To isolate economically meaningful rate refinancing activity, we further exclude households whose refinanced loan amounts deviated substantially (more than $5,000 above or $100,000 below) from their original mortgage. We then match these mortgages to voter registration records (L2 voter data), restricting attention to borrowers who were registered and eligible to vote in both the 2008 and 2012 elections and who remained in their homes until at least the 2012 election.[22]

A crucial aspect of our identification strategy exploits the quasi-random assignment of HARP eligibility based on the timing of loan securitization. To operationalize this, we merge anonymized CoreLogic loan-level data with investor-type information from ICE/McDash, identifying whether each mortgage was held by a government-sponsored enterprise (GSE) at two critical dates: May 31, 2009 – the HARP eligibility cutoff – and December 31, 2009. We restrict our final analytic sample to mortgages confirmed as GSE-held at the end of 2009, ensuring that HARP eligibility is determined solely by timing rather than unobservable loan or borrower characteristics. Thus, our instrumental variable (IV) for refinancing is a binary indicator equal to one for loans already held by a GSE on or before May 31, 2009. Our final analytic dataset comprises 56,550 borrowers who meet these criteria. This sample includes the same borrower and mortgage information as reported in our main analyses (Table 1), but now supplemented by an additional variable crucial for our HARP analysis:

[21] We also drop loans with purchase LTVs over 105% which are, as best we can tell, mostly data entry errors.

[22] Since voters often keep the same registration number if they move locally, it is possible to observe participation decisions at prior and future addresses. Our HARP sample, which consists of 2009 purchasers with non-missing 2008 participation is therefore a sample of local movers.

an indicator of HARP eligibility.[23]

The validity of our empirical strategy rests on the assumption that the timing of loan sales to government-sponsored enterprises (GSEs), and thus HARP eligibility, was effectively random with respect to borrower characteristics and voter participation. As discussed previously, this assumption is supported by prior literature, which shows no systematic bunching of loan originations or securitizations around the May 31, 2009 eligibility cutoff, nor evidence of strategic behavior by borrowers anticipating the program. Nonetheless, we explicitly test this assumption by examining balance on observable borrower characteristics.

Specifically, we conduct a series of balance tests by regressing key borrower and loan attributes on a dummy variable for HARP eligibility, using the following specification:

$$Y_i = \Psi \times \text{HARP Eligible}_i + \text{Controls}_i \times \Phi + \text{Party-by-County}_{p,c} \tag{2}$$

where Y_i is an attribute or characteristic of interest.

TABLE 7 HERE

Table 7 reports these balance results for our final sample of 56,550 borrowers. Crucially, column (1) shows that the likelihood of participation in the 2008 election is completely unrelated to HARP eligibility. Given that voter participation tends to exhibit substantial persistence, the absence of differences in 2008 voting behavior is particularly reassuring. We continue to find no significant differences between eligible and ineligible borrowers across a number of observable characteristics including borrower demographics, income, home value, loan amount, and property size. One exception is in the age of the property, with HARP eligible properties being roughly eight months older, on average, than HARP ineligible properties. We consider this difference economically insignificant and thus unlikely to explain the main results we document in the next subsection. Taken together, the results of these balance tests strongly support our identifying assumption that HARP eligibility, driven by the timing of loan sales to GSEs, was effectively orthogonal to borrower characteristics relevant for political engagement.

[23] As before, we visualize the creation of this dataset in Appendix B.

5.3 HARP Results

5.3.1 Full HARP Sample

Table 8 presents our primary results exploiting quasi-random variation in refinancing opportunities created by HARP eligibility.

TABLE 8 HERE

In this smaller subsample, the OLS results (column 1) continue to suggest that refinancing is positively associated with voting in the 2012 election. And consistent with the existing literature on HARP eligibility and refinancing, the first-stage estimates in column 2 reveal a strong positive effect of HARP eligibility on refinancing. HARP eligibility increases refinancing probability by roughly 13 percentage points, and the first-stage relationship is highly statistically significant (F-statistic = 203.42), indicating a robust instrument. The IV estimate, presented in column (3), directly shows that refinancing through HARP eligibility had an economically and statistically large effect on political participation. The IV estimate is roughly four times larger than the corresponding OLS estimate in column (1), and suggests that HARP refinancing increased participation by 4.4 percentage points among those that complied with this offer of refinancing.

The larger IV estimate in column (3) relative to the OLS result in column (1) is economically intuitive. OLS measures the average correlation across all borrowers who refinanced, whereas the IV estimate in column 3 identifies the local average treatment effect specifically among "compliers" – borrowers whose refinancing decisions depended critically on HARP eligibility. These borrowers were the most liquidity constrained, and thus the substantial positive IV estimate is consistent with the resource theory of political participation, suggesting that the causal impact of refinancing on voting is strongest among borrowers for whom refinancing most meaningfully relaxed financial constraints.

The reduced form, or intent-to-treat, estimate in column (4) further corroborates our approach. Specifically, we find that voters whose mortgages were HARP eligible were about 0.6 percentage points more likely to vote than otherwise similar voters whose mortgages just missed the HARP eligibility cutoff. Because HARP eligibility is a quasi-random highly targeted policy, and only affects refinancing frictions among eligible households, column (4) suggests that the debt savings from refinancing during this post-crisis period may have causally affected voter participation in the 2012 election.

6 Conclusion

This paper has shown that voters who refinanced their mortgages during the 2009-2012 period of relatively low interest rates were significantly more likely to vote in the 2012 Presidential election than otherwise similar borrowers who did not refinance. The treatment effects are larger for poorer voters and those who obtained the largest debt service reductions. This was a period marked by well-documented financial frictions that impeded mortgage refinancing and government programs like HARP, which (exogenously) relaxed some of these refinancing frictions and allowed the most constrained homeowners to refinance, seem to have engendered large increases in political participation.

Theoretically, there are arguments to be made that improved financial health decreases turnout, has no effect on turnout, or increases turnout. And the new causal evidence we bring to bear on this important question – evidence that points directly to a positive relationship between improved financial health and voter turnout – is our primary contribution. Ostensibly, we view our results as consistent with those theories that emphasize resource constraints as a key factor in political participation. These models suggest a number of potential channels through which improvements to voters' financial health might cause their participation to increase.

First, new access to greater resources allows people to overcome the cognitive costs of civic engagement. This argument observes that improved economic well being can provide voters with more time and mental resources to understand candidates' stances on salient issues and engage more with local civic groups, such as chapters of political parties, labor unions or churches, thereby increasing political participation (Mani et al., 2013; Putnam, 2000; Schaub, 2021; Verba et al., 1995; Wolfinger and Rosenstone, 1980). Another related possibility is that voters who become better off financially can more easily afford the explicitly pecuniary costs of voting. These costs include taking time off from work, which can be expensive for hourly workers, or paying for child care in order to queue at polling places or attend political rallies and civic meetings (Brady and McNulty, 2011; Pettigrew, 2021; Schafer and Holbein, 2020). More distinct from access to resources and relaxations to cognitive or pecuniary resource constraints is a third possibility related to internal political efficacy. Internal political efficacy refers to an individual's self-confidence in their ability to participate in and influence the political process. In our setting, this channel is in play if improved economic circumstances can foster increased civic pride and feelings of internal efficacy that in turn lead to increased political participation and civic engagement (Beaumont, 2011; Harder and Krosnick, 2008; Mintz et al., 2021).

There are thus many channels through which the additional resources obtained from refinancing can influence the voting decision. And while our research designs produce well-identified evidence consistent with a positive relationship between financial health and voter turnout, we leave it to future research to uncover the relative importance of each of the various mechanisms through which a relaxation of liquidity and balance sheet constraints might influence political participation.

Finally, our work speaks to the large body of research on political economy. If people without sufficient resources are less likely to be engaged in the process, then, as the argument goes, the average politician would be less likely to represent the interests of people for whom civic engagement is too costly (Avery, 2015; Franko et al., 2016). The implications for inequality of this state affairs are clear (Acemoglu et al., 2015; Bartels, 2016). We leave it to future research to study the effects of participation on the responsiveness of politicians.

References

ABEL, J. AND FUSTER, A. 2021. How do mortgage refinances affect debt, default, and spending? evidence from HARP. *American Economic Journal: Macroeconomics* 13:254–291.

ACEMOGLU, D., NAIDU, S., RESTREPO, P., AND ROBINSON, J. A. 2015. Democracy, redistribution, and inequality, pp. 1885–1966. *In* Handbook of income distribution, volume 2. Elsevier.

AGARWAL, S., AMROMIN, G., CHOMSISENGPHET, S., LANDVOIGT, T., PISKORSKI, T., SERU, A., AND YAO, V. 2023. Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinance program. *The Review of Economic Studies* 90:499–537.

AKEE, R., COPELAND, W., HOLBEIN, J. B., AND SIMEONOVA, E. 2020. Human capital and voting behavior across generations: evidence from an income intervention. *American Political Science Review* 114:609–616.

ANDERSEN, S., CAMPBELL, J. Y., NIELSEN, K. M., AND RAMADORAI, T. 2020. Sources of inaction in household finance: Evidence from the danish mortgage market. *American Economic Review* 110:3184–3230.

ARCENEAUX, K. AND NICKERSON, D. W. 2009. Who is mobilized to vote? a re-analysis of 11 field experiments. *American Journal of Political Science* 53:1–16.

AVERY, J. M. 2015. Does who votes matter? income bias in voter turnout and economic inequality in the american states from 1980 to 2010. *Political Behavior* 37:955–976.

BARTELS, L. M. 2016. Unequal democracy: The political economy of the new gilded age.

BAYER, P., CASEY, M. D., MCCARTNEY, W. B., ORELLANA-LI, J., AND ZHANG, C. S. 2022. Distinguishing causes of neighborhood racial change: A nearest neighbor design. Technical report, National Bureau of Economic Research.

BEAUMONT, E. 2011. Promoting political agency, addressing political inequality: A multilevel model of internal political efficacy. *The Journal of politics* 73:216–231.

BLAIS, A. 2000. To vote or not to vote?: The merits and limits of rational choice theory. University of Pittsburgh Pre.

BRADY, H. E. AND MCNULTY, J. E. 2011. Turning out to vote: The costs of finding and getting to the polling place. *American Political Science Review* 105:115–134.

BRADY, H. E., VERBA, S., AND SCHLOZMAN, K. L. 1995. Beyond ses: A resource model of political participation. *American political science review* 89:271–294.

BRÄNNLUND, A., CESARINI, D., LINDGREN, K.-O., LINDQVIST, E., OSKARSSON, S., AND ÖSTLING, R. 2024. Pocketbook politics: The impact of wealth on political preferences and participation. Technical report, National Bureau of Economic Research.

BROOCKMAN, D. E., RHODES, E., BARTIK, A. W., DOTSON, K., MILLER, S., KRAUSE, P. K., AND VIVALT, E. 2024. The causal effects of income on political attitudes and behavior: A randomized field experiment. Technical report, National Bureau of Economic Research.

BRUNNER, E., ROSS, S. L., AND WASHINGTON, E. 2011. Economics and policy preferences: causal evidence of the impact of economic conditions on support for redistribution and other ballot proposals. *Review of Economics and Statistics* 93:888–906.

BURDEN, B. C. AND WICHOWSKY, A. 2014. Economic discontent as a mobilizer: unemployment and voter turnout. *The Journal of Politics* 76:887–898.

CALOMIRIS, C. W. AND HABER, S. 2014. Fragile by design: The political origins of banking crises and scarce credit. Princeton University Press.

CAMPBELL, J. Y. 2006. Household finance. *The journal of finance* 61:1553–1604.

CANCELA, J. AND GEYS, B. 2016. Explaining voter turnout: A meta-analysis of national and subnational elections. *Electoral studies* 42:264–275.

CEBULA, R. J. 2017. Unemployment and voter turnout revisited: A brief note. *Electoral Studies* 48:149–152.

DESILVER, D. 2022. Turnout in us has soared in recent elections but by some measures still trails that of many other countries. *Pew Research Center* 1.

DI MAGGIO, M., KERMANI, A., KEYS, B. J., PISKORSKI, T., RAMCHARAN, R., SERU, A., AND YAO, V. 2017. Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging. *American Economic Review* 107:3550–88.

DOWNS, A. 1957. An economic theory of democracy. *Harper and Row* 28.

EHRLICH, G. AND PERRY, J. 2015. Do large-scale refinancing programs reduce mortgage defaults? evidence from a regression discontinuity design. *Evidence from a regression discontinuity design (October 22, 2015)* .

EMMENEGGER, P., MARX, P., AND SCHRAFF, D. 2015. Labour market disadvantage, political orientations and voting: how adverse labour market experiences translate into electoral behaviour. *Socio-Economic Review* 13:189–213.

ENOS, R. D., FOWLER, A., AND VAVRECK, L. 2014. Increasing inequality: The effect of gotv mobilization on the composition of the electorate. *The Journal of Politics* 76:273–288.

EVANS, R. J. 2005. The coming of the Third Reich, volume 1. Penguin.

FIELDHOUSE, A. J., MERTENS, K., AND RAVN, M. O. 2018. The macroeconomic effects of government asset purchases: Evidence from postwar us housing credit policy. *The Quarterly Journal of Economics* 133:1503–1560.

FRANKO, W. W., KELLY, N. J., AND WITKO, C. 2016. Class bias in voter turnout, representation, and income inequality. *Perspectives on Politics* 14:351–368.

FUNKE, M., SCHULARICK, M., AND TREBESCH, C. 2016. Going to extremes: Politics after financial crises, 1870–2014. *European Economic Review* 88:227–260.

FUSTER, A. AND WILLEN, P. S. 2017. Payment size, negative equity, and mortgage default. *American Economic Journal: Economic Policy* 9:167–91.

GANONG, P. AND NOEL, P. 2020. Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession. *American Economic Review* 110:3100–3138.

GERBER, A. S., GREEN, D. P., AND SHACHAR, R. 2003. Voting may be habit-forming: evidence from a randomized field experiment. *American journal of political science* 47:540–550.

GEYS, B. AND SØRENSEN, R. J. 2025. The causal effect of affluence on voter turnout: New evidence from lottery winnings. *British Journal of Political Science* 55:e7.

GILENS, M. 2012. Affluence and influence: Economic inequality and political power in America. Princeton University Press.

GREEN, R. K. AND WACHTER, S. M. 2005. The american mortgage in historical and international context. *Journal of Economic Perspectives* 19:93–114.

GYÖNGYÖSI, G. AND VERNER, E. 2022. Financial crisis, creditor-debtor conflict, and populism. *The Journal of Finance* 77:2471–2523.

HALL, A. B. AND YODER, J. 2022. Does homeownership influence political behavior? evidence from administrative data. *The Journal of Politics* 84:351–366.

HARDER, J. AND KROSNICK, J. A. 2008. Why do people vote? a psychological analysis of the causes of voter turnout. *Journal of Social Issues* 64:525–549.

HIRVONEN, S., SCHAFER, J., AND TUKIAINEN, J. 2024. Policy feedback and voter turnout: Evidence from the finnish basic income experiment. *American Journal of Political Science* .

KARAMON, K., MCMANUS, D., AND ZHU, J. 2017. Refinance and mortgage default: A regression discontinuity analysis of harp's impact on default rates. *The Journal of Real Estate Finance and Economics* 55:457–475.

KEYS, B. J., POPE, D. G., AND POPE, J. C. 2016. Failure to refinance. *Journal of Financial Economics* 122:482–499.

LEIGHLEY, J. E. AND NAGLER, J. 2013. Who votes now?: Demographics, issues, inequality, and turnout in the United States. Princeton University Press.

LIJPHART, A. 1997. Unequal participation: Democracy's unresolved dilemma presidential address, american political science association, 1996. *American political science review* 91:1–14.

LOEFFLER, H. 2023. Does a universal basic income affect voter turnout? evidence from alaska. *Political Science Research and Methods* 11:521–536.

MANI, A., MULLAINATHAN, S., SHAFIR, E., AND ZHAO, J. 2013. Poverty impedes cognitive function. *science* 341:976–980.

MARGALIT, Y. 2019. Political responses to economic shocks. *Annual Review of Political Science* 22:277–295.

MCCARTNEY, W. B. 2021. Does household finance affect the political process? evidence from voter turnout during a housing crisis. *The Review of Financial Studies* 34:949–984.

MIAN, A., SUFI, A., AND TREBBI, F. 2010. The political economy of the us mortgage default crisis. *American Economic Review* 100:1967–1998.

MINTZ, A., VALENTINO, N. A., AND WAYNE, C. 2021. Beyond rationality: Behavioral political science in the 21st century. Cambridge University Press.

PETTIGREW, S. 2021. The downstream consequences of long waits: How lines at the precinct depress future turnout. *Electoral studies* 71:102188.

PLUTZER, E. 2002. Becoming a habitual voter: Inertia, resources, and growth in young adulthood. *American political science review* 96:41–56.

POP-ELECHES, C., POP-ELECHES, G., ET AL. 2012. Targeted government spending and political preferences. *Quarterly Journal of Political Science* 7:285–320.

PUTNAM, R. D. 2000. Bowling alone: The collapse and revival of american community. *Simon Schuster* .

RAJAN, R. G. 2011. Fault lines: How hidden fractures still threaten the world economy. princeton University press.

RAJAN, R. G. AND RAMCHARAN, R. 2011. Land and credit: A study of the political economy of banking in the united states in the early 20th century. *The journal of finance* 66:1895–1931.

ROSENSTONE, S. J. 1982. Economic adversity and voter turnout. *American Journal of Political Science* pp. 25–46.

SCHAFER, J. AND HOLBEIN, J. B. 2020. When time is of the essence: A natural experiment on how time constraints influence elections. *The Journal of Politics* 82:418–432.

SCHAUB, M. 2021. Acute financial hardship and voter turnout: Theory and evidence from the sequence of bank working days. *American Political Science Review* 115:1258–1274.

SCHLOZMAN, K., SIDNEY, V., AND HENRY, B. 2018. Unequal and Unrepresented: Political Inequality and the People's Voice in the New Gilded Age. Princeton University Press.

SCHLOZMAN, K. L., SIDNEY, V., AND HENRY, B. 2012. The unheavenly chorus: Unequal political voice and the broken promise of american democracy.

SMETS, K. AND VAN HAM, C. 2013. The embarrassment of riches? a meta-analysis of individual-level research on voter turnout. *Electoral studies* 32:344–359.

STIGLER, G. J. 2021. The theory of economic regulation, pp. 67–81. *In* The political economy: Readings in the politics and economics of American public policy. Routledge.

VERBA, S. AND NIE, N. H. 1987. Participation in America: Political democracy and social equality. University of Chicago Press.

VERBA, S., SCHLOZMAN, K. L., AND BRADY, H. E. 1995. Voice and equality: Civic voluntarism in american politics. *Harvard UP* .

WOLFINGER, R. E. AND ROSENSTONE, S. J. 1980. Who votes? Yale University Press.

ZHANG, C. 2019. A shortage of short sales: Explaining the underutilization of a foreclosure alternative.

Figure 1: Treatment Effects vs Prevailing Rates

To create this figure, we estimate our main model as in Table 2 except instead of comparing those who refinanced to those who did not, we compare those who did not refinance to those who refinanced in the first month of 2009, to those who refinanced in the second month of 2009, and so on up to those who refinanced in the tenth month of 2012. The omitted group is those who did not refinance. We plot the effect of refinancing each month relative to not refinancing with a solid line. We also plot the prevailing rate each month with a dashed line. Data Sources: L2, CoreLogic, Freddie Mac, and FHFA.

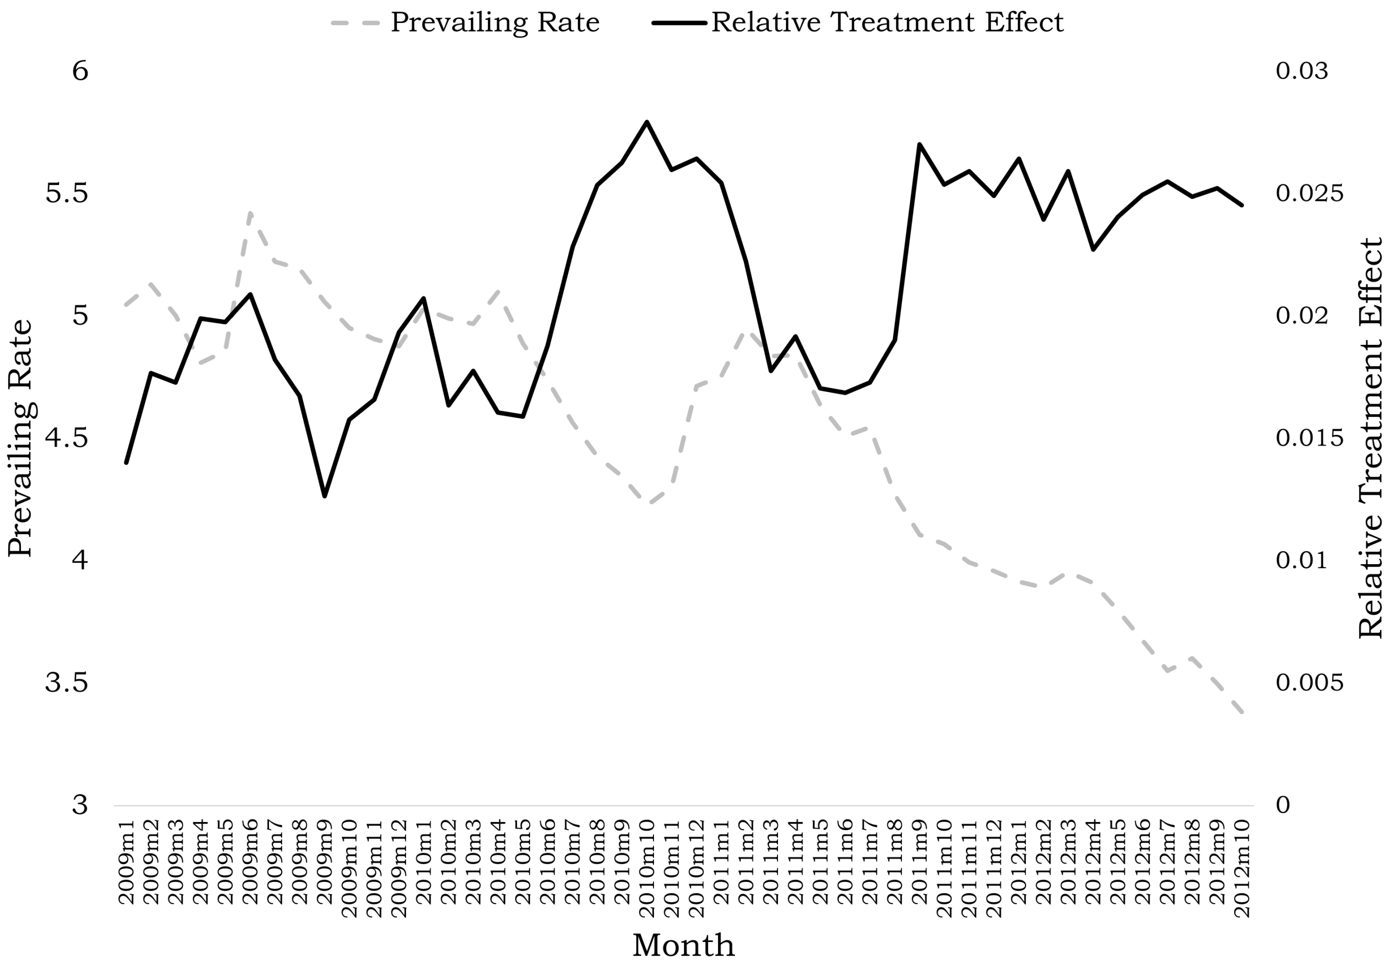

Table 1: Summary Statistics

	Mean	25$^{\text{th}}$ Pctile	75$^{\text{th}}$ Pctile	Count
Participation Rates				
2012 General Election (=1)	0.88	1.00	1.00	10,831,871
2008 General Election (=1)	0.89	1.00	1.00	10,831,871
Refinancing Activity				
Refinanced in 2009-2012 (=1)	0.45	0.00	1.00	10,831,871
Rate Drop (Ppts)	1.61	1.17	2.02	4,793,455
Voter Demographics				
Male (=1)	0.47	0.00	1.00	10,831,871
Birth Year	1961	1952	1970	10,831,871
Registration Year	1997	1992	2004	10,831,871
Democrat (=1)	0.36	0.00	1.00	10,831,871
Republican (=1)	0.37	0.00	1.00	10,831,871
Unaffiliated (=1)	0.27	0.00	1.00	10,831,871
Income at Origination ($s)	$94,383	$54,000	$110,000	6,484,835
Property Characteristics				
Year Built	1975	1959	1997	10,831,871
Square Feet	1,985	1,375	2,387	10,831,871
Outstanding Mortgage Characteristics				
Refi (=1)	0.64	0.00	1.00	10,831,871
Mortgage Amount ($s)	$208,276	$114,500	$265,500	10,831,871
Conventional (=1)	0.89	1.00	1.00	10,831,871
Under Conforming Loan Limit (=1)	0.95	1.00	1.00	10,831,871
Originated in 2005 (=1)	0.24	0.00	0.00	10,831,871
Originated in 2006 (=1)	0.22	0.00	0.00	10,831,871
Originated in 2007 (=1)	0.26	0.00	1.00	10,831,871
Originated in 2008 (=1)	0.28	0.00	1.00	10,831,871

Notes: This table describes our main sample which includes all individuals registered to vote in the 2012 election who (i) were also registered to vote in 2008, (ii) originated mortgages between 2005 and 2008, and (iii) had not refinanced those mortgages at any time before January 1, 2009. The sample and its construction is described in detail in Section 2.2. Participation data come from L2. Refinancing activity comes from CoreLogic. Rate drop is measured using data from Freddie Mac Data and is only defined for borrowers who refinanced. Sex, birth year, registration year, and party affiliation come from L2. Income at origination comes from HMDA. Property characteristics, loan purpose, loan type, loan amount, and year of origination are from CoreLogic. Author's calculations using data from FHFA defines loans under their county's conforming loan limit. Data Sources: L2, CoreLogic, Freddie Mac, HMDA, and FHFA.

Table 2: The Effect of Refinancing on Voting

Dependent Variable:	Voted in the 2012 General Election (=1)			
	(1)	(2)	(3)	(4)
Refinanced in the Post Period	0.0377***	0.0287***	0.0319***	0.0234***
	(0.00138)	(0.000892)	(0.00103)	(0.000718)
Prior Participation				
Voted in the 2008 General Election		0.373***		0.366***
		(0.00810)		(0.00792)
Demographics				
Male			-0.0149***	-0.00708***
			(0.000563)	(0.000361)
Birth Year: 1942 or Earlier (omitted)				
Birth Year: 1943 - 1958			0.00635***	0.00858***
			(0.00165)	(0.00146)
Birth Year: 1959 - 1974			-0.0162***	-0.0103***
			(0.00264)	(0.00215)
Birth Year: 1974 - 1990			-0.0377***	-0.0247***
			(0.00411)	(0.00312)
Registration Year: 1983 or Earlier (omitted)				
Registration Year: 1984 - 1991			-0.0145***	-0.00912***
			(0.00132)	(0.000977)
Registration Year: 1992 - 1999			-0.0361***	-0.0228***
			(0.00208)	(0.00123)
Registration Year: 2000 - 2008			-0.0644***	-0.0419***
			(0.00214)	(0.00124)
Property Characteristics				
Year Built: 1959 or Earlier (omitted)				
Year Built: 1960 - 1979			-0.00135*	-0.000441
			(0.000768)	(0.000609)
Year Built: 1980 - 1999			0.0000814	0.000293
			(0.000984)	(0.000802)
Year Built: 2000 or Later			0.00122	0.00220**
			(0.00121)	(0.000910)
Square Feet: 1249 or Less (omitted)				
Square Feet: 1250 - 1999			0.0221***	0.0148***
			(0.000892)	(0.000557)
Square Feet: 2000 - 2999			0.0361***	0.0241***
			(0.00148)	(0.000896)
Square Feet: 3000 or more			0.0382***	0.0265***
			(0.00215)	(0.00135)

table continued on next page...

...table continued from previous page

Outstanding Mortgage Characteristics				
Refi			-0.0219***	-0.0229***
			(0.00192)	(0.00118)
Log Mortgage Amount			0.0167***	0.0113***
			(0.00142)	(0.000871)
Conventional			0.0197***	0.0132***
			(0.00147)	(0.00100)
Under Conforming Loan Limit			-0.000296	-0.00107
			(0.00166)	(0.00117)
Fixed Effects				
Party × County	X	X	X	X
Quarter of Origination			X	X
Counts				
N	10,831,871	10,831,871	10,831,871	10,831,871
Sample Means				
Voted in the 2012 General Election	0.88	0.88	0.88	0.88
Refinanced in the Post Period	0.45	0.45	0.45	0.45

Notes: This table estimates the effects of refinancing between 2009 and 2012 on voting in the 2012 general election. The sample is described in detail in Table 1. Data Sources: L2, CoreLogic, and FHFA. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 3: Heterogeneity with Respect to Political Affiliation

Dependent Variable:	Voted in the 2012 General Election (=1)			
Sample Splits Based On:	Voter Party Affiliation			
Sample:	Republican	Democrat	Unaffiliated	Full
	(1)	(2)	(3)	(3)
Refinanced in 2009-2012	0.0212***	0.0192***	0.0318***	0.0180***
	(0.000928)	(0.000814)	(0.00113)	(0.00125)
Refinanced in 2009-2012 × Republican (omitted)				
Refinanced in 2009-2012 × Democrat				0.000738
				(0.00127)
Refinanced in 2009-2012 × Unaffiliated				0.0196***
				(0.00218)
Republican (omitted)				
Democrat				-0.0157***
				(0.00221)
Unaffiliated				-0.0775***
				(0.00489)
Controls	X	X	X	X
Fixed Effects				
Party × County	X	X	X	X
Quarter of Origination	X	X	X	X
Counts				
N	4,000,136	3,900,296	2,931,439	10,831,871
Sample Means				
Voted in the 2012 General Election	0.91	0.89	0.81	0.88

Notes: This table estimates the differential effects of refinancing between 2009 and 2012 on voting in the 2012 general election by party. The sample is described in detail in Table 1. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, a dummy for whether the outstanding mortgage is a purchase or refinance, log mortgage origination amount, a dummy for if the outstanding mortgage is a conventional loan, and a dummy for if the outstanding mortgage was under the county's conforming loan limit at the time of origination. Data Sources: L2, CoreLogic, and FHFA. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4: Heterogeneity with Respect to Outstanding Mortgage Purpose

Dependent Variable:	Voted in the 2012 General Election (=1)		
Sample Splits Based On:	Outstanding Mortgage Purpose		
Sample:	Purchase	Refinance	Full
	(1)	(2)	(3)
Refinanced in 2009-2012	0.0251***	0.0218***	0.0255***
	(0.000812)	(0.000767)	(0.000985)
Refinanced in 2009-2012 × Purchase Mortgage (omitted)			
Refinanced in 2009-2012 × Refi			-0.00332***
			(0.000898)
Outstanding Mortgage is Purchase Mortgage (omitted)			
Outstanding Mortgage is Refi			-0.0213***
			(0.00135)
Controls	X	X	X
Fixed Effects			
Party × County	X	X	X
Quarter of Origination	X	X	X
Counts			
N	3,874,717	6,957,048	10,831,871
Sample Means			
Voted in the 2012 General Election	0.87	0.88	0.88

Notes: This table estimates the differential effects of refinancing between 2009 and 2012 on voting in the 2012 general election by loan purpose. The sample is described in detail in Table 1. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, a dummy for whether the outstanding mortgage is a purchase or refinance, log mortgage origination amount, a dummy for if the outstanding mortgage is a conventional loan, and a dummy for if the outstanding mortgage was under the county's conforming loan limit at the time of origination. Data Sources: L2, CoreLogic, and FHFA. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 5: Heterogeneity with Respect to Voter Income

Dependent Variable:	Voted in the 2012 General Election (=1)		
	(1)	(2)	(3)
Refinanced in 2009-2012	0.0236***	0.0228***	0.0289***
	(0.000755)	(0.000744)	(0.00113)
Refinanced in 2009-2012 × Income: $0 - $49,999 (omitted)			
Refinanced in 2009-2012 × Income: $50,000 - $74,999			-0.00521***
			(0.000974)
Refinanced in 2009-2012 × Income: $75,000 - $124,999			-0.00845***
			(0.00107)
Refinanced in 2009-2012 × Income: $125,000 - $2,000,000			-0.00954***
			(0.00123)
Income: $0 - $49,999 (omitted)			
Income: $50,000 - $74,999		0.0134***	0.0151***
		(0.000508)	(0.000643)
Income: $75,000 - $124,999		0.0225***	0.0259***
		(0.000717)	(0.000941)
Income: $125,000 - $2,000,000		0.0272***	0.0314***
		(0.00110)	(0.00131)
Controls	X	X	X
Fixed Effects			
Party × County	X	X	X
Quarter of Origination	X	X	X
Counts			
N	6,483,547	6,483,547	6,483,547
Sample Means			
Voted in the 2012 General Election	0.88	0.88	0.88
Refinanced in 2009-2012	0.47	0.47	0.47

Notes: This table estimates the effects of refinancing between 2009 and 2012 on voting in the 2012 general election. The sample is described in detail in Table 1. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, a dummy for whether the outstanding mortgage is a purchase or refinance, log mortgage origination amount, a dummy for if the outstanding mortgage is a conventional loan, and a dummy for if the outstanding mortgage was under the county's conforming loan limit at the time of origination. Data Sources: L2, CoreLogic, HMDA, and FHFA. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 6: Heterogeneity with Respect to Voter Savings

Dependent Variable:	Voted in the 2012 General Election (=1)	
Sample:	Refinanced 0 or 1 times Between 2009 and 2012	
Subsample:	Full	Non-Missing Income
	(1)	(2)
Rate Drop		
Did Not Refinance (omitted)		
Rate Drop: 99 bpts or less	0.0172***	0.0172***
	(0.000796)	(0.000905)
Rate Drop: 100 bpts - 199 bpts	0.0213***	0.0208***
	(0.000652)	(0.000706)
Rate Drop: 200 bpts or more	0.0251***	0.0244***
	(0.000929)	(0.000935)
Income		
Income: $0 - $49,999 (omitted)		
Income: $50,000 - $74,999		0.0135***
		(0.000509)
Income: $75,000 - $124,999		0.0226***
		(0.000724)
Income: $125,000 - $2,000,000		0.0272***
		(0.00110)
Controls	X	X
Fixed Effects		
Party × County	X	X
Quarter of Origination	X	X
Counts		
N	9,872,407	5,879,967
Sample Means		
Voted in the 2012 General Election	0.87	0.87

Notes: This table estimates the effects of refinancing between 2009 and 2012 on voting in the 2012 general election. Different from our other analyses, refinance is not a dummy variable, but rather a categorical variable with four categories: did not refinance, refinanced and saved 99 basis points or less, refinanced and saved between 100 and 199 basis points, or refinanced and saved between 200 basis points or more. The sample is described in detail in Table 1 with the added restriction that borrowers who refinanced two times or more and omitted. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, a dummy for whether the outstanding mortgage is a purchase or refinance, log mortgage origination amount, a dummy for if the outstanding mortgage is a conventional loan, and a dummy for if the outstanding mortgage was under the county's conforming loan limit at the time of origination. Data Sources: L2, CoreLogic, Freddie Mac, HMDA, and FHFA. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 7: HARP Balance Test

Dependent Variable:	Voted in 2008	Log Sale Amount	Log Loan Amount	Male	Age	Years Registered	Income	Square Footage	Year Built
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
HARP Eligible	-0.003	0.001	0.000	0.001	0.182	0.001	500.166	-0.866	-0.630**
	(0.004)	(0.001)	(0.001)	(0.003)	(0.143)	(0.099)	(1453.820)	(9.616)	(0.295)
Controls	X	X	X	X	X	X	X	X	
Fixed Effects									
Party × County	X	X	X	X	X	X	X	X	X
Quarter of Origination	X	X	X	X	X	X	X	X	X
Counts									
N	56,550	56,550	56,550	56,550	56,550	56,550	45,212	56,550	56,550
Sample Means									
Dependent Variable	0.85	12.49	12.30	0.50	40.37	8.59	$105,593	2200.76	1980

Notes: This table estimates the effects of being HARP eligible on various borrower and loan characteristics. The sample consists of loans originated between January 2009 and June 2009 that were held by a GSE by the end of 2009 and is described in detail in Section 5.2. HARP eligibility equals 1 if a loan was held by a GSE by the end of May 2009, and 0 otherwise. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, four bins for income, log mortgage origination amount, and log sale amount. When the variable is used as the dependent variable, it is not included as a control and we use the raw amount, not a binned version of it. Data sources: CoreLogic, L2, HMDA, and ICE, McDash. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 8: HARP Results

Dependent Variable:	Voted in 2012	Refi'd in Post Period	Voted in 2012	Voted in 2012
	OLS	First Stage	2SLS	Reduced-form
	(1)	(2)	(3)	(4)
Refinanced in the Post Period	0.010***		0.044**	
	(0.003)		(0.020)	
HARP Eligible		0.130***		0.006**
		(0.009)		(0.003)
Controls	X	X	X	X
Fixed Effects				
Party × County	X	X	X	X
Quarter of Origination	X	X	X	X
First Stage F-Stat			203.42	
Counts				
N	56,550	56,550	56,550	56,550
Sample Means				
Dependent Variable	.90	.32	.90	.90

Notes: This table estimates the effects of refinancing between 2009 and 2012 on voting in the 2012 general election using HARP eligibility as an instrumental variable. The sample consists of loans originated between January 2009 and June 2009 that were held by a GSE by the end of 2009 and is described in detail in Section 5.2. HARP eligibility equals 1 if a loan was held by a GSE by the end of May 2009, and 0 otherwise. Column (1) presents the estimates from an OLS regression that estimates our main model, Equation 1, on this sample. Column (2) presents the results of the first stage regression of refinancing between 2009 and 2012 on HARP eligibility. Column (3) presents the IV regression estimate of voting in the 2012 general election on refinancing between 2009 and 2012 where HARP eligibility is the IV. Column (4) presents the estimates from the reduced form regression of voting in the 2012 general election on HARP eligibility. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, four bins for income, log mortgage origination amount, and log sale amount. Data sources: CoreLogic, L2, HMDA, and ICE, McDash. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Relieving Financial Distress Increases Voter Turnout: Evidence from the Mortgage Market

Haoyang Liu, W. Ben McCartney, Rodney Ramcharan, Calvin Zhang, and Xiaohan Zhang

Online Appendix

A A Difference-in-Differences Framework

In this online appendix, we leverage both the time series and cross section of our dataset and explicitly use a difference-in-differences research design. The first difference occurs along the time dimension and compares 2008 and 2012 participation decisions. The second difference compares voters who refinanced their mortgages in between those two elections and those who did not. Specifically, to estimate the impact of refinancing on voting, we use the following difference-in-difference specification:

$$\begin{aligned}\text{Voted}_{i,t} = \beta_1 \times \text{Refinanced}_i \times \text{Post}_t + \beta_2 \times \text{Post}_t + \beta_3 \times \text{Refinanced}_i \\ + \text{Controls}_i \times \Phi + \text{Party-by-County}\epsilon_{p,c} \quad \text{(A1)}\end{aligned}$$

where $\text{Voted}_{i,t}$ is a dummy indicating whether individual i voted in the general election in year $t \in \{2008, 2012\}$. Refinanced_i is a dummy indicating that an individual refinanced between January 2009 and October 2012 and Post_t is variable equal to 0 in 2008 and 1 in 2012. We include the same vector of controls as before that characterize voters, their homes, and their original mortgages. And we continue to absorb variation in participation explained by location and party affiliation by including party-by-county fixed effects. Finally, we replace the control variables and main effect of refinancing with individual-level fixed effects. Their inclusion further helps us address omitted variable bias related to latent individual factors affecting both refinancing and 2012 civic engagement decisions. The coefficient of interest, β_1, measures whether individuals that refinanced between elections were more likely to vote in 2012.

[TABLE A1 – Difference-in-Differences – HERE]

Table A1 presents the results. In column (1), we omit all control variables and fixed effects. Looking to the main effects first we see that borrowers who refinanced their mortgage at some point between 2009 and 2012 are more likely to vote, on average. This is consistent with the basic intuition that borrowers with the financial literacy to know that refinancing is a good idea and the time and resources to actually refinance are more likely to vote in elections as well. Crucially, this effect gives us a baseline and can be thought of like a pre-trend. The negative main effect on the post-period dummy tells us that participation turnout declined between 2012 and 2008 among this population.

That turnout was lower in 2012 than 2008 is a well documented fact. Finally, the interaction tells us that voters who refinanced in the years leading up to the 2012 election were especially likely to vote in 2012 compared to those who did not, relative to the baseline difference between these two groups. The estimate is a significant 0.0169 and is reassuringly similar to the 0.0234 we estimate using our alternative specification.

In column (2), we include our standard battery of control variables. And then in column (3), we replace the control variables, main effect of refinancing, and party-by-origination fixed effects with an individual voter fixed effect. Of course, since none of our controls are time varying, the estimate of the treatment effect does not vary across specifications. It is due to this limitation of the data combined with the parsimony of our previous research design that leads us to prefer the model described by Equation 1.

Table A1: The Effect of Refinancing on Voting

Dependent Variable:	Voted in the General Election (=1)		
	(1)	(2)	(3)
Interaction Effect			
Refinanced in 2009-2012 × 2012 General Election	0.0169***	0.0169***	0.0169***
	(0.00109)	(0.00109)	(0.00109)
Main Effects			
Refinanced in 2009-2012	0.0240***	0.0190***	
	(0.00133)	(0.00102)	
2012 General Election	-0.0213***	-0.0213***	-0.0213***
	(0.00178)	(0.00178)	(0.00178)
Controls		X	
Fixed Effects			
Party × County		X	
Quarter of Origination		X	
Voter			X
Counts			
N	21,665,404	21,665,404	21,665,404

Notes: This table studies the effect of refinances on participation in the 2008 and 2012 general elections through a set of differences-in-differences specifications. Column (1) is without fixed effects. Column (2) controls for party-by-county and quarter of origination fixed effects and includes a number of control variables. Controls include a dummy for participation in 2008, a dummy for if the voter is male, four birth year bins, four registration year bins, four year built bins, four bins for building square footage, a dummy for whether the outstanding mortgage is a purchase or refinance, log mortgage origination amount, a dummy for if the outstanding mortgage is a conventional loan, and a dummy for if the outstanding mortgage was under the county's conforming loan limit at the time of origination. Column (3) controls for voter fixed effects. Standard errors, adjusted for clustering at the state-by-year level, are reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

B Visualizing the Sample Creation Process

In this data appendix we illustrate and describe in greater detail how our two main samples are created. Complete data cleaning and sample creation code as allowed by the contracts with our data vendors will also be supplied to interest readers.

B.1 Full Sample

In Figure B1, we visually represent how the full sample is constructed. We use two main data sources. The first, shown in Panel A, is the deeds and mortgage data from CoreLogic. Our goal with this data is to create a sample of all mortgages outstanding as of January 1, 2009 that were "at risk" of being refinanced. To that end, we start with the sample of all fixed-rate mortgages originated between 2005 and 2008. Among that set, we then keep each household's most recently originated one. That is, if a household took out a purchase mortgage in 2005 and then refinanced that mortgage in 2007, we keep only the refinance in our sample. We then restrict this sample to just those mortgages tied to properties that were not sold before the 2012 election. This yields a sample of more than 14 million mortgages. Our data set thus focuses on "never-movers," at least over the time series of interest – 2005-2008 mortgage origination to 2012 general election. Focusing on this group allows for the cleanest measure of treatment status and estimate of said treatment.

Panel B describes our voter data. We use the snapshot of registered voters as of early 2014. This ensures that 2012 participation, which sometimes takes states a long time to record in the data, is fully documented in the voter files. Since voter files remain fairly stale except between major elections when voters go to update their registrations, using early 2014 snapshots does a good job describing the population that was registered to vote in the 2012 election. Our key sample restriction here is to drop voters whose 2008 participation decision we cannot observe. Since we eventually merge this data with the sample of homeowners who have lived in their homes since at least 2008, this effectively drops just those voters who did not register until after the 2008 election. The final sample of voters includes more than 123 million voters.

Panel C presents the match rates. We begin with the sample of voters and then ask if somebody with their same name and same address exists in our sample of mortgages. We allow some flexibility here to match potential spouses. The final data set that we use for our full sample analysis includes just under 11 million voters attached to just under 7 million unique mortgages.

Figure B1: Full Sample Creation Schema

This schema describes overviews the sample creation procedure used to build the full sample, discussed in the main text in Section 2.2.

A. Input – CoreLogic Real Estate Data

Description & Sample Restrictions
- Fixed-rate mortgages
- Originated between Jan 1, 2005 and Dec 31, 2008
- Property owner had not sold home prior to Oct 31, 2012

Key Variables
- Owner first name, Owner last name
- Street address of property
- Loan purpose, Date of origination, Amount
- House size, House Age

Sample Size
- 14,207,675 distinct mortgages

B. Input – L2 Voter Data

Description & Sample Restrictions
- Registered voters as of the 2012 November election
- Non-missing 2008 and 2012 participation

Key Variables
- Voter first name, Voter last name
- Street address of primary residence
- Party affiliation, Date of registration
- Date of birth, Sex
- Participation in 2008 and 2012 elections

Sample Size
- 123,071,909 distinct registered voters

C. Main Sample

Merge Steps
1. Exact match on street address
2. Same first and last name *or* likely spouse

Sample Size
- 6,840,781 distinct mortgages
- 10,885,440 distinct registered voters

B.2 HARP Sample

In Figure B2, we visually represent how the HARP sample is constructed. Panels A, B, and C are very similar to Panels A, B, and C in Figure B1. The exceptions come in the sample restrictions for the mortgages. We closely follow Abel and Fuster (2021) in our construction of mortgages eigible for the HARP experiment. Specifically, we limit to mortgages originated in the first half of 2009. Since the HARP policy was designed to help highly levered and even underwater borrowers refinance, we limit our sample to borrowers with high LTVs. We choose to use the LTV at the time of purchase since measuring current LTVs is notoriously difficult, especially during this period of rapidly falling house prices. In some sense, whether a borrower had a "true" current LTV of 93% or 98% or 107% is not knowable; three different appraisers could well come up with three different current home values. We therefore opt to include everybody who took out an economically meaningful loan, which we define as having an LTV at purchase, a time when both the loan and the house has a very well-defined value, of at least 75%.

The other major requirement to be in the HARP experiment is that the loan *could have* been purchased by a GSE. Since we cannot know this using just information available in the deeds data, we need to merge in information from ICE/McDash. In order to be in compliance with our contracts and protect privacy, the bridge between CoreLogic and McDash was made behind a firewall, independently of this project. We are then given a file with three variables: the CoreLogic loan ID, whether the loan was owned by a GSE as of December 31, 2009, and whether the loan was owned by a GSE as of May 31, 2009. Of course, we know that loans owned by a GSE as of Dec 31, 2009 *could have* been purchased by a GSE since we see that they were. When we then merge in these variables to the CoreLogic-L2 matched sample we are left with 163,000 distinct mmortgages that match to ICE/McDash.

We then make the conservative approach to limit our sample to just these loans. This likely omits some loans that could have been purchased by a GSE and for whatever reason were not, but it has the key advantage of ensuring that, within our final sample, the loans differ only in when they were purchased by a GSE. This final sample includes 43,500 mortgages that were owned by a GSE as of the end of 2009. On these mortgages are just over 60,000 distinct voters.

Our final sample is much smaller than the 219,407 in Abel and Fuster (2021). There are three main reasons for this difference. First, LTV is a crucial variable in this context, our decision to focus on purchase loans drop many mortgages from the sample. 27.6%, or approximately 60,500 of the loans in Abel and Fuster (2021), are purchase mortgages. Second, to be included in our sample requires, as before, that the loan in CoreLogic match to at least one voter in the L2 data. And finally, privacy considerations mean we cannot use Equifax CRISM data. Instead, we use a bridge between CoreLogic and ICE/McDash. The coarsness of the matching variables means that many loans in our main dataset fail to match to ICE, McDash (of the 294,168 mortgages that meet the first set of sample requirements, 162,964 match to McDash).

Figure B2: HARP Sample Creation Schema

This schema describes overviews the sample creation procedure used to build the HARP sample, discussed in the main text in Section 5.3.1.