The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXXV, NO. 1 • FEBRUARY 2020

Stimulating Housing Markets

DAVID BERGER, NICHOLAS TURNER, and ERIC ZWICK*

ABSTRACT

We study temporary fiscal stimulus designed to support distressed housing markets by inducing demand from buyers in the private market. Using difference-in-differences and regression kink research designs, we find that the First-Time Homebuyer Credit increased home sales by 490,000 (9.8%), median home prices by \$2,400 (1.1%) per standard deviation increase in program exposure, and the transition rate into homeownership by 53%. The policy response did not reverse immediately. Instead, demand comes from several years in the future: induced buyers were three years younger in 2009 than typical first-time buyers. The program's market-stabilizing benefits likely exceeded its direct stimulus effects.

IN THE AFTERMATH OF THE GREAT RECESSION, the U.S. housing market suffered extraordinary distress (Figure 1). As house price growth slowed, a shortage of prospective homebuyers caused vacancies to rise and housing inventory to double from 2004 to 2006 and remain at historic levels through 2008. The boom coincided with a rapid, widespread rise in household debt secured by real estate (Mian and Sufi (2015)). When house prices began to fall, defaults, foreclosures, and further downward pressure on prices ensued (Campbell, Giglio, and Pathak (2011), Mian, Sufi, and Trebbi (2015), Guren and McQuade (2015)). By mid-2008, the composition of home sales had shifted dramatically, with nearly 40% classified as distressed or foreclosure sales.

*David Berger is at Duke University and the National Bureau of Economic Research (NBER). Nicholas Turner is at the Federal Reserve Board. Eric Zwick is at Chicago Booth and NBER.
We thank Andrew Abel, Gene Amromin, Michael Best, Jediphi Cabal, Anthony DeFusco, Paul Goldsmith-Pinkham, Adam Guren, Erik Hurst, Anil Kashyap, Amir Kermani, Ben Keys, Henrik Kleven, Pat Langetieg, Adam Looney, Janet McCubbin, Matt Notowidigdo, Christopher Palmer, Jonathan Parker, Amit Seru, Isaac Sorkin, Johannes Stroebel, Amir Sufi, Joe Vavra, Rob Vishny, Owen Zidar, and seminar and conference participants for comments, ideas, and help with data. Tianfang Cui, Prab Upadrashta, Iris Song, and Caleb Wroblewski provided excellent research assistance. The views expressed here are ours and do not necessarily reflect those of the U.S. Treasury Office of Tax Analysis, the Internal Revenue Service (IRS) Office of Research, Analysis and Statistics, or the Federal Reserve Board. We all have no relevant or material financial interests that relate to the research described in this paper. The underlying individual-level tax data were accessed while Turner worked as staff economist in the Office of Tax Analysis in the U.S. Treasury. To comply with Internal Revenue Code (IRC) 6103(j) that defines permissible uses of tax return data, Treasury civil servants had the right to review preliminary drafts. These reviews focused on protecting taxpayers from risk of disclosure and did not influence the structure or content of the paper in a material way. Zwick gratefully acknowledges financial support from the Neubauer Family Foundation, Initiative on Global Markets, and Booth School of Business at the University of Chicago.

DOI: 10.1111/jofi.12847

Panel A. Inventory (NAR)

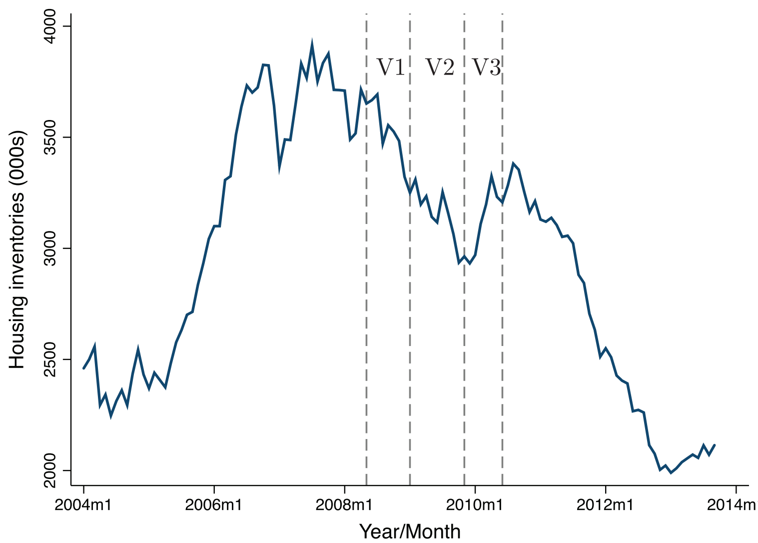

Panel B. Existing Home Sales Composition (DataQuick)

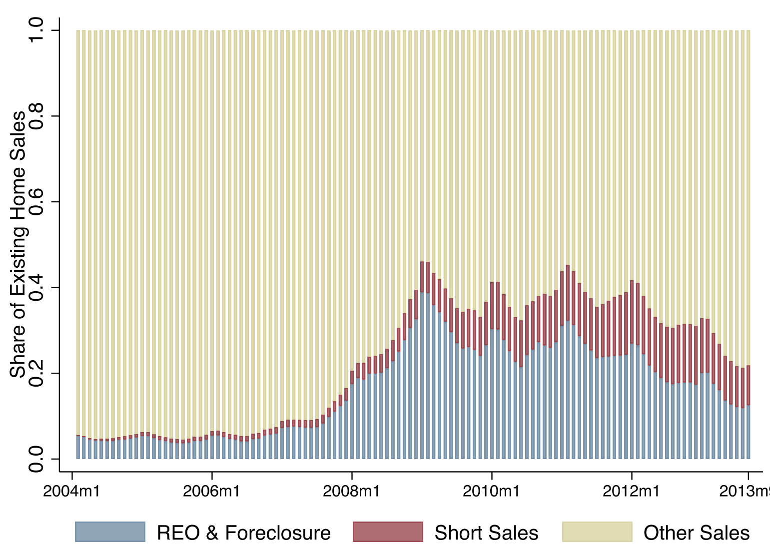

Figure 1. The state of the housing market. Panel A plots seasonally adjusted housing inventory, defined as the number of homes listed for sale, from the National Association of Realtors (NAR). The vertical markers correspond to the FTHC loan program (V1), the start of the FTHC grant program (V2), the scheduled expiration of the FTHC grant program, and the actual expiration of the FTHC grant program (V3), respectively. Panel B plots the month-by-month share of existing home sales in DataQuick in each of three categories: nondistress resales, short sales, and institution-owned sales (REO) or foreclosures. (Color figure can be viewed at wileyonlinelibrary.com)

The debt-induced overhang in the housing market prompted many policy responses, including debt renegotiation programs to repair household balance sheets, government asset purchases to support financial markets, and monetary and fiscal policy to spur demand growth.[1] Yet these policies do not directly target the problem of housing supply overhang, nor do they promote reallocation when houses are vacant or no longer held by high-value users.

This paper evaluates a complementary policy, the First-Time Homebuyer Credit (FTHC), which was a $20 billion stimulus program designed to support U.S. housing markets with a temporary tax incentive for new homebuyers between 2008 and 2010. We combine data from administrative tax records with transaction deeds data to measure program exposure and housing market outcomes for approximately 9,000 ZIP codes, which account for 69% of the U.S. population. We use difference-in-differences and regression kink research designs to estimate the effect of the policy on home sales, homeownership, and the housing market more broadly.

We present five main findings. First, the policy proved effective at spurring home sales. We estimate that the FTHC increased the number of home sales during the policy period by 490,000 units nationally. Second, the surge in home sales did not reverse immediately in the year following the policy period. Instead, demand appears to have come from several years in the future. Third, the policy induced transitions into homeownership. We estimate that receiving the FTHC increased the likelihood of being a first-time homebuyer by over 50%. Fourth, the policy response came mainly via existing home sales, implying that the direct stimulative effects of the program were small. Fifth, the health of the housing market, as reflected in house prices, improved. A back-of-the-envelope calculation suggests that the consumption response to the increase in house prices was likely larger than the policy's direct stimulus effect.

We first document the effect of the FTHC on home sales. Our difference-in-differences design compares ZIP codes at the same point in time whose exposure to the program differs. We define program exposure based on the number of potential first-time homebuyers in a ZIP code. ZIP codes with few potential first-time homebuyers serve as a "control group" because the policy does not induce many people to buy in these places. We measure exposure as the year-2000 share of people in a ZIP code who are first-time homebuyers.

The key threat to this design is the possibility that time-varying, ZIP-specific shocks are correlated with our exposure measure. We assess this threat in four ways. First, we present graphical evidence of parallel pre-policy trends, clear breaks during the policy period, and spikes at policy expiration. Second, we show that the results are robust to including city-by-time fixed effects, to using varying weighting schemes and sample definitions, and to adding explicit controls for exposure to the subprime bust and to all major contemporaneous

[1] Diamond and Rajan (2011), French et al. (2010), Shleifer and Vishny (2010a), Hanson, Kashyap, and Stein (2011), and Eberly and Krishnamurthy (2014) discuss potential policy solutions. A recent empirical literature evaluates some of the programs to address debt overhang during the Great Recession (Agarwal et al. (2017a, 2017b)).

policies. Third, we document shifts in the age distribution of first-time homebuyers that ZIP-by-time trends cannot explain. And fourth, we exploit differential program generosity based on house price levels and compare starter homes to large homes in a within-time placebo test.

Complementing our market-level, difference-in-differences research design, we estimate the effect of the FTHC at the individual level using detailed population-level administrative tax data and a sharp regression kink design (RKD). This research design exploits the income phase-out range of the FTHC. The causal effect of the FTHC on being a first-time homebuyer is the ratio of the slope change in the probability of being a first-time homebuyer and the slope change in the potential FTHC. We estimate that the FTHC increased the rate of transition into homeownership by 0.76% relative to a baseline rate of 1.43%. This estimate implies an aggregate effect of 520,000 additional purchases, matching the estimate from our difference-in-differences approach with a different identification strategy.

We next explore the role of the FTHC program as housing market stabilizer. We first examine the effect on house prices following the same empirical strategy used to analyze home sales. We find that the program increased house prices significantly. In our preferred specification, a one-standard-deviation increase in exposure to the program led the median home to appreciate by $2,400 during the policy window. We also show that aggregate repeat-sale price indices probably understate the true effect of the program because they smooth high-frequency price changes and exclude a large share of policy-relevant sales.

Last, we present suggestive evidence that the program likely accelerated the process of reallocation from low-value sellers to high-value buyers. Many transactions during the policy period involved sales by low-value homeowners, including investors, institutional sellers, financial institutions, government-sponsored entities, and builders and developers selling unsold inventory of recently built homes. Furthermore, many buyers induced by the program were constrained by down payment requirements and liquidity needs that the credit helped relax.[2]

Down payment constraints can also explain why we fail to find evidence of a sharp reversal after the policy expires: absent the policy, induced buyers must wait until they have accumulated the necessary down payment as savings. Although many policy-period buyers bought with high loan-to-value (LTV) ratios, they were not more likely to default in the subsequent three years than other cohorts of homebuyers. The fact that housing demand was being pulled from years rather than months in the future further supports the view that the program had stabilizing effects in the medium run.

Our paper contributes to the empirical literature on policy responses to distress in debt markets, especially policies motivated by the Great Recession. Relative to these papers, we focus on how policy can both stabilize prices during potential fire sales and address capital overhang by accelerating reallocation,

[2] We explore this fact and the implications for theories of intertemporal demand for durables in a follow-up paper (Berger et al. (2018)).

which is typically slow during periods of industry decline or macroeconomic weakness. Despite a large theoretical literature, relatively little empirical work evaluates policies that target overhang.[3] Our paper complements studies that estimate the effects of fiscal stimulus by analyzing an important durable goods stimulus program. Taken together, these studies demonstrate that the reversal of durable goods stimulus programs depends on which activity is targeted and who the marginal buyers are.[4]

A closely related paper by Best and Kleven (2017) studies the effect of fiscal stimulus on housing sales in the United Kingdom. The authors find similar effects on home sales that only reverse partially post-policy. While they study similar questions, the two papers employ different research designs with different strengths and weaknesses. In particular, they use a bunching research design and a difference-in-differences design based on house price cutoffs, while we use cross-market variation in program exposure, cohort analysis of the age distribution of first-time buyers, and a RKD based on income cutoffs. We note two substantive differences. First, their research designs do not permit study of the broader effects of these policies on housing market health. We document that the FTHC had powerful effects on market-level house prices and promoted reallocation of underused housing, which suggests that the policy played an important role as a housing market stabilizer. Second, the policy that we study focuses explicitly on first-time homebuyers. We can therefore identify the causal effect of a temporary tax change on transitions into homeownership—an important input for evaluating many government policies—in addition to the effect on overall transaction volume.

Our results largely align with other work on the FTHC program. Using a cross-city and cross-segment difference-in-differences strategy, Brogaard and Roshak (2011) find that quantity was not measurably affected and that prices rose by \$6K to \$11K initially but by only \$1K to \$5K after expiration. Dynan, Gayer, and Plotkin (2013) conclude that the credit had "at best, small and mostly temporary effects on housing activity," identifying small positive effects on home sales (2% higher during program months) and on prices (less than 1%) using cross-state variation in FTHC programs. They also find that grant programs have a larger effect than loan programs. In contrast, Hembre (2018) finds negligible price effects but large quantity responses—a 16% increase in purchases, or 255,000 new homeowners, during the loan program and the first

[3] A theoretical literature, going back to the classic debate between Hayek (1931) and Keynes (1936), studies how policy should respond to capital overhang following investment booms. When booms coincide with credit expansions, high-valuation potential buyers often cannot finance distressed asset purchases in the subsequent slump (Shleifer and Vishny (1992)). In this case, an overhang leads to fire sales and inefficient liquidation, creating a role for welfare-improving policy intervention (Fisher, 1933; Kiyotaki and Moore, 1997; Lorenzoni, 2008; Eggertsson and Krugman, 2012). Ramey and Shapiro (2001), Eisfeldt and Rampini (2006), and Rognlie, Shleifer, and Simsek (2018) analyze the effect of capital overhang amid frictions that slow reallocation.

[4] Adda and Cooper (2000), House and Shapiro (2008), Mian and Sufi (2012), Berger and Vavra (2015), Zwick and Mahon (2017), and Green et al. (2016) study other durable goods stimulus programs.

grant program. Overall, he estimates that the program led to 400,000 new homeowners. Relative to these papers, our approach yields somewhat stronger results, likely driven by more granular data and sharper research designs.

Section I provides background information on the FTHC program. Section II describes the data. Section III describes our main empirical strategy. Section IV presents results on home sales using cross-sectional exposure. Section V presents results on homeownership using the income phase-out. Section B presents results on house prices and reallocation. Section VII uses results from Sections IV, V, and B to estimate the aggregate direct and indirect effects of the policy. Section VIII concludes.

I. Policy Background

The FTHC was a temporary stimulus policy introduced in the United States between 2008 and 2010 with the aim of supporting weak housing markets. There were three versions of the program. The first version, enacted on July 30, 2008 as part of the Housing and Economic Recovery Act, provided an interest-free loan of up to $7,500 on qualifying home purchases made between April 9, 2008 and June 30, 2009. To be eligible for the maximum value of this version of the credit, a single (married) taxpayer needed a modified adjusted gross income (AGI) below $75,000 ($150,000) and must not have owned a principal residence during the three-year period preceding the purchase date.

The second version of the credit was enacted on February 17, 2009 as part of the American Recovery and Reinvestment Act. The policy window was extended to include purchases made through November 30, 2009. Importantly, the maximum credit was raised to $8,000 (specifically, 10% of sale price up to $8,000) and the credit now took the form of a refundable tax credit rather than an interest-free loan. The latter feature significantly increased the value of the credit to potential homebuyers.

The third version of the credit was enacted on November 7, 2009 as part of the Worker, Homeownership, and Business Assistance Act. The policy window was extended to include purchases closing before July 1, 2010.[5] The third version also raised the income limits so that eligibility began to phase out for a single (married) taxpayer with modified AGI above $125,000 ($225,000). For each version of the credit, the eligible amount phased out linearly over a $20,000 range above the income limit.

To claim the credit, tax filers needed to note the FTHC on their income tax returns (Form 1040) and attach an additional credit claim (Form 5405). Claimants also needed to document that the relevant purchase occurred during the policy window and submit evidence supporting the claim's eligibility.[6]

[5] The expanded policy also added a $6,500 Long-Time Homebuyer Credit (LTHC). To qualify for the LTHC, an individual must have owned and used the residence as his or her principal residence for five consecutive years during the eight years prior to the date of the new purchase.

[6] Such documents could include the settlement statement (typically Form HUD-1), executed retail sales contract (for mobile homes), or certificate of occupancy (for new construction).

To accelerate payment, filers could amend previously filed tax returns, for example, by amending the 2008 return for a home bought in 2009.

We focus on the second and third versions of this policy. First, these versions were considerably more generous and thus more likely to induce new purchases.[7] Second, these versions were broadly publicized at the time they were enacted and thus were more likely to induce changes in behavior than retrospective claims for past purchases. Third, unlike the first loan-based version of the credit, the second and third versions could contribute to down payments, following lender guidance by the Department of Housing and Urban Development (HUD; see Mortgagee Letter 2009-15).

Figure 2 plots time series that justify our focus on the second and third versions. Panel A plots existing home sales from the National Association of Realtors (NAR) and shows significant aggregate spikes at the end of the second and third policy windows. Panels B and C confirm these spikes within our analysis sample using data from DataQuick. Panel D plots Google search trend data for the terms "first-time home buyer" and "home buyer credit," along with vertical markers for policy events. Interest in these credits spiked at the beginning of the second version, remained elevated throughout both policy periods, and then declined after the end of the third version.

Congress passed the FTHC with the explicit purpose of inducing demand for homes at a time of unusual market weakness. As macroeconomic stimulus, the goal was to spur economic activity by inducing new home sales and the expenditures that accompany a home sale. As market stabilizer, the goal was to slow the rate of house price declines and promote the reallocation of underutilized homes to higher value buyers. In the respective words of Senators Cardin, Shelby, and Salazar, the program aimed to "help the housing market," "help get homebuilders and the housing industry back on track," and it would "help us get rid of the glut we currently have in the market."[8] The nonrandom timing of the policy motivates the cross-sectional approach we pursue to separate the effect of the program from other factors affecting housing markets at this time.

II. Data

In this section, we present an overview of our data sources, we discuss the construction of key variables used in our analysis, and we present summary statistics. Appendix I presents additional information on the data construction process, detailed variable definitions, and supplementary sample statistics.

[7] Assuming a 3% real rate of return, the interest-free loan was worth $1,400 in present value, while the later versions were worth 5.7 times as much. Table IA.I in the Internet Appendix estimates the effect of the first version and shows modest positive effects. The Internet Appendix is available in the online version of this article on *The Journal of Finance* website.

[8] See Congressional Record, Vol. 154, No. 52 (April 3, 2008) and Congressional Record, Vol. 154, No. 124 (July 26, 2008).

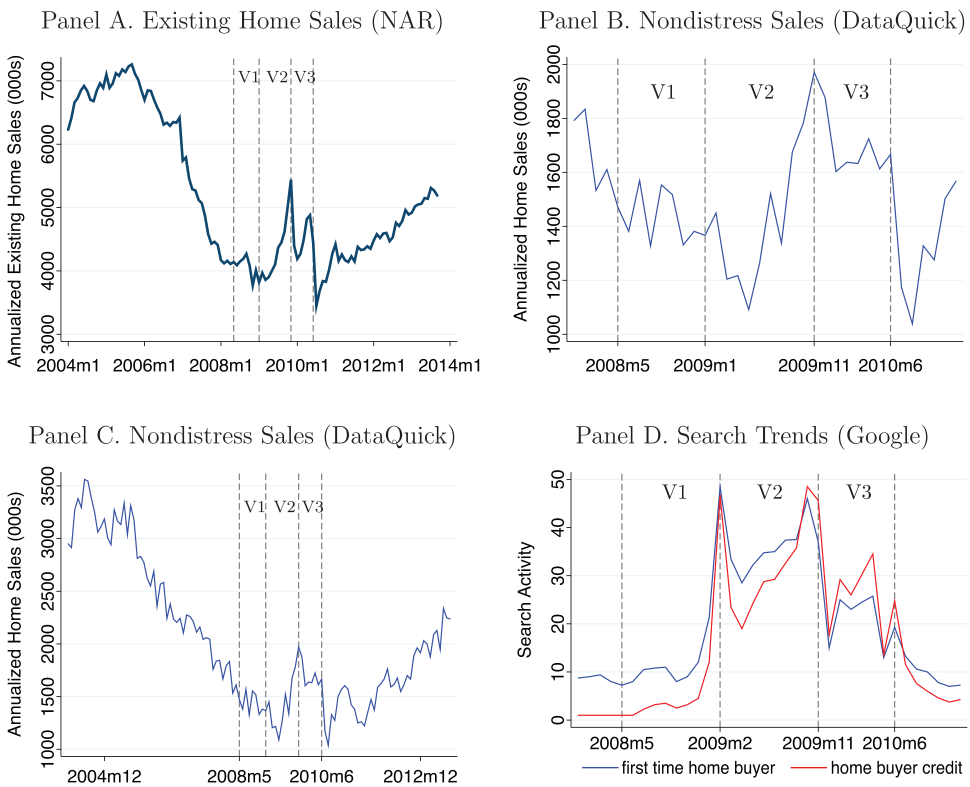

Figure 2. Aggregate home sales and the policy window. Panel A plots existing home sales on a seasonally adjusted annual basis from the National Association of Realtors (NAR). Panels B and C plot seasonally adjusted, annualized monthly home sales from DataQuick along with vertical markers for policy events. These data exclude distress transactions and new construction. Panel D plots Google search trend data for the terms "first time home buyer" and "home buyer credit" along with vertical markers for policy events. The vertical markers are as in Figure 1. (Color figure can be viewed at wileyonlinelibrary.com)

A. Data Sources

We capture program exposure using the population of de-identified individual tax return data, available between 1998 and 2013. We measure homeownership through itemized deductions of mortgage interest, mortgage insurance premiums, and property taxes on Form 1040, Schedule A, or through information return Form 1098 submitted by lenders (which includes interest payments and points paid).[9] The data's panel structure allows us to determine whether a taxpayer owned a home in the past. We also use tax data to measure claims of the homebuyer credit filed on Form 5405. This form records the date

[9] The information return helps identify homeowners who do not itemize their tax returns. Lenders are required to file Form 1098 for all borrowers who pay at least $600 of mortgage interest, points, or insurance premiums during the year.

of purchase, which we use to study the time series of claims. Masked identifiers allow us to link these claims to the individual's tax return to measure the ZIP code associated with that person's claim.

There are two potential issues with our approach to measure homeownership. First, we miss those individuals who own their homes outright and use the standard deduction or who do not file a tax return. These individuals likely make up a very small portion of first-time homebuyers, however, as first-time homebuyers typically buy with a mortgage,[10] and nonfilers consist largely of poor and elderly people. Second, we may mistakenly label refinance events as purchase events. This will be the case only for homeowners who previously owned their homes without a mortgage. This issue introduces measurement error in predicting program responses but is not an obvious confound.

Monthly home sales data come from deed-level recorder and assessor data from DataQuick (now owned by CoreLogic). The records provide details on each transacted home, including date of transfer, price, size, age, number of bedrooms, and number of bathrooms, as well as information on the type of transaction, including short sales, financial institution-owned sales (REO), foreclosures, and an indicator for whether the transaction is made between related parties or at arm's length.

We use information between 2004 and 2013, which yields a consistent sample of covered places over time. Figure 2 shows that the DataQuick housing data closely match the time series patterns for publicly available data published by the NAR. On average, the aggregate counts in our filtered data represent 40% to 50% of the levels reported by NAR.

House price data come from the Federal Housing Finance Agency (FHFA), CoreLogic, and DataQuick. FHFA's repeat sales price indices are available at the yearly level for the largest set of ZIPs in our sample.[11] CoreLogic's repeat sales price indices are available monthly for a smaller set of ZIPs. We compute median prices for ZIPs within our DataQuick home sales sample, which we use in cross-sectional tests based on pre-policy price levels and for back-of-the-envelope calculations.

B. Analysis Sample and Summary Statistics

We construct a ZIP-by-month panel by aggregating individual transactions from the deeds records into counts for various transaction types. The primary analysis sample begins with counts at the ZIP-month level for nondistress sales of existing homes. To ensure that estimates are not biased by changes in geographical coverage, we only include ZIPs with more than 90% of their transaction time series complete from 2006 onward. All other data sets are

[10] Based on survey evidence from 8,449 consumers who purchased a home between July 2009 and June 2010, 96% of first-time buyers used mortgage financing (National Association of Realtors (2017)).

[11] Bogin, Doerner, and Larson (2016) describe the construction and data sources of these price indices. They refer to these indices as "experimental" or "developmental" because FHFA uses the same microdata to produce official price indices at higher levels of aggregation.

Table I

Summary Statistics for Home Sales Analyses

This table presents summary statistics for analysis of the FTHC at the ZIP level. Statistics are presented at both the ZIP-by-month and ZIP levels. Monthly Home Sales include nondistress resales. Section I in the Internet Appendix provides a detailed description of the data sources and variable construction and presents additional statistics.

	Mean	10^{th}	Median	90^{th}	N
Housing transactions					
Monthly home sales (SA)	19.6	3.7	14.2	41.5	1,019,086
Home sales/average monthly sales, 2007	1.04	0.43	0.92	1.79	1,000,860
Program exposure (ZIP)					
First-time buyers/tax units, 2000 (IRS)	2.99	1.92	2.90	4.15	8,883
Cross-sectional characteristics (ZIP)					
Population, 000s (ACS)	23.26	5.58	20.31	45.06	8,883
Unemployment rate, 06-10 avg. (ACS)	7.83	4.30	7.20	12.20	8,883
Average gross income, 2005 (IRS)	62.45	32.12	50.43	99.20	8,883
Subprime card share, 1996 (Equifax)	0.30	0.16	0.28	0.46	8,733
Subprime share of originations (HMDA)	0.21	0.08	0.18	0.39	8,842
FHA expansion exposure (HMDA)	0.15	0.02	0.13	0.33	8,842
HARP exposure (HMDA)	0.21	0.10	0.22	0.30	8,842
HAMP exposure (HAMP/HMDA)	0.05	0.02	0.04	0.09	8,880
Median age, 06-10 avg. (ACS)	38.51	31.70	38.40	45.00	8,883
Median rent, 06-10 avg. (ACS)	970.76	637.00	910.00	1397.00	8,883
Fraction below poverty line (ACS)	12.05	3.60	9.80	23.80	8,883
Urban share of census blocks (Census)	83.30	39.80	99.10	100.00	8,883

filtered to restrict the analysis sample to the same set of ZIPs. The primary sample contains 1,019,086 ZIP-months for 8,883 ZIPs across 47 states. These ZIPs account for 69% of the U.S. population in 2007.

Table I presents summary statistics for the sample used in the home sales analysis. We seasonalize home sales counts using a within-ZIP transformation for each month. For each ZIP, we also compute the mean of monthly house sales in 2007, which we use as our primary scaling and weighting variable. Our main outcome variable is scaled monthly sales of existing homes, excluding distressed or forced sales, censored at the 99% level to remove outliers. The average observation has 19.6 sales per month. This varies from 3.7 sales at the 10^{th} percentile to 41.5 at the 90^{th}. The 10^{th} percentile of the scaled variable is 0.43, the median is 0.92, and the 90^{th} percentile is 1.79.

III. Main Empirical Approach

Our empirical strategy exploits cross-sectional variation across geographies in ex ante exposure to the FTHC program to isolate the effect of the program from aggregate macroeconomic shocks. Mian and Sufi (2012) and Chodorow-Reich et al. (2012) use this approach to estimate the effect of fiscal policy. Its main advantage is that it produces a counterfactual for estimating what would have happened absent the policy. Areas with few potential first-time homebuyers serve as the "control group," because buyers in these areas would likely be ineligible for the credit. The difference between treatment and control areas provides an estimate of the causal impact of the program.[12]

The policy targeted first-time homebuyers, so we expect the impact to be greater in areas where historically first-time homebuyers tend to buy. Accordingly, we measure program exposure by identifying locations with more first-time buyers in a given period prior to the policy. Higher exposure may reflect local amenities, such as schools or parks that attract first-time buyers, or it may reflect a local housing stock better suited to these buyers in terms of affordability, lot size, etc. We construct the exposure measure at the ZIP code level to study the effect of the policy on market-level outcomes such as house prices. These local general equilibrium or market effects would be missed using an individual-level identification strategy.[13]

We use individual tax and information returns to capture the number of first-time homebuyers in each ZIP. First-time homebuyers are people identified as homeowners in t but not in $t-1$ or $t-2$. Exposure is the number of first-time homeowners in a ZIP in 2000 scaled by the number of tax-filing units in 2000. There is significant variation in our measure of exposure at the ZIP code level (Table I). In particular, program exposure varies from 1.92% at the 10$^{\text{th}}$ percentile to 4.15% at the 90$^{\text{th}}$. Mean exposure is 2.99%.

Figure 3, Panels A and B, shows that there is significant variation in exposure to the program across areas. Figure 3, Panel A, maps county-level variation across the country, and Figure 3, Panel B, maps ZIP-level variation for Boston, Chicago, and San Francisco. Darker areas indicate more exposure to the program.

Exposure is relatively concentrated in suburban areas around cities where first-time homebuyers tend to locate. To confirm this observation, Table II presents bivariate regressions of program exposure on ZIP-level observables. ZIPs with high exposure have higher rents and fewer people below the poverty line. The populations are larger and somewhat younger. Income is weakly correlated with program exposure. Substantial variation in ex ante exposure within cities allows us to pursue a research design that conditions on city-time fixed effects.[14]

[12] This logic requires the assumption that induced transactions in one group do not systemically spill over via transactions induced by "real estate chains" into either treatment or control areas. We find limited evidence that real estate chains affect the results. Below we discuss how relaxing this assumption affects our interpretation.

[13] This approach permits estimation of the aggregate effect of the program under various assumptions about the macroelasticity. For aggregation we assume that this elasticity is a constant function of exposure, but we relax this assumption and estimate heterogeneous effects in robustness tests. Our approach does not naturally map into a person-level microelasticity of housing demand. To explore the latter, we deploy a regression kink research design that exploits the program's income eligibility threshold (see Section V).

[14] We use CBSAs to define city boundaries. Although exposure varies at the ZIP level, we cluster standard errors at the CBSA level to permit within-city correlation in error terms.

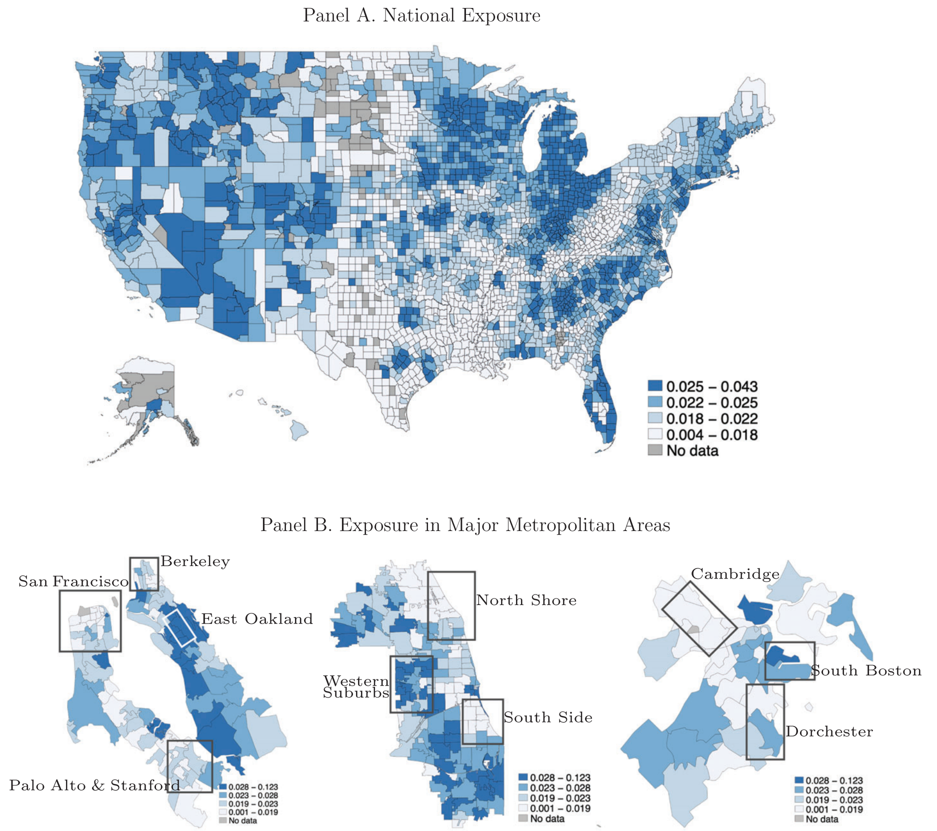

Figure 3. Maps of FTHC program exposure. Panel A presents a county map of program exposure, defined as the number of first-time homebuyers in an area in 2000 divided by the number of tax filers in 2000. Panel B presents ZIP code maps for three metro areas: from left to right, the San Francisco Bay Area, Chicagoland within Cook County, and Boston and Cambridge. Darker shadings reflect higher exposure. (Color figure can be viewed at wileyonlinelibrary.com)

To address the concern that we may not accurately capture program exposure, we show that areas with higher ex ante exposure did indeed see more individuals claim the credit. Figure 4, Panel A, plots binned bivariate averages ("binscatters") of FTHC claims from tax records against program exposure. Exposure is strongly correlated with take-up in the cross-section. The regression coefficient on the specification with Core-Based Statistical Area (CBSA) fixed effects and ZIP-level controls is 0.59 with a clustered t-statistic of 20.

Figure 4, Panels B and C, shows that our exposure measure also predicts time series variation in claims in these areas. We plot counts of FTHC claims by month of home purchase for purchases made between February 2009 and September 2010 along with vertical markers for policy events. The vertical markers correspond to the start of the FTHC loan program, the start of version two of the credit, the scheduled expiration of version two, and the actual

Table II

Correlates of Program Exposure

This table presents bivariate regressions of program exposure on ZIP-level observables. Variables have been normalized, so the coefficients can be interpreted as a one-standard-deviation change in x produces a β-standard-deviation change in exposure, where β is the reported coefficient. Standard errors are clustered at the CBSA level. *, **, and *** denote significance at 0.10, 0.05, and 0.01 levels, respectively.

	LHS is Exposure		
	Coefficient	R^2	N
Exposure correlates			
Median age	−0.052*	0.0027	8,883
	(0.030)		
Median rent	0.192***	0.0371	8,883
	(0.053)		
Fraction below poverty line	−0.280***	0.0784	8,883
	(0.033)		
Fraction classified as urban	0.078***	0.0062	8,883
	(0.022)		
Controls			
Log(population)	0.077***	0.0060	8,883
	(0.027)		
Unemployment rate	−0.102***	0.0103	8,883
	(0.033)		
Log(average gross income)	0.025	0.0006	8,883
	(0.034)		
Subprime cardholder fraction	−0.016	0.0003	8,733
	(0.039)		
Subprime in 2004 to 2007	0.069**	0.0048	8,842
	(0.033)		
FHA expansion exposure	−0.044	0.0019	8,842
	(0.039)		
HARP exposure	−0.002	0.0000	8,842
	(0.034)		
HAMP exposure	−0.029	0.0008	8,880
	(0.034)		

expiration of version three. Figure 4, Panel B, plots national claim counts month-by-month. Figure 4, Panel C, plots claim counts for high- and low-exposure quintiles of ZIP codes.[15] Our exposure measure predicts not only more FTHC claims in high exposure areas, but also the spikes in claims observed in the national claims data.

Numerous robustness tests confirm that our main results do not depend on the precise way in which we measure exposure. Consistent with an area's exposure being a slow-moving characteristic, Figure IA.I in the Internet Appendix shows that exposure is highly correlated over time and that alternative measures of exposure yield nearly identical take-up predictions.

[15] Quintiles are formed using weights that ensure each quintile has equal population in 2007.

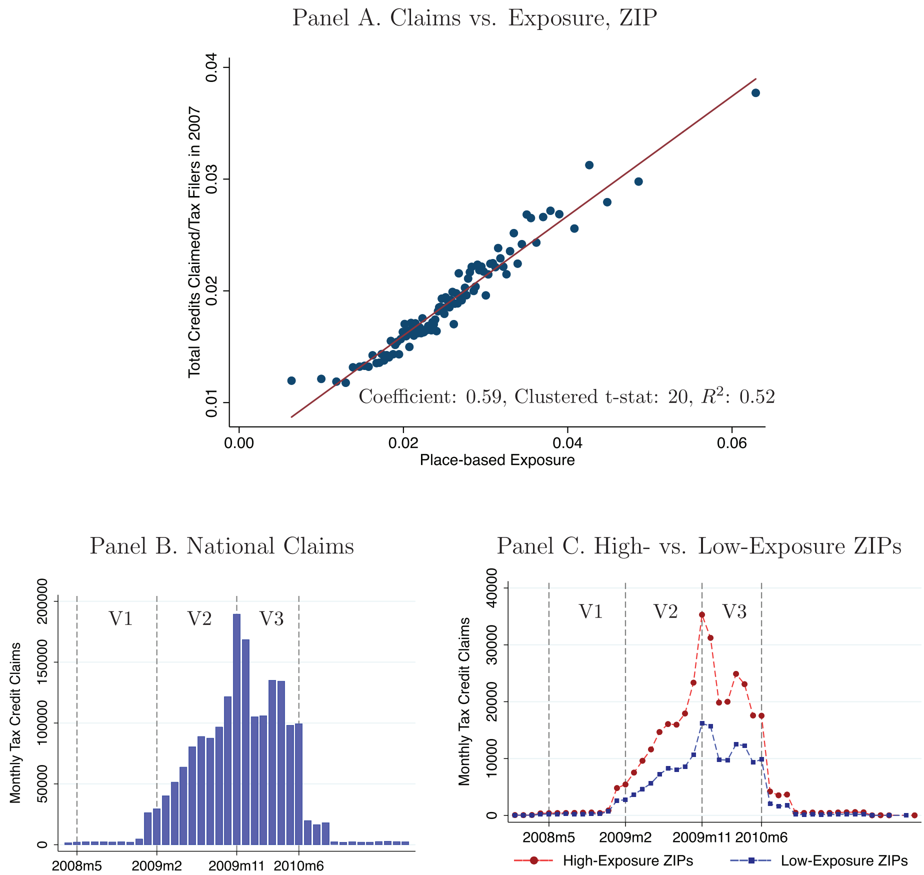

Figure 4. Program exposure and FTHC claims. Panel A plots a binscatter (100 bins) of ZIP-level FTHC claims from tax records scaled by the number of tax filers in 2007 versus program exposure. Exposure is defined as the number of first-time homebuyers in an area in the year 2000. Panel B plots national counts of FTHC claims by month of home purchase for purchases between February 2009 and September 2010. Panel C plots claim counts for high- and low-program-exposure quintiles of ZIPs sorted using program exposure. The quintiles are formed using weights that ensure each quintile has equal population in 2007. Panels B and C only include nonamended claims for versions two and three of the program, as our data do not include month of purchase information for amended returns or claims for version one of the program. Panels B and C include vertical markers for policy events defined as in Figure 1. (Color figure can be viewed at wileyonlinelibrary.com)

An additional concern with our measure is that, while it is strongly correlated with FTHC take-up, unobservable factors unrelated to the FTHC program may drive differential purchase patterns. After all, places where first-time homebuyers typically buy are not random. For example, a risk to our design is that our measure captures the expansion in subprime credit documented by Mian and Sufi (2009), leading to different ZIP-time trends within cities as the cycle corrected. To mitigate this risk, our preferred measure is the number of first-time homebuyers in 2000, a pre-subprime period, that is not driven by the increase in subprime purchases later in the decade.[16] We also control for the subprime share of borrowers in 1996 (Mian and Sufi's (2009) measure) and the subprime share of loans over the period 2004 to 2007.

To mitigate concerns about differential ZIP-time trends, our baseline analysis conditions on city-time fixed effects and we report results with and without observable controls. This approach removes many potential confounds from our analysis. In addition, we explicitly test for parallel trends in the pre-period, we perform within-ZIP placebo tests, we explore the effects of differences in credit generosity, and we control for contemporaneous housing market policies. We also exploit information in the age distribution of first-time homebuyers over time to show that the median age of first-time homeowners falls during the policy period and the age distribution reverts immediately after the policy expires. Moreover, the highest exposure ZIPs account for the largest share of the shift in the first-time homebuyer age distribution. Finally, we note that the short-lived nature of the policy and the sharpness of the time series responses weigh against alternative stories that operate at lower frequencies.

IV. The Effect of the FTHC on Home Sales

A. Home Sales

We begin with a graphical analysis that demonstrates our main finding on quantities: home sales respond strongly to the FTHC program but do not show a sharp, immediate reversal in the year following the policy period.

Figure 5, Panel A, plots monthly home sales between July 2007 and November 2011 for ZIPs divided into 100 quantiles and sorted based on ex ante program exposure. We present these data in the form of a calendar-time heatmap, which is analogous to the traditional two-group calendar-time graph but plots visually discernible time series for many groups. Columns correspond to months, and rows to groups of ZIPs sorted by exposure. Each cell's shading corresponds to the level of the outcome variable, which is monthly home sales scaled by average monthly home sales in 2007.

The heatmap yields four observations. First, high- and low-exposure series closely track each other each month prior to the policy, deviating only during the policy window. Note that each sequence of consecutive months in the pre-period provides a placebo test that fails to reject the design's core identification assumption of parallel trends (see the discussion of Figure 5, Panels B to D, for statistical tests). Second, the smoothly increasing gradient visible at each

[16] The year 2000 is the earliest year for which at least one year of information returns from lenders are available to classify past homeownership. Tables IA.V and IA.VI in the Internet Appendix show that our results do not depend on the year in which we measure exposure due to its high persistence over time. Because location-based exposure is persistent over time, choosing an early year does not fully address the subprime concern. Examining pre-trends is therefore key to evaluating the role of confounding factors.

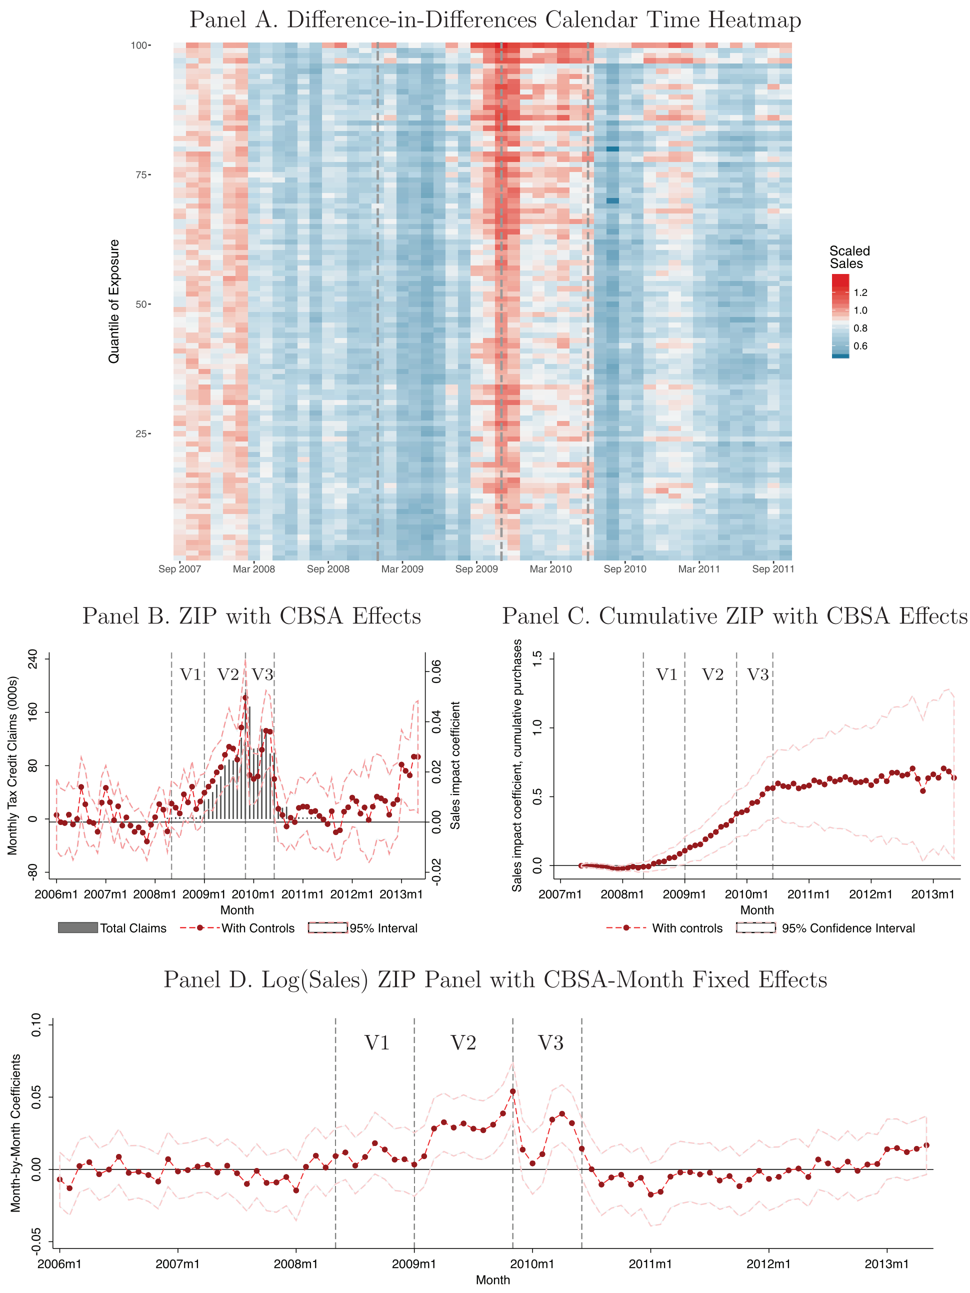

Figure 5. The effect of the FTHC on home sales. The panels in this figure plot the monthly and cumulative effects of the FTHC on nondistress resales at the ZIP level. Panel A plots a difference-in-differences calendar-time heatmap of monthly sales for ZIPs divided into 100 quantiles and sorted based on program exposure. Columns correspond to months and rows correspond to groups of ZIPs sorted by exposure. Exposure is the number of first-time homebuyers in a ZIP in 2000

policy expiration date shows that the policy response is monotone in ex ante exposure and not driven by a few outlier ZIPs. Third, the gradient does not reverse sign in the 15 months following the second policy expiration, but rather the series return to the baseline pattern of parallel trends, that is, the data do not indicate a sharp reversal of the policy response. Last, we use the lowest exposure quantile as a counterfactual to estimate the cumulative number of sales induced by the program. The heatmap shows that this group is a credible counterfactual, as it indicates no response to the program during the policy period. In Section IV.B, we find further support for the bottom quantile as a counterfactual based on the relative change in the age distributions of first-time buyers.

Figure 5, Panel B, plots coefficients from regressions that estimate the monthly effects of the program. We run month-by-month regressions of the form

$$\frac{Home\ Sales_i}{Average\ Monthly\ Sales_{i,2007}} = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i, \tag{1}$$

where $Exposure_i$ is the ZIP-level program exposure for area i and α_{CBSA} is a CBSA-specific constant.[17] In controls specifications, X_i includes log population, the average unemployment rate from 2006 through 2010, log average gross income in 2005, the subprime share in 1996, the average share of subprime originations from 2004 to 2007, exposure to the 2009 Federal Housing

[17] For the 129 ZIPs without an associated CBSA, we assign them a state-specific constant.

scaled by the number of tax filing units in 2000. Each cell's shading corresponds to a level of the key outcome variable, which is monthly home sales scaled by average monthly home sales in 2007. The quantiles are formed using weights that ensure each quantile has an equal number of home sales in 2007. Panel B plots coefficients for monthly home sales regressions with controls, and overlays national counts of FTHC claims by month of home purchase. Panel C plots coefficients for cumulative sales regressions. In Panels B and C, we run month-by-month regressions, weighted by total home sales in 2007, of the form

$$\frac{y_i}{Sales_{i,2007}} = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i,$$

where y_i is either monthly home sales in area i or cumulative monthly home sales in area i beginning 17 months before the program. X_i is a set of controls that include log population, the average unemployment rate from 2006 through 2010, the log of average gross income, the subprime share in 1996, the average share of subprime originations from 2004 to 2007, exposure to the FHA expansion, and exposure to the HARP and HAMP programs. In Panel D, we run a panel regression, weighted by total home sales in 2007, of the form

$$y_{it} = \alpha_i + \delta_{CBSA,t} + \sum_t \beta_t Exposure_i + \sum_t \gamma_t X_i + \varepsilon_{it},$$

where y_{it} is the log of monthly home sales in area i from 2004 onward (2004m1 through 2005m12 are omitted). Program exposure is normalized by its cross-sectional standard deviation. (Color figure can be viewed at wileyonlinelibrary.com)

Administration (FHA) expansion, and exposure to the Home Affordable Refinance Program (HARP) and Home Affordable Modification Program (HAMP) programs. All regressions are weighted by average monthly home sales in 2007. This approach is approximately equivalent to a panel regression with time-specific coefficients on exposure and the control variables, and with ZIP, month, and CBSA-month fixed effects.[18] To aid interpretation, we normalize exposure by its cross-sectional standard deviation.

Figure 5, Panel B, plots coefficients for these regressions with controls. The results confirm those from the heatmap. Exposure patterns do not predict differences in sales activity until the policy window begins, and the coefficients spike in accord with the aggregate series. The coefficient of 0.05 for November 2009 implies that a one-standard-deviation increase in program exposure produces a 5% increase in monthly home sales relative to the average level in 2007. This is approximately 0.12 standard deviations of the left-hand-side variable. During the policy period, the coefficients closely mirror the time series of national FTHC claims by month of home purchase (the correlation is 0.66).

Figure 5, Panel C, plots coefficients for regressions that replace monthly sales with cumulative monthly sales. The series is approximately flat prior to the policy window, increases monotonically over the window, and flattens in the post-period. The cumulative effects are between 50% and 60% relative to the average level of monthly sales in 2007. These figures should not be confused with aggregate estimates, which we provide below. Again, we see no evidence of a sharp negative relationship between sales and exposure in the 17 months following the policy. At longer horizons—between 1.25 and three years after the policy expired—our cumulative regressions lose statistical power as each subsequent month adds noise and increases standard errors. Still, we can statistically reject a full reversal through 2012.

Figure 5, Panel D, plots coefficients from a panel specification with CBSA-month and ZIP fixed effects. We estimate the following regression at the ZIP level from 2004 onward, weighted by total home sales in 2007:

$$\log(\textit{Home Sales})_{it} = \alpha_i + \delta_{CBSA,t} + \sum_t \beta_t \textit{Exposure}_i + \sum_t \gamma_t X_i + \varepsilon_{it}, \qquad (2)$$

where β_t is a month-specific exposure coefficient and γ_t is a vector of month-specific coefficients on ZIP-level controls. We plot coefficients on exposure from January 2006 onward, omitting January 2004 through December 2005. The coefficients closely match those in Figure 5, Panel B, both in pattern over time and in magnitude.

Table III presents the average monthly effects of the FTHC on home sales pooled over different policy windows for a variety of specifications. We run

[18] This cross-sectional approach follows Mian and Sufi's (2012) evaluation of the Cash for Clunkers program, which aids comparison with their findings. We also run the standard difference-in-differences specification via panel regression, as advocated by Bertrand, Duflo, and Mullainathan (2004) and employed by Best and Kleven (2017). Figure 5, Panel D, and Table IA.I in the Internet Appendix present estimates based on this approach. The results are very similar.

Table III

The Effect of the FTHC on Home Sales

This table presents the average monthly effects of the FTHC on home sales for ZIPs pooled over different policy windows with various specifications. We run cross-sectional regressions, weighted by average monthly home sales in 2007, of the form

$$\frac{\overline{Sales}_{i,t\to T}}{Sales_{i,2007}} = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i,$$

where y_i is average monthly home sales in area i over the relevant time period. In controls specifications, X_i is a set of controls that include log population, the average unemployment rate from 2006 through 2010, the log of average gross income, the subprime share in 1996, the average share of subprime originations from 2004 to 2007, exposure to the FHA expansion, and exposure to the HARP and HAMP programs. Exposure is normalized by its cross-sectional standard deviation. Column (3) includes CBSA fixed effects. Column (4) presents unweighted regressions. Column (5) excludes Arizona, California, Florida, and Nevada. Standard errors are clustered at the CBSA level. *, **, and *** denote significance at 0.10, 0.05, and 0.01 levels, respectively.

	No Controls (1)	Controls (2)	CBSA FE (3)	No Weights (4)	Ex Sand (5)
Pre-policy	−0.001	0.002	0.004	0.002	0.000
2007m9 to 2009m1	(0.004)	(0.004)	(0.002)	(0.003)	(0.003)
Observations	8,883	8,692	8,692	8,692	6,664
R^2	0.000	0.033	0.365	0.302	0.365
Policy	0.027***	0.024***	0.026***	0.030***	0.020***
2009m2 to 2010m6	(0.010)	(0.009)	(0.005)	(0.007)	(0.005)
Observations	8,883	8,692	8,692	8,692	6,664
R^2	0.013	0.131	0.475	0.451	0.446
Post-policy	0.017	0.016*	0.004	0.007	0.001
2010m7 to 2011m11	(0.012)	(0.010)	(0.005)	(0.008)	(0.005)
Observations	8,883	8,692	8,692	8,692	6,664
R^2	0.003	0.234	0.579	0.536	0.506
Early policy	0.016*	0.014*	0.021***	0.024***	0.016***
2009m2 to 2009m9	(0.009)	(0.008)	(0.005)	(0.007)	(0.005)
Observations	8,883	8,692	8,692	8,692	6,664
R^2	0.006	0.078	0.413	0.371	0.431
Spike 1	0.045***	0.039***	0.037***	0.044***	0.032***
2009m10 to 2009m12	(0.012)	(0.011)	(0.006)	(0.008)	(0.006)
Observations	8,866	8,675	8,675	8,675	6,661
R^2	0.019	0.123	0.437	0.409	0.359
Spike 2	0.037***	0.031***	0.034***	0.035***	0.029***
2010m4 to 2010m6	(0.011)	(0.010)	(0.007)	(0.009)	(0.006)
Observations	8,875	8,684	8,684	8,684	6,663
R^2	0.014	0.088	0.374	0.365	0.379
Controls	No	Yes	Yes	Yes	Yes
CBSA FE	No	No	Yes	Yes	Yes

cross-sectional regressions of the form

$$\frac{Average\ Monthly\ Sales_{i,t\rightarrow T}}{Average\ \ Monthly\ \ Sales_{i,2007}} = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i, \tag{3}$$

where the left-hand-side numerator is average monthly home sales in area i over the relevant time period. We use the same control set, weighting, and specification for exposure as in Figure 5, Panels B and C.

The results of the pooled regressions confirm the patterns from the figures. In the pre-policy period, there is little sign of differential trends. The 17-month policy period shows a significantly greater average effect on monthly sales, with this effect most pronounced during the two windows leading up to policy expiration. The CBSA fixed effects specification gives our preferred estimate of 2.6% higher sales per month (relative to average monthly sales in 2007) in response to a one-standard-deviation increase in exposure. The first spike shows a somewhat stronger but statistically indistinguishable effect relative to the second spike. One potential explanation for this result is that the second period included the LTHC, which our exposure measure is not designed to predict. The post-policy period yields coefficients that are approximately equal to 0 and statistically insignificant, indicating that there is little to no reversal in the 17 months after the policy ends.

While the policy was able to significantly increase sales, these sales did not reverse for at least two years. The lack of reversal is surprising, since standard intertemporal theory suggests that temporary price subsidies for durable goods simply reallocate demand across time. In line with this rationale, Mian and Sufi (2012) and Green et al. (2016), who study the Cash for Clunkers (CARS) program, find that while the program was able to stimulate demand for cars during the policy period, these sales reversed completely after seven to 12 months.

We believe that home sales do not reverse in the post-FTHC period for two key reasons. First, buyers were induced to purchase a home during the policy period instead of from one of many more distant future years. The ability to pair the credit with a low down payment loan enables earlier transitions into homeownership among otherwise-constrained buyers (see Section VI.B). In addition, the relatively long policy window (16.5 months versus two months for CARS) allows information about the program to spread beyond those already in the market. Both features allowed the FTHC to draw demand from farther into the future. To the extent that sales would have occurred over several years or were pulled from many years in the future, they will be difficult to detect in our data. However, in the next section, we provide evidence based on the change in the median age of first-time buyers that suggests sales were indeed pulled from several years in the future.

Second, the credit may have induced new buyers who would not have become homeowners absent the credit. In contrast to CARS, the FTHC targeted *new* potential homeowners, which allows a second, *extensive-margin* effect. Best and Kleven (2017) study the effects of similar programs in the United Kingdom and find that such extensive margin purchases can be substantial, even when the subsidy is small relative to the level of house prices. In our setting, it is difficult to differentiate extensive margin purchases from those accelerated from future years. In a follow-up paper, Berger et al. (2018) use an estimated structural model to explore the quantitative magnitudes of the retiming effects and extensive margin effects.

We explore two additional reasons for the nonreversal in sales: (1) buyers may move quickly and "trade up" their homes, and (2) the homebuying chain obscures reversion by generating additional sales. Conceptually, in both cases the additional transactions could continue to boost cumulative sales. If a first-time buyer quickly upgrades her home, then two additional sales could result—the purchase of the upgraded home and the sale of the first-time home to a new buyer. However, we do not find evidence that this factor provides a meaningful effect on the nonreversal of cumulative sales in the immediate post-policy period.[19] Likewise, in the case of the homebuying chain, the seller of the home to the first-time buyer is likely to subsequently buy a house. Unfortunately, our data do not permit us to track comprehensively whether and where sellers subsequently buy. Section VI.B presents evidence showing that the FTHC increased sales of vacant properties. To the extent that the credit effectively brought these homes back into use, there may not be an offsetting reversal or amplifying real estate chain effect.

B. The Age Distribution of First-Time Buyers

The nonreversal of the policy-period response following the program's expiration raises questions about the extent to which buyers pulled these purchases forward in time. In this section, we present direct evidence indicating that, in the absence of the program, many buyers would not have bought homes for several years.

Figure 6, Panel A, plots age distributions of first-time homebuyers identified between 2002 and 2013. We highlight the age distribution for 2009, which shifts substantially to the left relative to other years. The age distribution over the 2002 to 2008 period is considerably older than that for the 2009 cohort. In the four post-policy years from 2010 to 2013, the age distribution closely resembles that in the pre-policy years, although with more mass in the early and mid-30s and less mass between the ages of 40 and 50. The median age across all first-time buyers is 35 in the nonpolicy years and 33 in 2009. Among

[19] To test the idea that first-time buyers move out of their home quickly, we used a large random sample of credit report data that allow us to track the first-time buyer through 2018. Based on these data, we do not find evidence of a meaningful treatment effect on the likelihood of moving during the first five years of ownership. Over a longer horizon, we find small and statistically significant effects on the likelihood of moving, indicating a relatively slow reversal in sales. Section II in the Internet Appendix discusses evidence reported in Table IA.VII on whether some market segments experience short-run reversal. The results suggest that short-run reversal is more likely when induced buyers have higher income, are not liquidity constrained, and participate in markets where buyers are more mobile. However, such markets account for a small share of total sales.

Panel A. Distribution of First-Time Buyers, 2002–2013

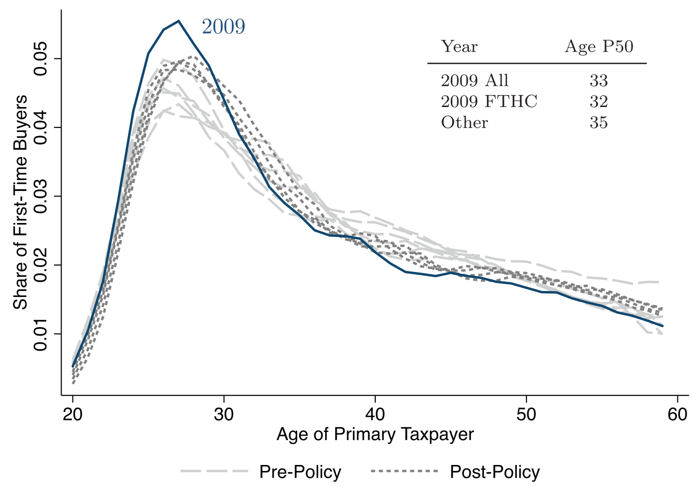

Panel B. Distribution of First-Time, High vs. Low Exposure, 2002–2013

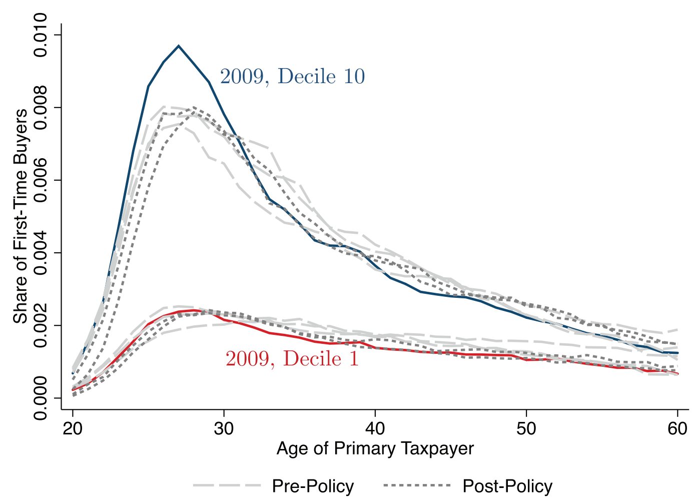

Figure 6. Policy shift in the age distribution of first-time buyers. Panel A plots age distributions of first-time homebuyers identified using income tax return and information return data for the period 2003 to 2013. Panel B shows the relationship between the shift in the age distribution in 2009 and program exposure. For readability, only odd-numbered years are plotted. As in Panel A, the denominator for these distributions is the total number of first-time homebuyers each year. The numerator is the number of buyers for each age in either the highest decile ZIPs (top group of curves) or the lowest decile ZIPs (bottom group of curves) of program exposure. The highest exposure ZIPs account for the largest share of the shift in first-time homebuyer age observed in the aggregate data. The lowest exposure ZIPs show a limited shift in first-time homebuyer age, as well as a much lower overall share of first-time homebuyers. The FTHC was primarily in effect in 2009, highlighted in each graph. All other years are in gray, with pre-policy years in long dashes and post-policy years in short dashes. (Color figure can be viewed at wileyonlinelibrary.com)

FTHC claimants, the median age is 32, three years younger than the typical first-time buyer in other years.

The comparability of the age distributions in the pre- and post-policy periods mitigate concerns that area-time cyclicality, pent-up demand, or secular trends explain the slow reversal in our research design. For example, because many first-time homebuyers buy with FHA-insured loans, the expansion of FHA's loan limits in 2009 might interfere with our approach. However, such a confound would predict the shift in the age distribution of first-time homebuyers to continue in the years after the FTHC expired, contrary to the temporary shift in the data. The results instead suggest that a noticeably younger cohort of first-time buyers appears only in 2009, driven by the temporary policy incentive to accelerate the transition into homeownership.

Figure 6, Panel B, provides evidence in line with the view that the FTHC explains this pattern. This figure illustrates the relationship between the shift in the age distribution in 2009 and program exposure by plotting age distributions over time for the top and bottom treatment deciles. The graph further validates our cross-sectional research design in two ways. First, there are more first-time homebuyers in higher treated areas, and these areas account for the largest share of the shift in homebuyer age observed in the aggregate data, validating our designation of treated ZIPs.[20] Second, the lowest exposure ZIPs show a limited shift in first-time homebuyer age, as well as a much lower overall share of first-time homebuyers, consistent with our assumption of a limited treatment effect in low treatment areas. The similarity of the age distributions among the lowest deciles over time also helps rule out confounds from within-city sorting (Landvoigt, Piazzesi, and Schneider (2015)) that may affect lower treatment areas.

C. Robustness and Placebo Tests

Table III presents results of several additional robustness tests of our key findings. The estimates are similar with and without CBSA fixed effects, although somewhat more precise in the former specification, and change little when excluding controls. Unweighted regressions lead to modestly larger estimates during the policy window. Column (5) excludes the sand states—Arizona, California, Florida, and Nevada—which weakens the policy-period estimates only slightly. In general, the estimates are very similar across states.[21]

[20] Figure IA.3, Panel A, in the Internet Appendix shows that this pattern generally holds, as the tilt in the age distribution toward young buyers increases monotonically in program exposure across ZIPs. Figure IA.3 in the Internet Appendix, Panel B, shows that the increase in the share of under-30 buyers in 2009 is strongly correlated with program exposure.

[21] Table IA.I in the Internet Appendix replicates the baseline results and these robustness checks using a log specification based on Best and Kleven (2017). The table also considers an alternative pre-period that separates version one of the FTHC program (i.e., the loan program spanning 2008m5 to 2008m12) and the two years prior to version one. The table suggests modest effects of version one that are approximately one-third the size on a monthly basis and 16% on a cumulative basis relative to versions two and three (i.e., the grant program).

Importantly, the parallel trends assumption test yields coefficients close to zero across specifications, and we find similarly null average post-policy effects.[22]

The pre-policy coefficients provide strong evidence that our design is valid, that is, that low exposure areas serve as a counterfactual to high exposure areas. The sharp timing of the policy addresses many concerns about omitted variables because most potential confounds move more slowly. Yet some concerns remain. One concern is that time-varying, area-specific shocks are correlated with our exposure measure. For example, suppose that our exposure measure is highly correlated with the share of subprime borrowers, which peaked during the 2004 to 2007 period. If true, then the increase in sales during the policy period could be driven by "pent-up" subprime demand and not the FTHC. While the inclusion of CBSA-month fixed effects helps mitigate this concern, there is still significant variation in subprime borrowing within cities. However, Table II shows that our exposure measure is essentially uncorrelated with the share of subprime borrowers—using both Mian and Sufi's (2009) subprime measure and a measure of subprime borrowing from 2004 to 2007—which suggests that our main results are not driven by a subprime, pent-up demand effect.[23]

A related concern is that area-specific trends beginning in 2009 might confound our estimates. Such trends could be driven by coincident policies designed to shore up the housing market or by area-specific cyclicality. We address this threat in three ways.

First, we construct controls designed to capture the exposure of ZIP codes to different housing market programs. The most relevant coincident program is the expansion of the FHA loan limits in 2009 as part of the Economic Stimulus Act of 2008, which persisted in the years following the FTHC program. We measure exposure to this FHA policy change using the average share of loan originations in HMDA data in 2004 to 2007 that would have been eligible under the new regime. Figure IA.4 in the Internet Appendix shows that this measure strongly predicts FHA take-up at the ZIP code level, but is weakly negatively correlated with FTHC exposure (the correlation of −0.04 is reported in Table II). Because the FHA policy change targeted larger mortgages, it affected buyers in higher price local areas that are generally too expensive

[22] Figure IA.2 in the Internet Appendix presents a placebo test that further confirms these findings. The test estimates month-by-month regressions and plots coefficients from the noncontrol specification in Figure 5, Panel B, emphasized with a bold line, along with equivalent regressions shifted both backward in time to start in 2005, 2006, and 2007 and forward in time to start in 2009 and 2010. The policy coefficients are unusually high while the pre- and post-policy coefficients coincide with the placebo series. The figure suggests that seasonal confounds not captured by our seasonality adjustment do not influence our estimates of the spikes.

[23] An additional concern comes from Landvoigt, Piazzesi, and Schneider (2015), who document that cheaper areas in San Diego had larger swings in the boom–bust cycle. This differential sensitivity could potentially explain our results since we are identifying effects from the difference between high- and low-exposure areas. The correlation between our quantity results and the sharp timing of the policy suggests that this pattern, which operates at lower frequencies, is not driving our results.

for first-time homebuyers.[24] We also follow Agarwal et al. (2017a, 2017b) to construct measures of exposure to the HARP and HAMP programs, which attempted to alleviate debt burdens among underwater borrowers. For HARP, we use the ZIP-level share of loan originations purchased by Fannie Mae or Freddie Mac over 2004 to 2007. For HAMP, we use an estimate of the ZIP-level share of mortgage modifications under the HAMP program. Table II shows that these programs also targeted different areas than the FTHC program, as correlations are close to zero. We include exposure measures for each of these programs in our control set, but weak correlations imply that our estimates do not depend on these controls.

Second, we explore whether the effects are larger in places where initial price levels are low. For homes with prices above \$80,000, the FTHC is fixed at \$8,000. Thus, the subsidy is relatively less generous in more expensive places. In the first row of Table IV, we estimate a differenced version of equation (3):

$$\frac{\Delta Average\ Home\ Sales_i}{Average\ \ Monthly\ \ Sales_{i,2007}} = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i, \tag{4}$$

where Δ*Average Home Sales* equals the average number of home sales in area i during the policy period minus the average number of home sales in area i during the 17-month pre-period. We first reproduce the results using specifications from Table III to confirm that the estimates are unchanged.[25]

Columns (6) and (7) in the first row of Table IV divide the sample of ZIPs into the bottom three ("Low p") and top three ("High p") deciles in median house prices during 2008. The effects are concentrated in the low-price ZIPs, which yield a coefficient of 0.026, while the high-price ZIPs show no discernible effect with a coefficient of 0.000. Figure 7 plots coefficients and confidence intervals for regressions from each decile of initial house prices. The coefficient declines monotonically as initial price levels increase and the corresponding generosity of the credit declines. These split-sample findings provide further evidence that our results are indeed due to the FTHC policy. In addition, because the FHA expansion targeted higher price areas, these results lend further support to the view that our results capture an FTHC program effect.

Third, we consider an alternative approach to validate our design with a within-time placebo test. First-time buyers are more likely to buy smaller homes than larger homes, so smaller homes should respond more strongly to the program. If ZIP-level shocks are driving our results, we should see similar patterns across all types of homes. Table IA.III in the Internet Appendix

[24] The more relevant part of the FHA expansion for our purposes came through loans that qualified under the pre-crisis regime. This expansion occurred as the private market for second liens to reduce down payments contracted. See Department of HUD Mortgagee Letter 2009-07 and Goodman, Seidman, and Zhu (2014) for more details on the FHA loan limit expansion. A distinct mortgage assistance program called FHA Secure, created in August 2007 and expanded in July 2008, is also unlikely to bias our results, as this program enrolled only a few thousand borrowers.

[25] Tables IA.V and IA.VI in the Internet Appendix confirm the results in Table IV using alternative measures of program exposure (respectively, the 2007 number of first-time homebuyers and the average number of first-time homebuyers over the period 2000 to 2007).

Table IV

The Effect of the FTHC on Home Sales

This table presents regressions of the average monthly effects of the FTHC on different categories of home sales. We run cross-sectional regressions, weighted by average monthly home sales in 2007, of the form

$$y_i = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i,$$

where y_i is a housing market outcome in area i over the relevant time period. In the first row, the outcome is the difference in average monthly nondistress home resales for the policy period versus the 17-month pre-period. In the second row, the outcome is the difference in average monthly new construction sales for the policy period versus the pre-period. In the third row, the outcome is the difference in average monthly foreclosures and short sales for the policy period versus the pre-period. Exposure is normalized by its cross-sectional standard deviation. All columns include CBSA fixed effects and a set of controls as defined in Table III. Column (2) presents unweighted regressions. Column (3) excludes Arizona, California, Florida, and Nevada. Column (4) trims the left-hand-side variable at the 5th and 95th percentiles. Column (5) restricts the sample to areas with average home sales in 2007 above the 10th percentile. Columns (6) and (7) divide the sample of ZIPs into the bottom three ("Low p") and top three ("High p") deciles in median house prices during 2008. Standard errors are clustered at the CBSA level. Tables IA.V and IA.VI in the Internet Appendix confirm these results using alternative measures of program exposure (respectively, the 2007 number of first-time homebuyers and the average number of first-time homebuyers over the period 2000 to 2007). *, **, and *** denote significance at 0.10, 0.05, and 0.01 levels, respectively.

	CBSA FE (1)	No Weights (2)	Ex Sand (3)	Trimmed (4)	$\overline{\text{Sales}} > \text{P10}$ (5)	Low p (6)	High p (7)
LHS is Long-Diff Sales							
Coefficient	0.021***	0.025***	0.019***	0.014***	0.021***	0.026***	−0.000
	(0.003)	(0.005)	(0.003)	(0.003)	(0.003)	(0.005)	(0.006)
Observations	8,692	8,692	6,664	7,834	7,825	2,431	2,375
R^2	0.446	0.464	0.362	0.445	0.451	0.625	0.434
LHS is Long-Diff Construction							
Coefficient	−0.002	0.006	−0.003	0.012**	−0.002	0.004	−0.003
	(0.008)	(0.008)	(0.010)	(0.006)	(0.008)	(0.022)	(0.018)
Observations	4,721	4,721	3,411	4,243	4,522	1,150	1,156
R^2	0.121	0.117	0.128	0.162	0.122	0.208	0.160
LHS is Long-Diff Foreclosures and Short Sales							
Coefficient	0.037	0.022	0.051*	0.028*	0.038	0.058*	−0.011
	(0.027)	(0.022)	(0.030)	(0.015)	(0.027)	(0.032)	(0.041)
Observations	8,506	8,506	6,478	7,663	7,705	2,392	2,316
R^2	0.343	0.319	0.233	0.433	0.349	0.433	0.272
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

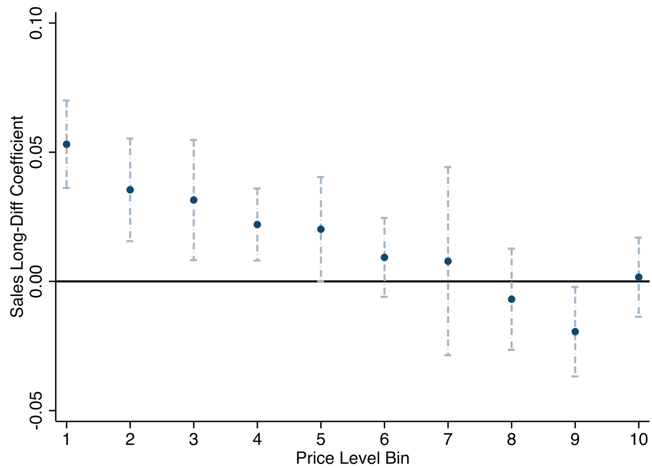

Figure 7. Heterogeneity by initial price level. This figure presents long-difference estimates of the average monthly effects of the FTHC on home sales for ZIP codes partitioned based on the level of median house prices during 2008. We run cross-sectional regressions, weighted by average monthly home sales in 2007, of the form

$$y_i = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i,$$

where y_i is the difference in average monthly nondistress home resales for the policy period versus the 17-month pre-period in area i. All regressions include CBSA fixed effects and X_i is a set of controls defined in Figure 5. Standard errors are clustered at the CBSA level. (Color figure can be viewed at wileyonlinelibrary.com)

presents regressions of the same form as those in Table III. We divide the home sales series into "starter" homes—those with one, two, or three bedrooms—and large homes—those with four or more bedrooms. We run the ZIP-level specifications separately for each series.[26] Estimates for the starter-home sample closely match those in our full sample, while those for larger homes are weakly negative and statistically insignificant. Thus, our main results are concentrated among starter homes, while larger homes show little response to the program.

D. New Construction

Our analysis thus far focuses on existing home sales. This category of transactions is the largest and most reliably recorded in the DataQuick database. Both of these features permit the high-frequency analysis that we use to validate our research design. However, in examining the policy as fiscal stimulus intended to spur GDP growth, existing home sales are not the ideal category to study, as they only contribute to output through transaction fees and complementary purchases.

[26] Because of incomplete reporting across places, the sample used here is a subset of the main sample where fewer than 5% of transactions between 2004 and 2013 have missing data on the number of bedrooms.

Table IV presents the effects of the program on new home sales, using new construction data recorded by DataQuick. To do so, we estimate a version of equation (4) with new home sales instead of existing home sales. As above, we seasonally adjust the new home sales series prior to averaging. The results indicate that the program had approximately no effect on new home sales. The point estimate is −0.002 and not statistically distinct from zero, as compared to 0.021 for existing home sales.

We confirm this finding in several robustness checks. In particular, we run unweighted regressions, we exclude sand states, and we use alternative samples or outcome-variable censoring to reduce the possible influence of outliers. All specifications suggest that the FTHC did not induce additional construction.

This finding is not surprising for a time when the national market suffered from significant overhang of recently built homes and elevated vacancies. Nevertheless, the result implies that the FTHC's stimulative effects through residential investment were likely of second-order importance, despite the substantial increase in existing home sales caused by the program.

V. The Effect of the FTHC on Homeownership

To supplement our earlier findings on the effect of the FTHC on home sales, we provide evidence from an RKD. This approach allows us to characterize the aggregate effect of the FTHC during the first year of the policy period, and to estimate the underlying response at the household level. To implement the RKD, we draw on administrative population-level tax data and exploit the phase-out range of the FTHC.

A. Policy Background, Data, and Estimation Method

The maximum value of the FTHC is $8,000 but it phases out for higher income taxpayers. In 2009, single (joint) taxpayers with AGI less than $75,000 ($150,000) were eligible for the maximum credit. The FTHC phased out linearly over the next $20,000, so that single (joint) taxpayers with an AGI of $95,000 ($170,000) or above were not eligible for the credit.[27] We refer to the point at which the FTHC begins to phase out as the "kink point" because the value of the FTHC has a kink, or slope change, at this point. Taxpayers below the kink can receive the maximum credit, while those above the kink are eligible for a smaller credit. The RKD exploits this kink point by relating the slope change in the FTHC to the slope change in the likelihood of being a first-time homebuyer.

To construct the sample for this analysis, we draw on population-level administrative tax data. These data provide a large sample size close to the kink point: our baseline sample includes over 3.8 million observations. A second advantage of these data is an accurate measure of income, so measurement error

[27] In 2010, the phase-out points increased to $125,000 for single taxpayers and $225,000 for joint taxpayers. We find no evidence that the credit increased homeownership at these higher income levels using an RKD. Because of smaller sample sizes, however, the estimates are not precise enough to exclude the estimates from the 2009 kink.

in exposure to the FTHC is negligible. We draw samples within ±\$8,000 of the kink point, located at \$75,000 (\$150,000) for single (joint) taxpayers. First-time homebuyers are taxpayers who pay mortgage interest in 2009, such that the primary taxpayer and the secondary taxpayer among joint filers did not pay interest in the prior three years.

We use a "sharp" RKD to identify the causal effect of the FTHC on the likelihood of being a first-time homebuyer. We estimate the intention-to-treat effect by imputing the value of the FTHC for all observations. The causal effect of the FTHC on being a first-time homebuyer equals the slope change in the probability of being a first-time homebuyer divided by the slope change in the imputed potential FTHC.

Following prior work (Card et al. (2015), Nielsen, Sørensen, and Taber (2010), Manoli and Turner (2018)), we focus on a constant effect additive model to examine the effect of the FTHC on homeownership at the household level:

$$\textit{First-Time Homebuyer}_i = \beta \textit{FTHC}_i + g(\textit{Distance}_i) + \varepsilon_i, \tag{5}$$

where $\textit{First-Time Homebuyer}_i$ ($\textit{FTHB}_i$) is an indicator equal to 1 if the household is a first-time homebuyer, $\textit{FTHC}_i$ is the value of the credit in hundreds of nominal dollars, and $\textit{Distance}_i$ is the distance in hundreds of nominal dollars to the kink point. The function g is a continuous function of kink distance, and the FTHC is assumed to be a continuous, deterministic function of kink distance with a slope change at zero. The average treatment effect is given by

$$\beta = \frac{\lim_{k\to 0^+} \frac{\partial E[\textit{FTHB}_i|\textit{Distance}_i=k]}{\partial k}|_{k=0} - \lim_{k\to 0^-} \frac{\partial E[\textit{FTHB}_i|\textit{Distance}_i=k]}{\partial k}|_{k=0}}{\lim_{k\to 0^+} \frac{\partial E[\textit{FTHC}|\textit{Distance}_i=k]}{\partial k}|_{k=0} - \lim_{k\to 0^-} \frac{\partial E[\textit{FTHC}|\textit{Distance}_i=k]}{\partial k}|_{k=0}}. \tag{6}$$

The numerator of this expression is the slope change around the kink in the probability of being a first-time homebuyer. The denominator is the slope change around the kink of the FTHC. We estimate the numerator using a regression of the form

$$\textit{First-Time Homebuyer}_i = \gamma \textit{Distance}_i + \delta D_i \textit{Distance}_i + \pi X_i + \nu_i, \tag{7}$$

where D_i is an indicator equal to 1 if the tax return has AGI above the kink point, and X_i is a vector of covariates including controls for AGI, ZIP code, and indicators for age and the number of dependent children. The sharp RKD uses the statutory slope change, so that the RKD estimator is given by $\beta = \frac{\delta}{-0.4}$, where -0.4 is the statutory slope change in the FTHC at the kink point.[28]

When estimating the regression in equation (7), we use a bandwidth of ±\$8,000. This bandwidth allows us to control for a nonlinear relationship between income and homeownership separate from the slope change at the kink point. We can separate income from kink distance because of the simultaneous existence of two kinks, one for single and one for joint filers. These kinks

[28] The FTHC is reduced from \$8,000 to \$0 over a \$20,000 interval of AGI, giving a slope of −0.4.

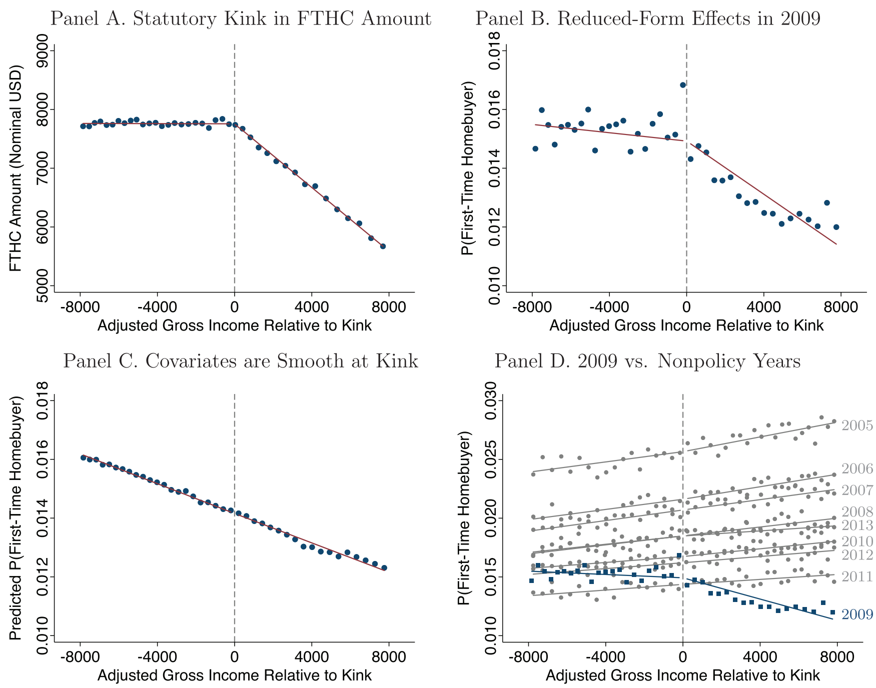

Figure 8. Regression kink analysis around income phase-out. The panels in this figure plot the effect of the FTHC on the first-time homebuying decision using the 2009 income phase-out of the credit. In each graph, we pool single and joint filers after recentering their AGI relative to their respective kinks and plot binned means for 50 equal density bins in a $8,000 bandwidth. Panel A plots the statutory kink for those who claim the credit. Panel B plots the reduced-form effect of the FTHC, where the outcome is an indicator for whether a household is a first-time homebuyer in 2009. Panel C plots predicted outcomes from a regression of the first-time homebuyer outcome on AGI and fixed effects for age, number of children, and ZIP code. The graph demonstrates continuity and smooth densities of the underlying covariates in one figure. Panel D provides reduced-form plots for all years from 2005 through 2013. The figure provides a visualization of an RK difference-in-differences design, in which the nonkink years (2005 through 2008 and 2010 through 2013) provide placebo tests. (Color figure can be viewed at wileyonlinelibrary.com)

also allow for placebo tests with filer groups placed at the alternative kink. We conduct further placebo tests using nonpolicy years.

B. Results

Figure 8 depicts results from the main specification and robustness checks. In each graph, we pool both single and joint filers after recentering their AGI relative to their respective kinks and plot binned means for 50 equal-density bins. Panel A plots the statutory kink for those who claim the credit, revealing the sharp kink induced by the phase-out region. The average level of credit to the left of the kink is slightly below $8,000 because some claimants are restricted by the purchase price of the house or apartment that they buy. Panel

B plots the reduced-form effect of the FTHC, revealing a kink in the propensity of homeownership that closely matches the kink location in the FTHC schedule.

Table V presents estimates of the reduced-form effect, which correspond to the paramater δ in equation (7). Given low baseline first-time homebuyer rates relative to the entire population close to the kink, we multiply the estimates by 10^5 for ease of readability. Standard errors are clustered at the ZIP code level. Columns (1) through (6) present regressions that vary the control set, which includes a linear control for AGI, a three-knot cubic spline in AGI, ZIP code fixed effects, primary taxpayer age fixed effects, and number of children fixed effects. The controls have essentially no effect on the estimated kink, which varies from -3.7 to -4.0 with t-statistics between 6 and 7.

Three additional tests confirm the robustness of this result. First, Figure 8, Panel C, plots predicted first-time homeowner rates as a function of income, age fixed effects, ZIP code fixed effects, and number of children fixed effects. This panel demonstrates the continuity and smoothness of the underlying densities of these covariates, which are necessary for the validity of the RKD. The predicted homeownership rates show no kink at the statutory kink.

Second, Figure 8, Panel D, provides reduced-form plots for all years from 2005 through 2013. The panel depicts an RK difference-in-differences design in which the nonkink years (2005 through 2008 and 2010 through 2013) provide placebo tests. As the housing cycle evolves, the average first-time homebuyer rate falls and recovers. Only in 2009, do we observe the pronounced kink in homebuying rates that coincides with the statutory kink. This panel rules out most alternative explanations for the observed kink in 2009.

The third robustness check is a within-2009 placebo test. We divide the sample into single and joint filers in 2009 and estimate placebo kinks for single filers at the joint filer kink and vice versa. Table V, columns (7) and (8), presents the single filer estimates and columns (9) and (10) present the joint filer estimates. For both groups, we find sharp and precise estimates at the true statutory kink and no evidence of a similar kink at the placebo. In summary, we find strong evidence of a causal effect of the FTHC on home purchases at the individual level, which is unlikely to be driven by confounding factors.

Our preferred estimate of -3.8×10^{-5} for \$100 in kink distance implies that the effect of the full \$8,000 of FTHC is an increase in first-time homebuyer propensity of 0.76 ($= (-3.8 \times 10^{-5})/(-0.4) \cdot (80)$) percentage points. This effect increases the within-sample baseline rate of 1.43 percentage points by 53%, which is economically significant. As is standard in discontinuity designs, such an estimate relies on extrapolation of the local treatment effect identified around the kink to larger changes in the credit amount. However, the local effect is identified under quite weak assumptions (DiNardo and Lee (2011)). We also report a separate estimate that restricts the treatment population to those filers who are eligible for the FTHC because they were not homeowners in the prior three years. Table V, column (11), reports this estimate, which equals -1.6×10^{-4} or 3.2% for the full credit (relative to a baseline rate of 5.4%). While not the primary focus of this paper, this microelasticity with respect to

Table V

The Effect of the FTHC on Homeownership

This table presents estimated effects of the FTHC on the first-time homebuying decision using the 2009 income phase-out of the credit. We report estimates of δ from regressions of the form

$$\textit{First-Time Homebuyer}_i = \gamma \textit{Kink Distance}_i + \delta D_i \textit{Kink Distance}_i + \pi X_i + \nu_i,$$

where *First-Time Homebuyer*$_i$ is an indicator equal to 1 if a household was a first-time homeowner in 2009, *Kink Distance*$_i$ is measured in hundreds of dollars, D_i is an indicator equal to 1 if the tax return has AGI above the kink point, and X_i is a vector of covariates including controls for AGI, ZIP code, and indicators for age and number of dependent children. Columns (1) through (6) include all single and joint filers within the $8,000 bandwidth around the filer groups' respective kinks and estimate the model with various combinations of controls. Columns (7) through (10) divide the sample into single and joint filers in 2009 and estimate placebo kinks for single filers at the joint filer kink and vice versa. Column (11) restricts the treatment population to those filers who are eligible for the FTHC because they were not homeowners in the prior three years. Coefficients are multiplied by 10^5 for readability. Standard errors are clustered at the ZIP code level to account for geographic covariance in the error term. *, **, and *** denote significance at 0.10, 0.05, and 0.01 levels, respectively. Table IA.IV in the Internet Appendix reports summary statistics for the estimation sample.

	All, Single, and Joint Filers						Single		Joint		
							Placebo	Real	Placebo	Real	Eligible
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Kink	−3.8***	−3.8***	−4.0***	−3.9***	−3.7***	−3.8***	10*	−6.8***	−0.5	−2.1***	−16***
	(0.53)	(0.53)	(0.53)	(0.53)	(0.54)	(0.53)	(4.9)	(1.2)	(0.36)	(0.44)	(2.0)
Observations (millions)	3.858	3.858	3.858	3.858	3.858	3.858	0.121	1.55	9.102	2.306	1.016
R^2	0.0001	0.008	0.008	0.014	0.019	0.025	0.13	0.029	0.023	0.02	0.047
Linear AGI	No	Yes	No	Yes	Yes	Yes	No	No	No	No	Yes
Spline AGI	No	No	Yes	No	No	No	No	No	No	No	No
Age and Kid FX	No	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
ZIP FX	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes

the credit is of independent interest for considering the microeconomic effects of housing market subsidies.

VI. The Effect of the FTHC on House Prices and Reallocation

A. House Prices

To explore the effect of the FTHC on house prices, we use data from the FHFA and CoreLogic. These data sets rely on a repeat sales methodology to estimate price indices at the ZIP level. The FHFA indices use all mortgages guaranteed by Fannie Mae and Freddie Mac and offer the greatest geographic coverage, but they are only estimated annually. CoreLogic's index uses its proprietary database to estimate monthly values. We augment this analysis by using microdata from DataQuick to explore the effect of the policy using raw, transaction-level prices.

We follow the same empirical strategy as for the home sales regressions in Table IV, exploiting within-CBSA variation in program exposure. For the FHFA data, the left-hand-side variable is the cumulative annual log price difference during 2009 and 2010 minus the cumulative annual log price difference during 2007 and 2008. We use a long-difference specification because, unlike housing transactions, month-to-month changes in house price indices at the ZIP level are quite noisy. We present estimates for both raw changes in price growth and for market-adjusted changes.[29] This adjustment allows us to control for differential exposure of high exposure ZIPs to the national cycle driven by higher risk in these areas. For the CoreLogic data, the left-hand-side variable is the raw cumulative monthly log price difference during the policy period minus the cumulative monthly log price differences during the 17-month preperiod. In all cases, we multiply the left-hand side by 100, so the treatment effect units are percentage points of growth per standard deviation change in program exposure.

Table VI presents results from these regressions. In our preferred specification, which uses the market-adjusted FHFA series, we find that the program led to an increase in cumulative price growth of 1.1% per standard deviation increase in exposure. At the median initial price level of \$222,000 in our sample, this implies a price increase of \$2,400 ($\approx 0.011 \times 222{,}000$).[30] This figure is plausible given the credit size of \$8,000 and considerable excess inventory in the market. This result also implies that even the highest exposure areas did not see house prices increase by more than the credit.[31]

[29] In the case of market-adjusted changes, we first estimate ZIP-specific housing market betas in the 10-year window from 1997 to 2006 and then subtract beta times the market return to compute a ZIP-level excess return. For the market return, we use the national annual FHFA house price index, which is estimated using a similar methodology as for the ZIP-level indices.

[30] Alternatively, after standardizing exposure, the 10^{th} percentile is 1.92 and the 90^{th} is 4.15. This implies that going from the 10^{th} to the 90^{th} percentile increases exposure by 2.23 standard deviations and increases house price growth by 2.4%, or \$5,331 at the median house price.

[31] House prices were falling on average during this time, so these effects may be interpreted as indicating that the program slowed the rate of price declines.

Table VI
The Effect of the FTHC on House Prices

This table presents regressions of the cumulative effects of the FTHC on different measures of house price growth. We run cross-sectional regressions, weighted by average monthly home sales in 2007, of the form

$$y_i = \alpha_{CBSA} + \beta Exposure_i + \gamma X_i + \varepsilon_i,$$

where y_i is a housing market outcome in area i over the relevant time period. In the first row, the outcome is the market-adjusted, cumulative annual log price difference from FHFA price index data during 2009 and 2010 minus the cumulative annual log price difference during 2007 and 2008. In the second row, the outcome is the unadjusted version of the price series from the first row. In the third row, the outcome is the raw cumulative monthly log price difference from CoreLogic during the policy period minus the cumulative monthly log price difference during the 17-month pre-period. In all cases, we multiply the left-hand side by 100, so the treatment effect units are percentage points of growth per one-standard-deviation change in program exposure. All series are seasonally adjusted prior to aggregation. Exposure is normalized by its cross-sectional standard deviation. Each column presents estimates based on specifications and samples as defined in Table IV. Standard errors are clustered at the CBSA level. *, **, and *** denote significance at 0.10, 0.05, and 0.01 levels, respectively.

	CBSA FE (1)	No Weights (2)	Ex Sand (3)	Trimmed (4)	$\overline{\text{Sales}}$ > P10 (5)	Low p (6)	High p (7)
	LHS is Long-Diff Price Growth (FHFA Mkt Adjusted)						
Coefficient	1.077***	1.149***	0.782***	0.685***	1.087***	1.073***	0.628**
	(0.183)	(0.222)	(0.201)	(0.125)	(0.184)	(0.297)	(0.278)
Observations	8,363	8,363	6,388	7,523	7,570	2,304	2,261
R^2	0.615	0.580	0.632	0.623	0.620	0.617	0.653
	LHS is Long-Diff Price Growth (FHFA Raw)						
Coefficient	1.116***	1.190***	0.802***	0.714***	1.127***	1.107***	0.689**
	(0.193)	(0.230)	(0.213)	(0.126)	(0.194)	(0.304)	(0.289)
Observations	8,363	8,363	6,388	7,523	7,570	2,304	2,261
R^2	0.619	0.585	0.633	0.619	0.623	0.610	0.669
	LHS is Long-Diff Price Growth (CoreLogic)						
Coefficient	0.581***	0.597***	0.575**	0.409**	0.586***	1.200***	0.245
	(0.190)	(0.194)	(0.242)	(0.186)	(0.193)	(0.296)	(0.465)
Observations	5,748	5,748	4,074	5,175	5,607	1,336	1,666
R^2	0.680	0.714	0.483	0.632	0.675	0.701	0.708
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CBSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

The estimate above is robust to different weights, alternative sample definitions, and censoring of the left-hand-side variable. In addition, the estimates vary little between the FHFA and CoreLogic samples and do not depend on the market adjustment for ZIP-specific cyclicality. As with the quantity results, we estimate more precise and qualitatively larger effects in the ZIPs with lower

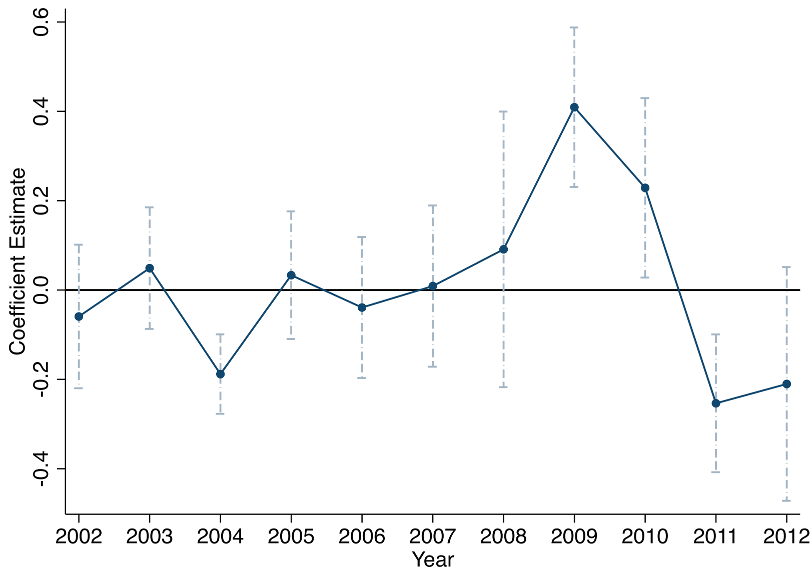

Figure 9. The effect of the FTHC on house prices. The figure plots coefficients for yearly house price growth regressions from market-adjusted house price indices at the ZIP level from FHFA. We run year-by-year regressions, weighted by total home sales in 2007, of the form

$$\Delta\tilde{r}_i = \alpha_{\text{CBSA}} + \beta Exposure_i + \gamma X_i + \varepsilon_i,$$

where $\tilde{r}_i$ is the first difference in market-adjusted house price growth in area i, and X_i is a set of controls as defined in Figure 5. Program exposure is normalized by its cross-sectional standard deviation. The left-hand side is multiplied by 100, so the treatment effects are percentage points of growth per standard deviation in program exposure. (Color figure can be viewed at wileyonlinelibrary.com)

initial house prices, although these differences are statistically significant only in the CoreLogic sample.[32]

Figure 9 allows us to explore the extent to which the price effects reverse in the years following the policy. Each year, we estimate cross-sectional regressions of first differences in market-adjusted price growth from the FHFA. The coefficients show no pre-trends in years prior to the program and strong trend breaks during the two program years, which match both qualitatively and quantitatively the positive long-difference effects in Table VI. In the year immediately following the program, price growth retreats somewhat, undoing approximately one-quarter of the increase caused by the program. This evidence is consistent with an incomplete reversal of the home sales response in the post-policy period.

[32] Section II in the Internet Appendix presents additional results using aggregate price data and a hedonic index constructed from our DataQuick data. That appendix also examines how repeat sales indices may obscure the effect of the program on aggregate prices, either by smoothing sharp changes in the time series or by excluding transactions more likely to see significant price increases during this time.

These results support a key rationale for stabilizing the housing market, which is to address the pecuniary externality that distressed sales impose on nearby homeowners by depressing house prices.[33] A related spillover concerns how house price changes affect the real economy via a balance sheet channel. Considerable evidence suggests that these effects were large during the Great Recession, as falling house prices led to a large decline in consumption and employment during the Great Recession (Mian and Sufi (2011, 2014)).[34]

B. Reallocation

In this section, we investigate the extent to which the FTHC helped reallocate underutilized homes to financially constrained buyers. The goal is to provide suggestive evidence of one mechanism through which the FTHC helped stabilize the housing market, namely, by mitigating a credit market failure due to the simultaneous presence of constrained buyers (as many natural buyers could not borrow because of the weak economy and credit disruptions) and elevated vacancies (as many homes had already been built yet remained unoccupied).[35] Section II in the Internet Appendix describes these calculations in detail.

B.1. Underutilized Homes on the Supply Side

In the aggregate DataQuick data, 24% of home sales during the policy period were from developers or builders, two-thirds of which were not new construction. Instead, these sales were from unsold inventories of homes built during the housing boom. Furthermore, 33% of home sales were either distressed sales, including both short sales and foreclosure auctions, or sales to reduce financial institutions' housing portfolios. Thus, many of the transacted homes did not involve transfers from one homeowner to another, but instead transitions of vacant homes to more productive use.

Additional evidence comes from de-identified tax returns of FTHC claimants. From the 2009 claimants, we find that 42% move into an address that had no tax filers in 2007, and 33% transition into a single-tax-filer address from a multiple-filer address in 2007. The data further suggest that FTHC claimants are more likely to form new households relative to first-time buyers in other

[33] Campbell, Giglio, and Pathak (2011) show that prices of houses within 0.05 miles of a foreclosure decline by about 1%. Similarly, Whitaker and Fitzpatrick IV (2013) find that an additional property within 500 feet that is vacant or delinquent reduces a home's sale price by 1% to 2%. Guren and McQuade (2015) show in a quantitative general equilibrium model that these effects can be large.

[34] Falling house prices lead to financial accelerator effects by reducing household net worth, which affects in turn whether firms can borrow to invest and whether households can borrow to consume (Kiyotaki and Moore (1997); Iacoviello (2005)). House prices can also affect bank balance sheets as losses realized by banks reduce their ability to borrow and lend (Shleifer and Vishny (2010b)).

[35] The fact that vacant homes depreciate faster due to lack of maintenance (Gerardi et al. (2015)) and may enable crime (Ellen, Lacoe, and Sharygin (2013), Cui and Walsh (2015)) further strengthens this case.

years, and that they disproportionately purchase recently vacated homes as opposed to new construction.

In the last row of Table IV, we focus on foreclosures and short sales and use program exposure to study the effect of the program on distressed sales within CBSAs. The point estimate is 0.037 and marginally statistically significant in some specifications. The magnitude suggests that the program induced a modest shift in the composition of sales toward distressed sales. Columns (2) through (7) confirm the robustness of this finding, as do specifications in Tables IA.V and IA.VI in the Internet Appendix using alternative exposure measures. We note that this effect is in addition to the level effect on distressed transactions due to the increase in demand caused by the program.

B.2. Financially Constrained Buyers on the Demand Side

According to HUD, the FHA supported 781,000 first-time homebuyers during 2009 and 882,000 during 2010, or approximately 56% of the first-time buyer market during these years.[36] FHA buyers receive low down payment loans in exchange for higher subsequent interest payments plus required mortgage insurance premiums. Over the first five years of an FHA-insured loan, households pay 73% more in interest and insurance premiums relative to interest on a conventional mortgage with a 20% down payment, which amounts to an extra \$27,000 for a \$200,000 home. That so many households chose an FHA mortgage despite the higher future cost suggests that down payment constraints were likely binding. In line with this claim, the majority of FTHC claimants used amended returns to accelerate receipt of the credit toward the purchase date, with younger buyers using this mechanism more aggressively than older buyers.

Given the high origination LTV ratios of policy-period homebuyers and evidence suggesting such LTVs can lead to distress, it is important to ask what happened to these buyers in the post-period. Figure 10 plots cumulative distress cohorts for purchases made during the policy period and compares these to cohorts based on 2006, 2007, and 2008 sales as well as cohorts based on 2011 sales. Both the 2009 and 2010 policy cohorts show no difference in default rates relative to the 2011 post-policy cohort. Furthermore, all three of these groups display considerably lower rates of transition into distressed sales than the pre-policy groups. Note that because these buyers purchased at the bottom, the path of aggregate prices implies that even riskier borrowers will have lower default rates post-crisis (Palmer (2015)). Nevertheless, the data do not indicate that the FTHC program drew unusually risky buyers into the market, despite the very high LTVs at which these buyers entered. In this sense, the reallocation of homes appears to have been stable.

[36] See Figure 6 in HUD's "Annual Report to Congress" (Department of Housing and Urban Development (2011)).

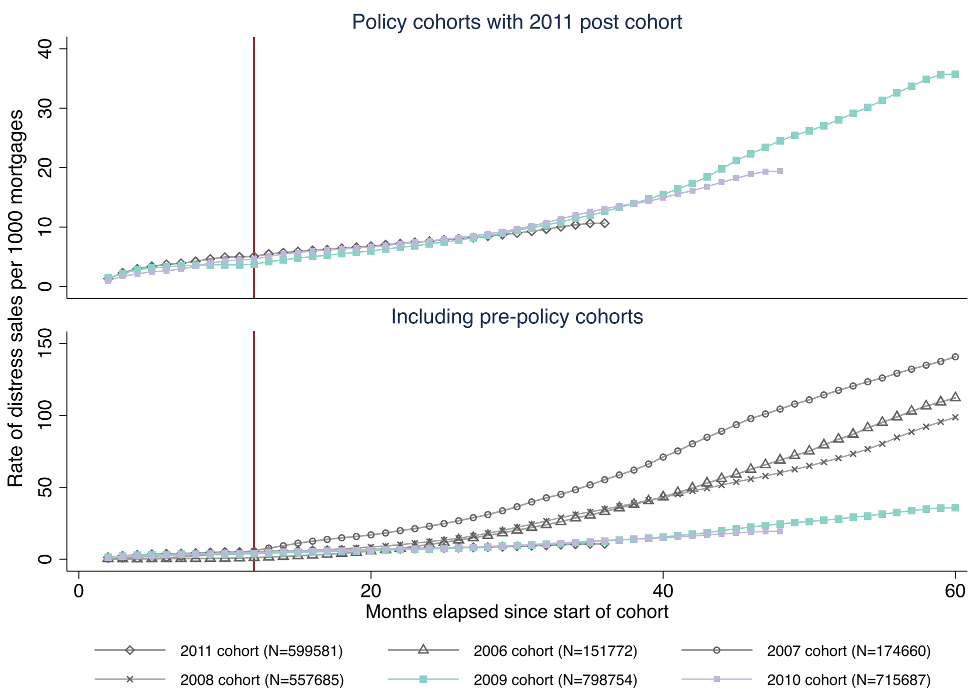

Figure 10. Default rates for policy period FHA buyers versus other cohorts. The figure plots cumulative distress cohorts for FHA-insured purchases made during the policy period and compares these to cohorts based on 2006, 2007, 2008, and 2011 sales. We measure transitions into distress using DataQuick by following properties purchased in a given year and computing the share of properties that become distressed sales. (Color figure can be viewed at wileyonlinelibrary.com)

VII. The Aggregate Effect of the FTHC

A. The Aggregate Effect on Home Sales

An important goal of the FTHC policy was to stimulate real economic activity, particularly in the housing sector. For new home sales, the effect can be seen in the value of the transacted homes; we have shown that this effect was likely small. For existing home sales, the effect can be seen in complementary purchases and transaction fees. To estimate this effect, we first measure the total number of existing sales induced by the program. We do so using three different methodologies based on the three quasi-experimental research designs in Sections IV and V. Reassuringly, the three approaches yield consistent results (see Section III in the Internet Appendix for details).

Table VII summarizes the results. The cross-sectional approach exploits differences in cross-sectional exposure and constructs an estimate of the aggregate effect using the group receiving the smallest shock as a counterfactual. We estimate that the FTHC increased existing home sales by 169,000 units

Table VII

The Aggregate Effect of the FTHC under Alternative Estimation Methods

This table summarizes our estimates of the aggregate effect of the FTHC program. The aggregate home sales effect is based on different research designs and aggregation assumptions. Direct multipliers estimate GDP effects from the increase in home sales via realtor fees and complementary purchases. Indirect multipliers estimate GDP effects from the response of consumption to increasing housing prices, either for the starter (S) home market or for all (A) homes. Multipliers in the cross-sectional approach are relative to the in-sample cost of FTHC claims. Multipliers in the regression kink approach are relative to the total cost of FTHC claims. See Section III in the Internet Appendix for details.

			Multipliers		
Estimation Method	Aggregate (%)	Sample	Direct	Indirect (S)	Indirect (A)
Cross-sectional					
A. Relative to bottom 1%	412K (8.2%)	169K	0.179	0.302	0.525
B. Relative to zero	568K (11.3%)	233K	0.241	0.427	0.733
C. Average of A and B	490K (9.8%)	201K	0.208	0.354	0.629
Regression kink					
Full population	520K (10.3%)	n.a.	0.371	n.a.	n.a.
Eligible only	610K (12.1%)	n.a.	0.435	n.a.	n.a.
Age distribution					
Young buyers	n.a.	261K	n.a.	n.a.	n.a.

within-sample during the policy period, which represents 8.1% of all sales during this period. In 2007, our sample covers 41% of the national existing home sales market. Extrapolating our estimates to the national market yields an estimated increase of approximately 412,000 units during the policy period. A less conservative approach that aggregates estimates relative to a zero-exposure baseline yields an estimate of 233,000 units within-sample, which amounts to 11.2% of all sales during the policy period. The extrapolated aggregate estimate in this case is 568,000 units. Our preferred estimate equals 490,000 units, the average of these alternative approaches.

To produce this aggregate estimate, we have not modeled general equilibrium effects, which are subsumed into time fixed effects. Regarding concerns about general equilibrium effects, it is comforting that aggregate home sales clearly demonstrate a policy response. Furthermore, the heatmap does not indicate that home sales fell below pre-policy levels in low-exposure areas, which would be predicted by binding aggregate resource constraints. In addition, because the policy was implemented with interest rates at the zero lower bound, any mitigating effect from rising rates was likely small. Nevertheless, without a full model, our aggregate estimate should be considered an imperfect approximation of the total effect.

A related concern is that spillovers between treatment and control ZIPs might bias our aggregate calculation. This could happen if a purchase in a high-exposure ZIP triggers a simultaneous purchase by the previous seller in a low-exposure ZIP. Three features of our analysis suggest that spillovers induced by such "real estate chains" are likely not quantitatively important. First, the inclusion of CBSA fixed effects means that only transactions that take place within a given city would bias our estimates. Second, many sellers were banks or developers, sellers for which these spillover effects are irrelevant. Finally, the most plausible sign of the bias is negative since the second transaction would artificially inflate sales in lower exposure ZIPs. Our aggregate estimates are therefore likely conservative relative to the true aggregate effect.

We obtain a second estimate of the total size of the FTHC program using the microlevel estimates from our RKD. Our preferred specification implies that the full FTHC increased the first-time homebuyer propensity by 0.76% among all tax filers (69 million total) and 3.2% among eligible tax filers (19 million total). These estimates imply that 520,000 or 610,000 households, respectively, were induced by the FTHC to purchase homes. Although the cross-sectional and RKD methods use different sources of identification and different approaches to aggregation, the final estimates are quite close. This fact lends further credibility to each research design and our aggregate estimates.

A third estimate of the aggregate effect of the FTHC comes from using variation embedded in the age distribution of first-time buyers, which shifts substantially toward young buyers in 2009. Under the assumption that older buyers are unaffected by the program, this procedure yields an aggregate estimate of 261,000 induced buyers under 35, or approximately 22% of first-time buyers in our data in this age group. This estimate provides a lower bound on the overall estimate as the assumption that older buyers are unaffected is too strong.

B. The Direct and Indirect Stimulative Effects

We now use our aggregate estimates to provide a back-of-the-envelope calculation for the program's direct and indirect fiscal impact. Section III in the Internet Appendix contains additional details.

There are two sources of direct effects. The first is the income generated for realtors. On average, the realtor fee ranges from 5% to 5.5% of the purchase price, so we use 5.25% in our calculations. The second is the complementary furniture, home improvement, and related expenditures at the time of a new purchase. We draw on the literature that estimates these expenditures to be between 1.9% (Benmelech, Guren, and Melzer (2017)) and 2.8% (Best and Kleven (2017)) of the purchase price and use the average of 2.35%.

Table VII presents the results. Under a variety of assumptions, the overall direct contribution to GDP of the FTHC within our sample is \$2.5 billion using our conservative aggregate estimate and \$3.3 billion using our more aggressive estimate, both of which are significantly below the program's within-sample \$14.0 billion cost.[37] These estimates imply that the FTHC generated fiscal impact multipliers of 0.18 in the conservative case and 0.24 in the aggressive case,

[37] The total cost of the FTHC program was \$20.3 billion and our sample contains 69% of the FTHC claims.

both much smaller than the typical tax multiplier during recessions. Taken together, these calculations lead us to conclude that the direct GDP effects of the program were relatively modest and significantly below its overall cost.

The policy also generated indirect stimulative effects by increasing house values. A recent empirical literature documents large causal responses of nondurable consumption to house price movements. Because house price appreciation affects existing homeowners, it is possible that these indirect effects are large, as housing wealth is the largest component of net worth for most households.

To derive a ballpark estimate of this effect, we apply the rule-of-thumb formula in Berger et al. (2017) to estimate the aggregate impact of our cross-sectional house price results. In contrast to the program's modest direct effects, we estimate that the indirect contribution to GDP is \$4.2 billion using a conservative baseline and \$6.0 billion under a more aggressive baseline. To calculate this estimate, we conservatively assume that the marginal propensity to consume (MPC) out of housing wealth is five cents for households with LTVs below 80%, as they have enough equity to borrow against even in tight credit markets, and zero for households with little or negative equity.[38] We further assume that only the housing stock for starter houses (one to three bedrooms) benefits from the price effects.

Despite these conservative assumptions, we nevertheless find that the program's indirect effects likely exceeded the direct effects. Our estimate implies an indirect effect multiplier of 0.30 in the conservative case and 0.43 in the aggressive case, a range that is significantly larger than the direct effect. Were we to apply the indirect effect to the entire housing stock, these multipliers would rise to 0.53 and 0.73, respectively. Combining both yields a total impact multiplier of the FTHC of 0.48 in the conservative case and 0.67 in the aggressive case, much of which comes from the indirect effect—the component usually ignored in policy evaluations.

VIII. Conclusion

In this paper, we examine whether policy can accelerate the reallocation process by spurring demand for housing in times of market weakness. We study temporary tax incentives targeted at marginal buyers in the housing market. Unlike debt renegotiation programs, financial market support, and fiscal and monetary stimulus, the policy we study directly targeted the overhang of distressed and vacant homes while aiming to keep them in private hands.

The program proved effective at spurring home sales, and these effects did not immediately reverse once the program ended—while the research design does not permit us to determine whether there is a permanent component to the increase in home sales, we show that these sales were not reversed over at least

[38] As a point of comparison, the average MPC out of housing wealth of the nine low-LTV empirical studies discussed in Ganong and Noel (2018) (see Figure 3) is 7.6 cents, and the average MPC out of housing wealth of the studies discussed in Carroll, Otsuka, and Slacalek (2011) is 5.5 cents.

the two years following the program. In addition, house prices increased, and buyers induced by the program were not more likely to default than previous or subsequent cohorts of buyers. This stable demand shock to the market likely accelerated the reallocation of vacant homes from the portfolios of institutional investors and banks and from the unsold inventories of home builders into the hands of higher value and possibly constrained first-time homebuyers.

Because the increase in housing transactions occurred largely in the existing home market, the effect on GDP was likely below the cost of the program. Yet by stabilizing house prices and thereby increasing housing wealth, the policy produced indirect effects that rivaled and likely surpassed the program's direct stimulus effects. This feature made the FTHC complementary to other principal and payment renegotiation programs, such as HAMP and HARP, that aimed to repair household balance sheets and improve mortgage affordability. A key difference is that the FTHC addressed a later stage in the foreclosure chain, namely, when the house is owned by the bank.

The policy also stimulated homeownership. This is notable because the U.S. government spends at least $70 billion a year on the mortgage interest deduction, in part to encourage homeownership. While the mortgage interest deduction may induce some marginal households into homeownership, it also induces households already planning to buy a home into buying larger homes, which has limited social benefits (Glaeser and Shapiro (2003), Kirker, Floetotto, and Stroebel (2016)). One lesson from the FTHC is that, if increasing homeownership rates is a policy goal, then directly targeting potential homeowners and the constraints they face may be a more cost-effective way to achieve this goal. More research into this question is needed.

Initial submission: March 14, 2018; Accepted: March 17, 2019
Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

REFERENCES

Adda, Jérôme, and Russell Cooper, 2000, Balladurette and juppette: A discrete analysis of scrapping subsidies, *Journal of Political Economy* 108, 778–806.

Agarwal, Sumit, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, Tomasz Piskorski, and Amit Seru, 2017a, Policy intervention in debt renegotiation: Evidence from the home affordable modification program, *Journal of Political Economy* 125, 654–712.

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tomasz Piskorski, Amit Seru, and Vincent Yao, 2017b, Mortgage refinancing, consumer spending, and competition: Evidence from the home affordable refinancing program, NBER Working Paper No. 21512.

Benmelech, Efraim, Adam Guren, and Brian T. Melzer, 2017, Making the house a home: The stimulative effect of home purchases on consumption and investment, NBER Working Paper No. 23570.

Berger, David, Tianfang Cui, Nicholas Turner, and Eric Zwick, 2018, Stimulating durable purchases: Theory and evidence, University of Chicago, Working paper.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, 2017, House prices and consumer spending, *Review of Economic Studies* 85, 1502–1542.

Berger, David, and Joseph Vavra, 2015, Consumption dynamics during recessions, *Econometrica* 83, 101–154.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, 2004, How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* 119, 249–275.

Carlos Best, Michael, and Henrik Jacobsen Kleven, 2017, Housing market responses to transaction taxes: Evidence from notches and stimulus in the uk, *Review of Economic Studies* 85, 157–193.

Bogin, Alexander N., William M. Doerner, and William D. Larson, 2016, Local house price dynamics: New indices and stylized facts, FHFA Working Paper No. 16-01.

Brogaard, Jonathan, and Kevin Roshak, 2011, The effectiveness of the 2008-2010 housing tax credit, SSRN Working Paper No. 1882599.

Campbell, John, Stefano Giglio, and Parag Pathak, 2011, Forced sales and house prices, *American Economic Review* 101, 2108–2131.

Card, David, David S. Lee, Zhuan Pei, and Andrea Weber, 2015, Inference on causal effects in a generalized regression kink design, *Econometrica* 83, 2453–2483.

Carroll, Christopher D., Misuzu Otsuka, and Jiri Slacalek, 2011, How large are housing and financial wealth effects? A new approach, *Journal of Money, Credit and Banking* 43, 55–79.

Chodorow-Reich, Gabriel, Laura Feiveson, Zachary Liscow, and William Gui Woolston, 2012, Does state fiscal relief during recessions increase employment? Evidence from the American Recovery and Reinvestment Act, *American Economic Journal: Economic Policy* 4, 118–145.

Cui, Lin, and Randall Walsh, 2015, Foreclosure, vacancy and crime, *Journal of Urban Economics* 87, 72–84.

Department of Housing and Urban Development, 2011, Annual Report to Congress Regarding the Financial Status of the FHA Mutual Mortgage Insurance Fund Fiscal Year 2011, http://portal.hud.gov/hudportal/documents/huddoc?id=FHAMMIFundAnnRptFY11No2.pdf. Accessed June 15, 2016.

Diamond, Douglas, and Raghuram Rajan, 2011, Fear of fire sales, illiquidity seeking, and credit freezes, *Quarterly Journal of Economics* 126, 557–591.

DiNardo, John, and David Lee, 2011, Program evaluation and research designs, in Orley Ashenfelter, and David Card, eds.: *Handbook of Labor Economics* (Elsevier Inc., Amsterdam).

Dynan, Karen, Ted Gayer, and Natasha Plotkin, 2013, *An Evaluation of Federal and State Homebuyer Tax Incentives* (The Brookings Institution, Washington, DC).

Eberly, Janice, and Arvind Krishnamurthy, 2014, Efficient credit policies in a housing debt crisis, *Brookings Papers on Economic Activity* 2014, 73–136.

Eggertsson, Gauti, and Paul Krugman, 2012, Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach, *Quarterly Journal of Economics* 127, 1469–1513.

Eisfeldt, Andrea, and Adriano Rampini, 2006, Capital reallocation and liquidity, *Journal of Monetary Economics* 53, 369–399.

Ellen, Ingrid Gould, Johanna Lacoe, and Claudia Ayanna Sharygin, 2013, Do foreclosures cause crime? *Journal of Urban Economics* 74, 59–70.

Fisher, Irving, 1933, The debt-deflation theory of great depressions, *Econometrica* 1, 337–357.

French, Kenneth, Martin Baily, John Campbell, John Cochrane, Douglas Diamond, Darrell Duffie, Anil Kashyap, Frederic Mishkin, Raghuram Rajan, and David Scharfstein, 2010, *The Squam Lake Report: Fixing the Financial System* (Princeton University Press, Princeton, NJ).

Ganong, Peter, and Pascal Noel, 2018, Liquidity vs. wealth in household debt obligations: Evidence from housing policy in the Great Recession, NBER Working Paper No. 24964.

Gerardi, Kristopher, Eric Rosenblatt, Paul S. Willen, and Vincent Yao, 2015, Foreclosure externalities: New evidence, *Journal of Urban Economics* 87, 42–56.

Glaeser, Edward L., and Jesse M. Shapiro, 2003, The benefits of the home mortgage interest deduction, *Tax Policy and the Economy* 17, 37–82.

Goodman, Laurie, Ellen Seidman, and Jun Zhu, 2014, *FHA Loan Limits: What Areas are the Most Affected?* (Urban Institute, Washington, DC).

Green, Daniel, Brian T. Melzer, Jonathan A. Parker, and Arcenis Rojas, 2016, Accelerator or brake? Cash for Clunkers, household liquidity, and aggregate demand, NBER Working Paper No. 22878.

Guren, Adam M., and Timothy J. McQuade, 2015, How do foreclosures exacerbate housing downturns? Working paper, Harvard.

Hanson, Samuel, Anil Kashyap, and Jeremy Stein, 2011, A macroprudential approach to financial regulation, *Journal of Economic Perspectives* 25, 3–28.

von Hayek, Friedrich, 1931, The "paradox" of saving, *Economica* 32, 125–169.

Hembre, Erik, 2018, An examination of the First-Time Homebuyer Tax Credit, *Regional Science and Urban Economics* 73, 196–216.

House, Christopher, and Matthew Shapiro, 2008, Temporary investment tax incentives: Theory with evidence from bonus depreciation, *American Economic Review* 98, 737–768.

Iacoviello, Matteo, 2005, House prices, borrowing constraints and monetary policy in the business cycle, *American Economic Review* 95, 739–764.

Maynard Keynes, John, 1936, *The General Theory of Employment, Interest, and Money* (Macmillan, London).

Kirker, Michael, Max Floetotto, and Johannes Stroebel, 2016, Government intervention in the housing market: Who wins, who loses? *Journal of Monetary Economics* 80, 106–123.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, 211–248.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider, 2015, The housing market(s) of San Diego, *American Economic Review* 105, 1371–1407.

Lorenzoni, Guido, 2008, Inefficient credit booms, *Review of Economic Studies* 75, 809–833.

Manoli, Day, and Nicholas Turner, 2018, Cash-on-hand and college enrollment: Evidence from population tax data and the earned income tax credit, *American Economic Journal: Economic Policy* 10, 242–271.

Mian, Atif, and Amir Sufi, 2009, The consequences of mortgage credit expansion: Evidence from the U.S. mortgage default crisis, *Quarterly Journal of Economics* 124, 1449–1496.

Mian, Atif, and Amir Sufi, 2011, House prices, home equity-based borrowing, and the U.S. household leverage crisis, *American Economic Review* 101, 2132–2156.

Mian, Atif, and Amir Sufi, 2012, The effects of fiscal stimulus: Evidence from the 2009 Cash for Clunkers program, *Quarterly Journal of Economics* 127, 1107–1142.

Mian, Atif, and Amir Sufi, 2014, What explains the 2007-2009 drop in employment? *Econometrica* 82, 2197–2223.

Mian, Atif, and Amir Sufi, 2015, *House of Debt: How They (and You) Caused the Great Recession, and How We Can Prevent it from Happening Again* (University of Chicago Press, Chicago, IL).

Mian, Atif, Amir Sufi, and Francesco Trebbi, 2015, Foreclosures, house prices, and the real economy, *Journal of Finance* 70, 2587–2634.

National Association of Realtors, 2017, Copyright ©2016 Profile of Home Buyers and Sellers, Technical report. https://www.nar.realtor/reports/highlights-from-the-profile-of-home-buyers-and-sellers.

Skyt Nielsen, Helena, Torben Sørensen, and Christopher Taber, 2010, Estimating the effect of student aid on college enrollment: Evidence from a government grant policy reform, *American Economic Journal: Economic Policy* 2, 185–215.

Palmer, Christopher, 2015, Why did so many subprime borrowers default during the crisis: Loose credit or plummeting prices? Working paper, MIT.

Ramey, Valerie, and Matthew Shapiro, 2001, Displaced capital: A study of aerospace plant closings, *Journal of Political Economy* 109, 958–992.

Rognlie, Matthew, Andrei Shleifer, and Alp Simsek, 2018, Investment hangover and the Great Recession, *American Economic Journal: Macroeconomics* 10, 113–153.

Shleifer, Andrei, and Robert Vishny, 1992, Liquidation values and debt capacity: A market equilibrium approach, *Journal of Finance* 47, 1343–1366.

Shleifer, Andrei, and Robert Vishny, 2010a, Asset fire sales and credit easing, *American Economic Review: Papers and Proceedings* 100, 46–50.

Shleifer, Andrei, and Robert Vishny, 2010b, Unstable banking, *Journal of Financial Economics* 97, 306–318.

Whitaker, Stephen, and Thomas J. Fitzpatrick IV, 2013, Deconstructing distressed-property spillovers: The effects of vacant, tax-delinquent, and foreclosed properties in housing submarkets, *Journal of Housing Economics* 22, 79–91.

Zwick, Eric, and James Mahon, 2017, Tax policy and heterogeneous investment behavior, *American Economic Review* 107, 217–248.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.
Replication code.