Cracking Down, Pricing Up: Housing Supply in the Wake of Mass Deportation*

Troup Howard† Mengqi Wang‡ Dayin Zhang§

October 2024

Abstract

US housing markets have faced a secular shortage of housing supply in the past decade, contributing to a steady decline in housing affordability. Most supply-side explanations in the literature have tended to focus on the distortionary effect of local housing regulations. This paper provides novel evidence on the interplay between residential construction, labor supplied to the construction industry, and immigration policy. We exploit the staggered rollout of a national increase in immigration enforcement to identify negative shocks to construction sector employment that are likely unrelated to local housing market conditions. Treated counties experience large and persistent reductions in construction workforce, residential homebuilding, and increases in home prices. Further, evidence suggests that undocumented labor is a complement to domestic labor: an indirect outcome of deporting undocumented construction workers is net job loss for US-born workers, especially in higher-skilled occupations. We find that any demand-side downward pressure on home prices linked to increased deportations is temporary and quickly dominated by the supply-side impact.

JEL Codes: R31, J60

*We would like to thank Darren Aiello, Jason Cook, Chloe East, Andra Ghent, Lu Han, Elena Patel, Adam Looney, Alvin Murphy, Christopher Palmer, Albert Saiz, Christopher Timmins, and Matthew Turner; as well as seminar participants at the AREUEA National Conference and AREUEA Virtual Seminar, BYU, Conference on Frictions in Real Estate and Infrastructure Investment, Federal Reserve Board of Governors, Federal Reserve Bank of Chicago, University of Colorado at Boulder, the Ivory-Boyer Innovations in Housing Affordability Summit, University of Utah, and the Pre-WFA Real Estate Symposium for their helpful comments. All remaining errors are our own. This paper was previously circulated under the title *How Do Labor Shortages Affect Residential Construction and Housing Affordability?*

†University of Utah. Email: troup.howard@eccles.utah.edu.

‡Amherst College. Email: mewang@amherst.edu.

§University of Wisconsin-Madison. Email: dayin.zhang@wisc.edu.

1 Introduction

General consensus holds that historical declines in housing affordability over the past decade are rooted in a secular decline in housing supply (Glaeser and Gyourko 2018). Since 2011 the US has added an average of 1.1 million new housing units every year (see Figure 1). This is 30% lower than the long-run equilibrium before the Great Recession and 34% lower than the annual new construction demand (Khater et al. 2018). An extensive academic literature exploring supply-side drivers of low homebuilding has tended to focus on the distortionary effects of stringent zoning codes and local building covenants (Molloy 2020). However, little attention has been paid to the most important production factor in construction—labor supply. Ironically, a major source of construction labor – immigrants – has recently been central to policy debates, with prominent politicians attributing soaring housing prices to an increase in (illegal) immigration.[1] Therefore, critical questions remain: how does labor supply affect home building and housing prices, and to what extent does immigration enforcement exacerbate construction workforce shortages and housing affordablity?

This paper leverages an increase in immigration enforcement arising from a federal program called Secure Communities (SC), which began in 2008 and eventually rolled out to all counties nationwide by 2013.[2] According to US Immigration and Customs Enforcement (ICE) records, this program was associated with the deportation of more than 300,000 undocumented immigrants during this time period. As the residential construction sector is well known to be a large source of employment for undocumented immigrants (Svajlenka 2021), Secure Communities provides a laboratory for exploring the sensitivity of output in the residential construction sector to labor supply (and in particular labor supplied by undocumented workers), and also for understanding the demand-side impact of mass deportations.

We find that immigration enforcement affects housing markets primarily through the supply channel. We exploit the staggered spatial rollout of SC at the county level, along with microdata from the American Community Survey to show that treated counties experience large reductions in overall construction employment. These reductions last for at least three years after SC implementation. In addition, we also show that from the standpoint of the overall construction industry, domestic labor and immigrant labor appear to be complements rather than substitutes. Not only does domestic labor flowing into the construction industry (either from other industries or from outside the labor force) fail to fully offset employment

[1] E.g., in the vice-presidential debate of the 2024 election, the Republican candidate said "...you have got housing that is totally unaffordable because we brought in millions of illegal immigrants to compete with Americans for scarce homes."

[2] This program underwent several iterations between 2008 and 2021. We detail full timing in Section 2.

losses, but in fact, immigration enforcement leads to reductions in total construction workforce and in average hours worked for US-born workers also.

We then show that these net reductions in the construction labor force are associated with a slowdown in residential construction. We use two measures of construction activity. Data on homebuilding permits allows us to measure a decrease in planned residential construction. Because myriad factors – including labor shortages – may drive a wedge between intended and realized construction, we also use housing transaction microdata to show a reduction in completed homebuilding by measuring the total quantity of new construction entering local housing markets. Using both measures, we find reductions in homebuilding that increase over time and are large relative to baseline. Three years after SC rollout, the average county of approximately 500,000 residents has foregone the equivalent of an entire year's worth of additional residential construction: 2,578 fewer housing units are permitted, and 1,994 fewer newly constructed homes enter the market.

Next, we empirically document the anticipated link between reduced quantities and increased prices. In order to distinguish between supply-side and demand-side channels, we conduct analysis separately between new construction and existing housing stock. Price impacts in the new construction segment are more likely to reflect supply shifts than demand, because undocumented residents are relatively less likely to purchase new construction homes which are typically built at the higher end of the market. By constrast, our analysis of price impact within existing housing stock is better suited to study the general equilibrium effects of increased immigration enforcement on home prices—most saliently, the extent to which the mechanical link between increased deportations and fewer residents demanding housing services places downwards pressure on home prices. Long-standing theories of filtering in housing markets suggest that this partitioning can not entirely separate supply-side price drivers from demand-side as shocks at the highest end of the market will eventually percolate across the quality spectrum (Ohls 1975, Rosenthal 2014). Nonetheless, we argue that the separation helps to disentangle the two channels, especially over the near-term horizon which our setting enables.

In the new construction segment, we find evidence of endogenous adjustment on home attributes: workforce shortages appear to lead builders to produce slightly smaller homes. To ensure an apples-to-apples comparison, we control semiparimetrically for hedonic characteristics of housing and find a large increase in the quality-adjusted price of new residential construction. Three years after SC implementation, the average new construction parcel is 18% more expensive relative to the baseline (after controlling for home attributes).

Our analysis of existing housing stock suggests both short-term and medium-term effects. Within three years, the supply side impact seems to dominate. Using a repeat-sales estimation, we find that average home prices increase by 10% by the 3nd year after implementation. We do find evidence of moderate price declines in the two years following rollout: quality-adjusted price for existing stock dips by 2.5% – 3.75% during this initial period. Consistent with a demand-side mechanism, we find these declines are concentrated in census tracts with above-median share of LEFB residents. Our results suggest that any price declines arising from deportation are quite short-lived: within two years, upwards pressure caused by declines in homebuilding appear to dominate any near-term reductions in price in all tracts.

In the last section of our paper, we explore wage responses in the construction sector. If SC reduces labor supply but does not meaningfully change the optimal level of homebuilding for the industry as a whole, we would expect to see wages adjust upwards in order to attract replacement workers to the sector. We use two different wage measures to explore this relationship, including a region-by-year index of journeyman construction wages provided directly by local unions. We find relative wages drift upwards slowly, rising to 4.6% above baseline by the 3rd year of SC implementation. Our findings overall show that these increases are not enough to fully offset the SC shock to construction labor supply, nor are they sufficient to prevent contractions in construction sector output.

Contributions and Related Literature Our paper's contributions should be viewed along two margins.

First, we provide novel evidence that US homebuilding is highly sensitive to labor supply. Negative shocks to the construction workforce appear to be highly persistent and appear to have very meaningful effects on real economic output in the housing sector. In equilibrium, builders do not fully offset workforce shocks by raising wages to attract replacement workers. This adds to a literature that explores the apparent underproduction of housing, and the subsequent sharp increase in home prices. Several papers (for example, Molloy 2020, Albouy et al. 2016) highlight the secular rise of the housing cost in US household budgets in the recent period and associate it with limited housing supply. Glaeser and Gyourko (2018) attribute the shortage of housing supply to regulations on land use and building since it has been well documented both theoretically (Glaeser and Gyourko 2003, Ortalo-Magné and Prat 2014, Helsley and Strange 1995) and empirically (Malpezzi and Green 1996, Ihlanfeldt 2007, Zabel and Dalton 2011, Jackson 2018) that housing regulations reduce supply and lead home prices to exceed the marginal cost of construction. While there is widespread recognition on behalf of policymakers and academics that a growing housing affordability crisis fundamentally has its roots in restricted housing supply, most explanations have tended to assume policy-based barriers to expanding supply. Our findings suggest that factor constraints may play an important role as well. One implication is that policies designed to address housing affordability may be less effective unless they also help increase labor supplied to the construction industry.

Second, this paper contributes to the extensive literature that aims to understand the effects of immigration on local economies. Pioneered by the seminal work by Card (1990), many papers study the labor market effect using geographic variations of immigration flows but come to different conclusions (Altonji and Card 1991, Hunt 1992, Card 2001, Friedberg 2001, Cohen-Goldner and Paserman 2011, Borjas 2017, Borjas and Monras 2017, Monras 2020). All these studies focus on a sudden influx of immigrants, like Cuban immigration into the US during the Mariel boatlift (Card 1990, Borjas 2017), Mexican immigration to the US during the Mexican Peso Crisis (Monras 2020), and Jewish immigration into Israel after the collapse of the Soviet Union (Friedberg 2001, Cohen-Goldner and Paserman 2011). Lach (2007) and Cortes (2008) show immigration flows are associated with lower prices for nontradable goods and services.

Our paper is most directly connected to a subset of papers that study the impacts of immigration on the housing market. Previous studies mainly focus on the demand channel. Saiz and Wachter (2011) and Sá (2015) find growing immigration settlement is viewed as a negative amenity and leads to native flight and slower housing value appreciation. But at the MSA level, Saiz (2003, 2007) show more immigrants are associated with inflated housing rents.

Our setting allows us to focus on the housing market impact of outflows rather than inflows, which is especially critical to understand during a period when proposals for mass-deportations have become a part of mainstream political discourse. We are also able to leverage microdata on output in housing markets to understand how immigration flows affects housing supply as well as house prices.

A crucial issue faced by many papers that relate immigration to real economic outcomes is that immigration decisions are almost certainly correlated with the economic trend of local economies. Many papers partially solve this endogeneity bias by using previous immigration labor share as instrumental variables; however, the concern that immigration share is correlated with persistent economic shocks remains.

Our paper avoids this concern by leveraging quasi-experimental regulatory variations across US counties introduced by the gradual rollout of the SC program. Several other papers have used this same setting. East et al. (2018) show that SC leads to reductions in employment for likely-undocumented residents but does not lead to local increases in either employment or wages for domestic workers. Miles and Cox (2014) show that SC has no meaningful impact on local crime rates. Alsan and Yang (2022) show that SC leads to reduced uptake of federal social service programs for Hispanic residents, even among those not eligible for deportation. Grittner and Johnson (2021) find that SC decreases safety in workplaces while also making workers less likely to submit reports to regulators, especially in industries with relatively high Hispanic share.

The rest of this paper proceeds as follows. Section 2 describes the institutional details of the SC immigration shock. Section 3 describes our empirical approach, and Section 4 outlines the key sources of data. Section 5 presents our results. Section 6 concludes.

2 SC Background

Secure Communities (SC) was a US Immigration and Customs Enforcement (ICE) program that launched at the end of 2008. The central pillar of SC was enhanced information sharing between local law enforcement and federal immigration databases. Prior to SC, local policing authorities would not, in general, investigate a detained individual's immigration status as this required the physical presence of a federal officer (Miles and Cox 2014, Alsan and Yang 2022). Under SC, fingerprint information (already collected by local law enforcement pursuant to an arrest) began to be automatically shared with the Department of Homeland Security (DHS).[3] DHS would then match those fingerprints against an internal database of foreign-born individuals. A subset of individuals appearing in this database would be potentially eligible for deportation: (i) those who have been previously deported, (ii) noncitizens without any record of entry into the country, (iii) those with expired visas, and (iv) individuals identified as potential national security threats. Given a fingerprint match, ICE would validate that the individual is removable under immigration law and upon validation would coordinate with local law enforcement to take custody and begin deportation.

Because coordinating information and logistics across more than 31,000 booking locations nationwide was highly resource-intensive (assuming custody, for instance, requires arranging both transportation and bed space), it was clear from the onset that the program could not simultaneously launch at all locations nationwide (Alsan and Yang 2022). The initial launch included five counties in the last months of 2008. The program gradually expanded nationally, with the last set of untreated counties adopting SC at the beginning of 2013. While we have an exact date for the official start of SC in each county, our empirical analysis necessarily uses annual aggregates. Therefore we code counties as initially treated during the first year in which they have implemented SC for at least half the year. This means, for instance, that the initial set of counties launching SC in October to December of 2008 are coded as a 2009

[3] Specifically, fingerprints sent to the FBI to check an individual's criminal history (the existing standard), would then be forwarded by the FBI to DHS. Miles and Cox (2014) and Alsan and Yang (2022) provide extensive detail on the tactical implementation and respective roles of local police, the FBI, and DHS.

treatment-cohort. Our results are not sensitive to this choice. Figure 2 maps the expansion of SC by year, depicting the treatment indicator used in all regressions. Online Appendix Figure S3 maps treatment cohorts using the actual date of SC launch, without consideration for when in the year that initial date falls.

The last set of untreated counties adopted SC in January 2013.[4] The policy remained in place for the next 22 months. Beginning in late 2014, US immigration policy continued to shift on margins of both policy and branding. In November 2014, the Secretary of DHS announced the discontinuation of SC, and (on the same day) announced a new policy called the Priority Enforcement Program (PEP). The major difference between the two programs was the severity of offense that would lead to engagement with DHS: while all encounters with local law enforcement fell under the umbrella of SC, PEP applied only once an individual had been convicted of a relatively serious crime or if ICE believed national security interests to be at stake. In 2017, President Trump signed an order reinstituting SC, and in January 2021 President Biden signed an executive order revoking that reauthorization.[5]

Typically this period between late 2014 and early 2021 would complicate empirical analysis, as it is somewhat unclear whether this should be deemed a 'treatment' period or not. In our setting, however, state-of-the-art techniques in difference-in-differences analysis dictate that we use only variation through January 2013, at which point all counties become treated with the original iteration of SC. We are, therefore, not using any variation from the more difficult-to-interpret period from 2014 onward. We elaborate on this issue at length in Section 3.

3 Empirical Strategy

The phased rollout of SC between 2008 and 2013 allows us to run a county-level staggered difference-in-differences design. In its canonical form, this research design recovers a causal impact of some intervention by comparing the gap in outcomes between treated and untreated units before and after treatment. Our main empirical specification is:

$$y_{it} = \alpha_i + \gamma_{rt} + \beta \mathbb{1}(treatment_{it}) + \epsilon_{it} \tag{1}$$

[4] There are nine counties for which we do not have an adoption date. Each appears to be a very small county with an atypical governance structure. Our sense is that these counties are each likely folded into the administrative governance of a larger neighboring county and therefore do not represent non-treatment regions.

[5] https://trumpwhitehouse.archives.gov/presidential-actions/executive-order-enhancing-public-safety-interior-united-states/ and https://www.whitehouse.gov/briefing-room/presidential-actions/2021/01/20/executive-order-the-revision-of-civil-immigration-enforcement-policies-and-priorities/ respectively.

where i denotes counties and γ_{rt} is a set of census-region-by-year fixed effects. We partition the time fixed effects by census region because population trends differ sharply by region, as Figure A1 shows. Long-term population growth will affect both housing demand and housing supply (to the extent that developers are forward-looking), and so this choice ensures that identifying comparisons are between counties on similar growth paths.

In recent years, several papers have shown a potential for bias in the estimated causal treatment effect, β, that arises specifically in staggered-rollout designs (De Chaisemartin and D'Haultfoeuille 2022, De Chaisemartin and d'Haultfoeuille 2020, Goodman-Bacon 2021, Callaway and Sant'Anna 2021, Borusyak et al. 2021). De Chaisemartin and D'Haultfoeuille (2022) survey several papers finding that bias is more likely in settings where most units are eventually treated. The scarcity of untreated units in later periods means that the two-way fixed-effects (TWFE) model necessarily places greater weight on potentially problematic pairings that use already-treated units as the control observation. A related issue arises if treatment effects increase in treatment duration. Sun and Abraham (2021) show that the resultant time-heterogeneity in treatment effects can lead to spurious violations of the parallel trends assumption that underlies causal interpretations of DiD estimators.

Both potential drivers of bias are present in the SC setting. First, essentially every county is treated at some point. From East et al. (2018) we have an activation date for 3,126 counties across 50 states and the District of Columbia. The remaining counties without an activation date are very small regions, either with an atypical governance structure or which are grouped with another statistical reporting unit. Therefore, these counties do not comprise an appropriate counterfactual region for the period after full rollout of SC. Second, the local impact of increased immigration enforcement is likely to be heterogeneous over time. An individual's choice to emigrate is presumably a function of: (i) expected economic payoff to residing in the United States while undocumented, (ii) the available payoff to remaining in the home country, and (iii) expected costs due to immigration enforcement actions. The first two of these factors are time-varying, which suggests that a shock to the expected costs of immigration enforcement will have different effects on immigration flows at different periods in time. In addition, it seems very possible that increasing immigration enforcement could have an impact that increases over time: if network effects are important for generating a payoff to migrating to any region, then increased deportation may make future inflows less appealing or less likely.[6]

Several papers have presented estimators that can address these concerns with the stan-

[6] Alternatively, one could tell an opposite story as well: if local demand for immigrant residents is static and capped for any reason, then increased deportations could make immigration more attractive. Either pattern would lead to a causal effect that shifts as a function of treatment time.

dard TWFE model. Our preferred specifications all use the approach of Gardner (2022). This is a two-stage estimation technique that first estimates both sets of fixed effects (county and region-year) from untreated units. Practically, this means that increasing treatment effects over time will not erroneously shift cross-sectional averages, nor will region-by-year fixed effects late in our sample rest primarily on outcomes in treated regions. Purged of problematic variation, the estimated fixed effects from the first stage are used to produce fitted values which residualize the dependent variable for use in the second-stage estimation. This approach is appealing for both its transparency and computational simplicity. In particular, analytic inference is possible, which is advantageous for our analysis of prices which draws upon transaction microdata spanning tens of millions of observations. While the use of bias-robust estimators does induce meaningful differences from a standard (and incorrect) TWFE DiD estimator, our results are not sensitive to the specific estimator selected.

3.1 Is SC Rollout Predictable?

The rollout maps shown in Figure 2 quickly suggest that SC did not launch in a purely exogenous manner: there is a clear pattern of rollout from the Southern border upward. Endogenous treatment in a DiD design requires a stronger identifying assumption: that divergence of outcomes after treatment is not driven by some factor that correlates with selection into treatment. This assumption is partially testable, and as usual, empirical evidence of parallel pre-trends is a necessary condition for having confidence that it holds.

Additionally, in this setting, because SC begins largely in Southern areas only a couple of years after the peak of the 2000s housing boom, we also want to explore the possibility of confounding long-run trends that correlate with selection into treatment. The concerning story would be something along the following lines: SC rolls out initially in Southern states; however, many of these are also the so-called "Sand States" which saw the largest run-up in construction and home prices during the 2000-2007 period, suggesting the potential for a cyclical collapse in building after the Great Recession that happens to coincide with SC rollout.

Other scholars have explored determinants of SC rollout. Cox and Miles (2013) consider measures of crime, income, non-Hispanic immigrant share, and political attitudes. Despite clear rhetoric from SC leadership implying a focus on jurisdictions facing high levels of crime, the authors find that only two factors strongly predict county rollout: (i) sharing a border with Mexico and (ii) Hispanic population share.

In Table 2, we explore several additional factors that are directly related to housing demand, housing markets, and dynamics of the Great Recession cycle. Each column shows the results from regressing an indicator for SC rollout in a given year on county-level characteristics in a stacked dataset, where each stack consists of all counties that either launch SC in a given year or counties that are not yet treated. (So the 2010 stack, for instance, would code counties launching in 2010 as 1, and code counties launching in 2011, 2012, or 2013 as 0, and exclude counties which already launched in 2009.) This regression, therefore, gauges whether a given county-characteristic predicts the specific year of rollout.

Column 1 confirms the existing finding from the literature that *ex-ante* Hispanic population share is a strong predictor of rollout. This county-level feature will be absorbed by the county-fixed effect in our regressions, and so from column 2 onward, we retain Hispanic share as a control variable, and test for the marginal effect of other predictors. Column 2 shows that county size does not meaningfully predict rollout. Columns 3–6 test the predictive ability of three-year lagged population growth trends. SC rollout is not associated with a three-year growth trend in US-born population. SC rollout is weakly associated with trends in Hispanic population, however the effect is small and statistically marginal. The standard deviation of Hispanic 3-year growth in the sample is 29.2%, meaning that an increase of one standard deviation correlates with a 1.4pp decrease in likelihood of SC launch in a given year. Column 5 shows that growth in LEFB population does not predict rollout.

In columns 6 and 7 we test directly for a link between SC and housing boom-bust dynamics. In column 6, we compute the growth in new construction (total square footage) in the three years prior to launch. We do not find any evidence that prior building predicts SC rollout either economically or statistically.[7] This means that our results are unlikely to be driven by cyclical fluctuations in building around the Great Recession. In column 7, we test whether the total price run-up between 2001 and 2007 predicts SC rollout. Again, we find no predictive power, suggesting that any results we find are unlikely to be driven by a correlation between treatment and house price patterns in the years prior to the Great Recession.

In total: the literature already suggests that SC launches earlier in localities with higher Hispanic share, and our findings confirm this. Apart from a weak association with lagged Hispanic growth, other relevant factors do not appear to be a strong predictor of rollout. Growth within the LEFB population, which we will show is particularly impacted by SC, does not hold predictive power. And we find no reason to be concerned that our results may be driven by some correlation between SC rollout and overbuilding during the housing boom, leading to a subsequent collapse in residential construction.

[7] The sample reduction in column 6 is due to a small number of counties that do not show any new construction in the base year, making a growth calculation impossible.

4 Data

This section introduces datasets used in this study.

4.1 Secure Communities

Information on the rollout dates of SC comes from East et al. (2018), who gather the implementation date of SC at county level from ICE.[8] Based on the rollout dates, we construct a county-year-level dummy variable indicating whether SC has been implemented. Due to additional concerns about bias in DiD designs that use continuous treatment variables, we use a binary indicator for SC treatment rather than a continuous variable capturing partial treatment in the implementation year. We code a county as treated in the year of launch if SC was introduced for at least six months of the year, and untreated otherwise. Once treated a county remains treated throughout the sample.[9] SC was implemented in all counties by 2013.[10] Among 3,126 counties, 2% adopted SC in 2009, 11% in 2010, 31% in 2011, 53% in 2012, and 3% in 2013.[11]

4.2 American Community Survey

We gather county-year level population and employment information from the 2005-2020 American Community Survey (ACS) and merge it with the SC rollout data. County-year level variables are aggregated from individual-level information with individual weights. Because SC rollout occurs at the county level, we need to create ACS measures also at the county-level. ACS microdata is released with geographic granularity at the Public Use Microdata Area (PUMA) level: regions with at least 100,000 people. This means that PUMA-to-county links are possible only for relatively large counties. We are able to create the relevant measures for 331 counties, which form the sample for our analysis of employment and population. For other results which do not rely on the microdata, we are able to produce both national estimates as well as estimates based only on the ACS-covered sample of counties.

[8] We thank Chloe East for generously sharing this data with us.

[9] As discussed in Section 2, SC underwent periods of suspension and/or rebranding starting in 2014. As a consequence of our empirical design, our estimates come from identifying variation between 2009 and 2013, and so this post-2014 period does not affect our results.

[10] We exclude the following counties in our analysis due to missing SC implementation information: Hoonah-Angoon Census Area, Alaska (FIPS code: 02105), Kalawao County, Hawawi (15005), Shannon County, South Dakota (46113), Emporia City, Virginia (51595), Fairfax City, Virginia (51600), Manassas City, Virginia (51683), (51685), Poquoson City,Virginia (51735), Doddridge County, West Virginia (54017).

[11] If using the calendar year of SC implementation, among 3,126 counties, 0.4% adopted SC in 2008, 3% in 2009, 25% in 2010, 35% in 2011, 33% in 2012, 3% in 2013.

4.3 Permits

County-level permits are from the Building Permits Survey (BPS) from the Census Bureau. The number of permits represents the amount of new privately-owned residential construction planned in each county in each year. The permits data contains information on both the number of permitted buildings and the number of units represented by those buildings. Both buildings and total units are reported in bins by building size: buildings with one unit, two units, three or four units, and five and more units. In regressions, we normalize the number of permits by the county-level population in 2005.

4.4 CoreLogic

CoreLogic compiles deed transaction records and property tax roll information from US county assessor and recorder offices. This data spans the near-universe of properties in the US, including variables on property characteristics, geographical locations, ownership changes, transaction date, and sales prices. A sale is flagged as a "new construction" transaction if the property is sold from the builder to the first owner. We construct county-year level measures of new construction sold by aggregating the number of properties sold, as well as the total square footage of properties sold. Importantly, CoreLogic allows us to observe both the year in which newly constructed homes are sold, as well as the year that the home was built. This allows us to be sure that any decrease in the amount of new construction sold into a given market is not simply arising from longer delays between completion and sale.

Our final dataset contains 4.22M observations of newly constructed homes between 2005 and 2012. We have arms-length market prices along with a full set of hedonic characteristics for approximately 59% of the observations. We provide complete detail for each step of the CoreLogic data build in Section 1 of our Online Appendix.

4.5 Wages

We use two measures of wages: (i) quarterly data on construction labor costs for more than 700 urban regions from 2007 onward compiled by RSmeans, a leading private supplier of benchmarking data for the construction industry, and (ii) wage indexes constructed directly from ACS microdata. The RSMeans data is obtained via survey: for each locality and in each year, RSMeans contacts local unions to directly solicit journeyman wage rates for each of 21 different trades. That wage information is aggregated by the weight of that trade's usage within the industry and normalized to be a region-state-year index of total construction wage costs relative to the national average. The regions surveyed tend to be the largest 10-15 urban areas in a given state. Most regions directly correspond to a county in our baseline dataset.

Some counties in the ACS data do not contain a region that RSMeans surveys; in these cases we match that county by hand to the closest region that RSMeans does cover (in the vast majority of such cases, the match is between some urban area and a close suburb).

As the wage data reported by RSMeans does not correspond to any particular individual, we cannot use this series to produce wage measures for subpopulations of interest. However, we can construct measures for LEFB workers and US-born workers by using the ACS data. That said, the ACS data is potentially subject to a meaningful source of measurement error: while wages earned are reported as a continuous variable, weeks worked are reported in fairly broad bins. This means that smallish intensive margin adjustments in time worked will tend to be unobservable. In a setting where reductions in work are prevalent, this may lead to a downward bias in observed measures of average wages. When we look at wage impact within the ACS data, we consider the ratio between construction-sector wages and all-industry wages. This will reduce scope for any bias arising from data structure to the extent that mis-measurement is the same across industries.

Table 1 provides summary statistics for key variables and datasets, including population, workforce, permits, new construction microdata, and wages.

5 Results

Our preferred specifications are event study versions of equation 1:

$$y_{it} = \alpha_i + \gamma_{rt} + \sum_k \beta^k \mathbb{1}(time\ since\ treatment_{it} = k) + \epsilon_{it}. \tag{2}$$

As described in Section 3, we use the bias-robust estimator of Gardner (2022) for all estimations. Because this approach relies on estimating fixed effects from pretreatment data and because every county is eventually treated, we face a mechanical limitation on the number of posttreatment coefficients that can be identified: this cannot exceed the number of periods separating first treatment from last treatment. In our empirical execution, SC dummies turn on between 2009 and 2013, which implies a maximum horizon of four periods: contemporaneous effect plus 3 future years. However, our use of census-region-by-year fixed effects reduces this horizon an additional year whenever we are using ACS data (e.g., for any results that use LEFB status). Within the set of counties that can be separately identified in ACS microdata, in 2009 there are treated counties in the Northwest (3), South (55), and West (6) regions. By 2012, the only set of untreated counties are in the Midwest, meaning that no identifying variation for four years of treatment remains within census region. As a result, we can estimate an impact for three periods: the contemporaneous effect plus two subsequent years. That estimate comes from comparing 2009-treated counties to within-region untreated counties in 2011, plus Midwest counties treated in 2010 ($n = 13$) compared to the untreated Midwest counties in 2012.

In the full sample of counties (that is: not restricting attention to counties which can be separately identified in ACS microdata) we are able to estimate a four-year horizon, because there are a set of counties ($n = 30$) in the South region which remain untreated in 2012 and therefore serve as a counterfactual for the earliest-treated set of counties. Therefore when considering quantity and price outcomes, which come from CoreLogic or the Census Building Permits Survey, we can validate the extent to which trends evident over the $T \in [0, 2]$ horizon appear to continue through $T = 3$.

5.1 Workforce Impact

The Department of Homeland Security maintains administrative records of individuals deported under Secure Communities. These records obviate the need to validate an overall first stage relationship via estimation: we know that SC did lead to individuals being removed from the county. Figure 3 shows the number of removals under SC by calendar year (top panel), and relative to the launch year (bottom panel). This data comes from administrative ICE records, and county-year aggregates are made available by the Transactional Records Access Clearinghouse at Syracuse University. Mirroring the time-window used in our empirical approach, these figures include removals from SC inception through the end of federal fiscal year 2013 (Sept 30th, 2013), but do not include any removals attributed by ICE to SC from fiscal year 2014 onward. The data reflect a total of 306,000 individuals removed. One salient feature, evident in the bottom panel, is that the pattern of removals in a given county is increasing and concave. Treated counties do not experience a single spike of removals, but rather see increasing numbers of individuals removed for multiple years. This suggests that treatment effects may build over time, and many of our results will be consistent with this.

To understand how SC impacts construction workforces, we cannot replicate the prior analysis because the administrative data contains no information on occupation for those deported under SC.[12] As a result, an industry-level analysis requires estimation using ACS microdata. ACS contains reported occupation as well as information on time worked over the prior year. Our measure of workforce includes those of working age who report construction occupations, regardless of employment status. Informally, we are counting those who report themselves to be construction workers, even if they didn't work over the prior year. This is

[12] To the best of our knowledge, ICE does not solicit or record such information as part of removals under SC.

the best measure of labor supply that we can extract from the ACS data, but it is important to recognize that this is an equilibrium outcome. Physical removal under SC is one way that construction workers in a given county can be reduced. However, in the ACS data, switching to another occupation will also reflect a reduction of labor supplied to the construction sector. Therefore, our estimates of SC impact will include direct removals plus other spillover impacts of immigration enforcement.

Because documentation status is not asked in the ACS data, we use three demographic groupings as proxies for undocumented immigrants: noncitizens, low-education and foreign-born (LEFB), and those indicating Hispanic ethnicity. All three proxies are imperfect. Noncitizens will include not only undocumented residents but also conditional and permanent residents, as well as those holding nonimmigrant status.[13] In addition, misreporting may be high if undocumented residents are hesitant to respond truthfully to questions about citizenship (Van Hook and Bachmeier 2013). The LEFB grouping is a standard designation used in the immigration literature. Of course this grouping will also include naturalized immigrants with low education. Finally, we consider respondents who indicate Hispanic heritage.[14] Although this grouping will certainly include a large number of US citizens, approximately 30% of the US construction workforce is Hispanic[15] and an estimated 25% of the construction workforce is undocumented (Svajlenka 2021). The grouping of Hispanic respondents, therefore, is likely to include a nontrivial share of those potentially impacted by SC. For all results, we include a fourth grouping of US-born residents as a natural comparison set.

While using ACS microdata allows us to differentiate between US citizens and likely undocumented immigrants, it does also constrain us to examine relatively larger counties, as smaller counties are grouped together into a single Public Use Microdata Area (PUMA). In these latter cases, we cannot appropriately assign a SC start date to each county within the PUMA. As a result, the population and employment data used in this section spans 331 counties with a total population of 156 million (in 2005). This represents just over half of the total 2005 US population.

Figure 4 shows that SC leads to a reduction in construction workers for treated counties. We normalize the number of workers by 2005 county-wide population. The evidence is highly consistent across all three groupings that encompass undocumented workers. In all cases, there is a sharp decline beginning in the year of treatment and ongoing reductions during

[13] Source: https://travel.state.gov/content/travel/en/us-visas/visa-information-resources/all-visa-categories.html

[14] The ACS data permits us to exclude respondents indicating Puerto-Rican heritage

[15] The Construction Industry: Characteristics of the Employed, 2003-2020; Bureau of Labor Statistics, available at https://www.bls.gov/spotlight/2022/the-construction-industry-labor-force-2003-to-2020/home.htm.

the horizon we can examine. Magnitudes are quite similar. Taking the estimates for the LEFB population, the contemporaneous implementation of SC leads to a reduction in LEFB construction employment equivalent to 5.3bps of the county population. For the average county in our data, this is equivalent to 271 fewer workers. The net impact of SC increases over time. At $T = 2$, we estimate the total reduction to be 26.4bps of total population—1,296 workers for the average county. Because our measure of workforce is not conditioned on employment, these estimations show that SC leads to fewer individuals indicating that they work in construction in a given county. Given the setting of immigration enforcement, it is clear that some of this shift must come from deporting construction workers. However, the LEFB estimate for $T = 2$ applied to the total population of the data sample (162M) suggests a reduction of slightly more than 425,000 workers. This is roughly the order of magnitude of *all* individuals removed under SC. However, only approximately 9-12% of LEFB respondents report construction occupations in the ACS data. Therefore, absent a belief that SC is sharply more likely to affect construction workers than those in other occupations, a meaningful share of net reductions in workforce must also include non-deportation channels: for instance, workers choosing to shift into other sectors, workers leaving the labor force, or residents voluntarily returning to country of origin (perhaps accompanying a deported family member).

As described in Section 3, our event studies are all estimated using the two-stage estimation technique of Gardner (2022). This technique has an important empirical consequence: the control group of observations used to identify time fixed-effects shifts at each time period (because only units remaining untreated are used). In Figure 5, we show the impact of SC on LEFB construction workers by treatment cohort using a balanced (and constant) control group of the last counties to be treated. Cohort 1 compares counties treated in 2009 with only counties untreated through 2012. Cohort 2 compares counties treated in 2010 with the same control group, and likewise for cohorts 3 and 4. For feasibility, it is necessary to run this cohort-by-cohort analysis using year fixed effects instead of the more rigorous region-by-year fixed effects which we employ in the pooled specification (this is just a consequence of the geographic rollout pattern). We include this evidence to demonstrate that the unusual use of a control group that changes in composition through time does not drive the patterns we find. Each cohort, compared to a fixed control group, shows evidence of a similar decline in LEFB construction workers. Each subsequent cohort loses an estimation period because all counties are treated in 2013. This prevents us from seeing a full four years of evidence for each cohort, but based on what can be examined, patterns look very similar across cohorts.

Returning to Figure 4, we also find that SC leads to the decline of US-born construction workers. If, as a first-order effect, increased immigration enforcement does not impact demand for construction services – an issue which we revisit at length in Section 5.4 – then the natural prediction would be for US citizens (or legal residents) to fill the vacant positions, increasing employment share for that population. Yet the event study shows declines. The contemporaneous effect, which is statistically significant at 5%, suggests a decline of 599 workers for the average county. This decline appears to persist for the next two years, however these estimates are not statistically significant.

One potential explanation is that construction labor markets are segmented and that undocumented labor supply acts as a complement to domestic labor rather than being a substitute. We test this theory according to skills-based segmentation. Studies have shown that undocumented immigrants are more likely to hold lower-skilled jobs than domestic workers.[16] One possibility is that higher rates of unionization in skilled trades create additional barriers to undocumented workers holding these jobs. If a shortage of lower-skilled labor makes it more difficult to find workers to finish framing a house, this will also reduce demand for electricians and plumbers required at the subsequent stage of construction.

We classify the following ACS occupation codes as "lower skill": (i) construction laborers, (ii) helpers in construction trades, (iii) painters and maintenance, (iv) drywall installers, (v) carpenters, and (vi) roofers. All remaining categories within the construction subcategory are classified as higher skill. In addition to management occupations (supervisors) and frequently unionized occupations (e.g., electricians and plumbers), this classification includes several occupations that may have significant skill heterogeneity (e.g., sheet metal workers or hazardous materials removal workers). As a result our partitioning should be regarded only as a high-level separation between occupations which are likely to include the lowest-skilled workers, and a set of occupations that are, on average, higher skilled.[17]

Figure 6 shows event-study results by skill classification. Within lower-skill occupations, domestic labor appears to be a substitute for immigrant labor. SC leads to a large reduction in low-skilled LEFB construction workers: a peak loss of 953 workers for the average county at $T = 2$. However, these losses are partially, but not totally, offset by increases in US-born employment for lower skilled occupations. The rate at which US workers replace lost LEFB labor, measured as the ratio of coefficients, is between 21% and 63% in the periods following SC adoption. At $T = 2$, the net effect suggests a loss of approximately 419 workers in these lower-skilled occupations for the average county.

An opposite pattern holds within higher-skilled occupations. SC also causes reductions

[16] https://www.pewresearch.org/fact-tank/2020/02/24/the-share-of-immigrant-workers-in-high-skill-jobs-is-rising-in-the-u-s/

[17] The authors would also like to acknowledge some level of discomfort with describing any of these jobs as lower-skilled; we suspect we'd find all of them quite difficult.

in high-skill LEFB employment; however, the impact is substantially smaller within this population: a peak impact of 356 fewer workers. Rather than experiencing any offsetting increase, higher-skilled US-born workers see employment declines that are larger than those in the high-skilled LEFB segment. Immigration enforcement appears to reduce the overall quantity of low-skilled construction labor supplied because domestic labor only partially offsets the shock to immigrant labor supply. In turn, this reduction in low-skilled labor supply appears to be associated with an overall shrinkage in higher-skilled labor supplied by both domestic and undocumented workers. At $T = 2$, our estimates imply a reduction of just over 1,300 higher-skilled workers for the average county; and 73% of that reduction comes from US workers.

Aggregate statistics on labor supplied to the construction industry confirm an industry-wide slowdown. From the ACS data, we compute total hours worked in each county by taking the reported number of weeks worked and multiplying by the reported number of hours per week the respondent usually worked in the preceding 12 months. We also compute the average annual hours worked *per construction worker remaining* in a county in a given year – this measure captures intensive margin adjustments within worker. Figure A2 shows how SC affects total labor supplied. The top panels suggest that both LEFB and US-born workers tend to work approximately the same amount. The LEFB point estimates show reductions on the order of 5% without any statistical significance. For US-born workers, the estimated impact is a fairly precise zero. The bottom panels show a reduction in total hours supplied to the construction sector in equilibrium. Given the lack of large intensive-margin adjustment, this reflects fewer people employed in the construction sector. In the LEFB population, as discussed, a large share of this reduction must come from SC removals. The complementarity between LEFB workers and US-born workers is evident in the bottom right panel. In total, the reduction in labor supplied by US-born workers is 3.4 million hours two years after SC launch (this estimate is only significant at 10%).

5.2 Direct Evidence on Homebuilding

In this section, we document the slowdown in output associated with the reduction in equilibrium labor supply. We focus on two measures of residential construction activity: intended construction (using residential permits) and completed new construction transactions (using administrative tax-roll microdata).

5.2.1 Permits

We begin by examining permitting intensity from the US Census Building Permits Survey. This survey contains data for nearly all counties nationwide, allowing us to estimate an effect both within the subsample of 331 counties that are separably identifiable in ACS and which therefore underlie the workforce analysis of Section 5.1 (henceforth the "ACS subsample"), as well as the full national sample. We estimate event studies, as per equation 2, with total permitting per 1,000 residents as the outcome variable. Figure 7 shows the results for permitted units (top) and total buildings units (bottom). These two measures are very highly correlated ($\rho = .91$), and so results are quite similar between the two specifications.[18] SC leads to a sharp reduction in both permitted units and permitted buildings. This effect is quite large. Focusing on the top panel of Figure 7, SC leads to .44 fewer units per 1,000 residents in the launch year, .6 fewer at $T = 1$, and 2.26 fewer units at $T = 2$. For the average county, this implies a total reduction of 1,624 units over the three-year period. While the Great Recession complicates measurement of baseline activity, the average number of units permitted across counties in 2005 was 3,464, and the average number permitted across all county-years prior to SC implementation (this includes a number of post-housing-boom observations) is 1,896. This is a very large reduction relative to either baseline. Relative to the top panels, the slightly smaller effects evident in the bottom panels on permitted buildings show that SC also induces builders to make intensive-margin adjustment on project size.

Patterns in the national sample (right) are extremely similar to the ACS subsample. In this larger estimation sample, all estimates are strongly statistically significant. In addition, using the national sample allows us to estimate impact for an additional year (as discussed at the beginning of Section 5), and we find that the permitting declines associated with SC persist. In the national sample, our estimates suggest a total reduction through $T = 3$ of 2,578 fewer units.

Figure S5 in our Online Appendix disaggregates the effect by building size in the national sample. Single-family homes are the chief driver of the decline, which is unsurprising as single-unit buildings also represent the majority of housing stock in most places. The statistically significant decline in units within three- or four-unit buildings is small: the cumulative impact across the four-year period translate to a reduction of 68 units for a county with half-a-million residents (comparable in size to the average county in the ACS subsample). In the national sample, there is no evidence of declines in units within the largest buildings. In the ACS subsample (not shown here), the corresponding point estimates suggest that declines in

[18] In addition, we use the exact reported number of buildings and units, rather than a model-based imputation of totals which the Census also provides. This choice does not meaningfully affect results.

largest-building units may be slightly more consequential, however the point estimates are still insignificant.

5.2.2 Observed New Construction

For several reasons, changes in permitting activity might not correspond to changes in actual construction. We use administrative tax-roll data from CoreLogic to test directly for new construction supplied to the housing market.

CoreLogic's deeds records contains a flag for new construction sale, which lets us aggregate a measure of new construction by county-year. Using home sales still leaves a possible wedge between total construction activity and our econometric measure: homes may be built but fail to sell. However, we observe both the date of sale and the year in which the house is reported built. Therefore, we can test directly for new construction amounts by using sale dates. But, crucially, by focusing on the built date, we can also test new construction completed – as long as that property sells at some point before 2022 (the last year reflected in CoreLogic). This means that our analysis will only miss new construction that fails to sell for more than 10 years. This window is sufficiently long that it is likely to include only a small and highly idiosyncratic number of properties.

Figure 8 shows the results from an event study regression following equation 2, where the dependent variable is the aggregate square footage of new construction (per 1,000 residents) entering the local market. Like our analysis of building permits, we find very similar results between the subset of 331 ACS-identifiable counties (left) and the full national sample (right). Our preferred specification is based on built-year (top row) as this is when the construction actually occurs. We include the sale-year results (bottom row) to demonstrate that changes in time between completion and sale are not large enough to meaningfully affect these aggregate quantity results when estimated at annual frequency.

In both the ACS subsample and the national sample, SC is clearly associated with reductions in homebuilding. These reductions increase over time, and in both samples, the $T = 2$ estimate is approximately 2,500 sq. ft. per 1,000 residents. As before, using the full national sample (right figures) allows us to estimate an additional year's worth of impact. The evidence suggests that the magnitude of reduction continues to increase. This is a flow measure of building, so to understand total magnitude over our estimation window we cumulate all post-treatment estimates and multiply by average county population. In the ACS and national samples, the total reduction through $T = 2$ is 2.25M and 2.37M sq. ft. respectively. In the national sample, the impact through $T = 3$ is 4.34M sq. ft. The median new home (prior to SC) is 2,232 sq. ft., which implies slightly more than 1,000 homes built in the average county over a three-year horizon, and 1,944 fewer homes in the national sample over the four-year horizon. These figures are somewhat smaller than the reductions we find in permitting activity, which may suggest that the speculative nature of some permitting (one could imagine builders filing despite being less than certain a given project will proceed, for reasons of bureaucratic or timing efficiency) leads permitting to be more sensitive to shocks than realized output.

Our preferred specifications in Figure 8 are based on reported square footage in the CoreLogic microdata. It is a stylized fact of real-estate microdata that hedonic attributes are unevenly recorded across properties. Happily for our purposes, square footage is recorded much more frequently than other attributes, however approximately 8% of observations still lack this information. For observations missing square footage, we impute size from the transaction price using the national average price per square foot. Although imperfect, we expect that this imputation will only serve to reduce classical measurement error, as we know of no reason to be concerned that the fidelity of attribute reporting to the local tax assessor would be correlated with any driver of homebuilding (especially conditional on county-fixed effects, which would absorb potentially-reduced administrative capacity in smaller or more rural counties). Figure S4 in our Online Appendix shows the results of repeating our quantity estimation without any imputation. As anticipated, we find very similar results. Slightly smaller magnitudes are consistent with increased measurement error from omitting any observation missing square footage.

5.2.3 Does Building Decline Because Immigration Enforcement Reduces Population?

One possible mechanical channel by which immigration enforcement could affect homebuilding is through the direct impact on population: every individual removed is, all else equal, one less person demanding housing services in a given region. In addition, immigration enforcement likely has direct population impacts beyond the focal individual subject to deportation. Family members may elect to leave the country alongside a deported individual, some people may choose to voluntarily return to home countries rather than accept exposure to future enforcement actions, and *ex-ante* decisions to immigrate may shift as a consequence of higher baseline enforcement. Builders, facing declining populations, might find a reduced level of optimal building. We explore the extent to which the data seems to support this type of mechanical channel for reduced construction output.

Figure A4 shows the results of estimating SC population impact within the ACS subsample (left) and the national sample (right). Within ACS counties, SC appears to have no impact on total population. Estimates are relatively precise zeros, with a single point estimate corresponding to any economically meaningful deviation from zero suggesting population *increases* at $T = 2$ rather than declines. In the national sample, we do find evidence of a treatment effect, but the data suggest population increases rather than decreases. Total county-level population is a complex equilibrium outcome, and so while these increases are perhaps surprising, there are a range of mechanisms that could push in this direction. One could imagine that deported individuals are disproportionately single adults (with families remaining in the country of origin, perhaps), and that deportation results in a vacated job. Some of these jobs may be filled us US citizen with families who relocate from another county. While this is an example of a pattern that would generate net population gains, it is by no means the only possibility. We do not try to trace out and empirically disentangle the various channels that could lead to population increases. This finding is consequential in this paper in two ways. First, it suggests that mechanical population *declines* are not a plausible channel generating the reductions in homebuilding that we find. To the extent that SC increases population – either as a true casual impact, or due to some confounding factor that we have been unable to identify – it biases us *against* finding declines in homebuilding. Second, however, an increasing population does present a confounding factor when we explore the price impact of reduced homebuilding. Any positive correlation between population and prices will bias us *towards* finding increases, and so as we explore prices in the next section, we will be careful to assess the extent to which our findings could arise from changes in population (see 5.3.3).

5.3 House Prices and Changes in Home Characteristics

So far we have shown a quantity response. An immigration shock reduces the number of workers in an industry that draws significantly upon undocumented labor. This reduction is persistent in time and leads to construction slowdown which we observe in both permitting patterns and in the supply of new construction entering housing markets. We now connect this quantity response with prices.

We segment our analysis of prices into two components in order to distinguish between supply- and demand-side drivers. Our first analysis focuses on new construction in an attempt to capture supply-side impacts of SC. The housing literature has documented that new construction tends to be added at the upper end of local house price distributions, and so to the extent that undocumented immigrants are unlikely to be purchasing above-median housing stock within a given region, the demand-side impact of SC is less likely to affect price estimates within this segment. Our second analysis considers the price impact on existing housing stock. There is no reason to believe that SC represents any meaningful supply shock in this segment.[19] As a result, estimated impacts on existing housing stock seem more likely to arise from demand shocks caused by immigration enforcement. We acknowledge that this separation sharper in the near-term. Filtering theories in housing markets suggest that we should expect price spillovers from new construction to existing housing stock, allowing supply-side drivers to affect existing house prices, however this link is both indirect and potentially realized on a longer timescale. Our preferred interpretation is that price impact in new construction is most likely generated through supply-side channels; and that price impacts in existing housing stock reflect both supply- and demand-side channels, but are likely most responsive to demand-side channels in the near term. In both analyses, we focus on single family homes.

5.3.1 Price Impact in New Construction

We estimate event studies using (logged) transaction prices as the dependent variable. We test for impact both in raw transaction prices and in specifications that include a rich set of hedonic controls. We also provide evidence on endogenous changes in housing stock attributes that generate a wedge between these two specifications. We consider the most relevant estimate to be the impact on quality-adjusted prices. This measures how reduced homebuilding affects prices, imagining that we can compare identical homes. Another standard approach to provide this analysis would use a repeat-sales sample along with a property-level fixed effect to control for all (time-invariant) property unobservables. However, due to our focus on new construction, we cannot make use of a repeat-sales technique. Instead we use home characteristics reported in the CoreLogic data to control for quality. We use three major attributes to capture size: square footage, number of bathrooms, and number of bedrooms. For square footage, we convert the continuous variable into small discrete bins and then employ a fixed effect for each bin. We include fixed effects for the integer number of bedrooms and bathrooms, along with fixed effects for age and census tract.

Figure 9 shows the result. The top row shows the impact on quality-adjusted home prices. As usual, the left column focuses on the ACS subsample, and the right uses the national sample. We find that SC increases house prices. This occurs at a delay: there is no meaningful price response until two years following implementation. The delay in price impact seems likely to reflect the slow-moving nature of the homebuilding industry; it is not unreasonable, for instance, to think homes already under construction or in the final stage of planning when SC was implemented would be more likely to be completed and that the

[19] That said, spillover effects are possible: SC's impact on labor supply could slow down residential redevelopment, indirectly preventing existing stock from being converted to new construction.

largest impact would be on very early-stage projects. Such timing also parallels the gradually intensifying effects on permitting and new construction completion that we find. The results of Figure 9 suggest that it takes about two years for shortages to become salient enough to have a large impact on market prices. At $T = 2$, the increase in price is approximately 6% in both the ACS subsample (top left) and the national data (top right). The extra estimation year available with the national sample suggests that these price impacts persist and increase: at $T = 3$, the average quality-adjusted new construction property has become 18.5% more expensive.

For completeness, the bottom-panels of Figure 9 show the impact of SC on raw prices. While the same broad pattern holds, the increases in years two and three are slightly smaller without controlling for housing attributes. This appears to be coming from endogenous shifts in the characteristics of homes that are built. Figure A3 shows that SC leads to reductions in overall square footage, number of bedrooms, and number of bathrooms for the average new construction project. All else equal, these changes will reduce transaction prices, counterbalancing shortage-induced increases. But without careful quality controls, the comparison is not apples to apples: paying the same amount for a smaller home is also, of course, an increase in the price of housing. It is also possible that builders reduce quality on unobservable margins. If this is an important channel, then the increases we estimate are understated.

5.3.2 Price Impact in Existing Housing Stock

Panel A of Figure 10 estimates the price impact within existing housing stock in the national sample. At the farthest time-horizon, we see increases that are similar to those estimated within the new construction segment. However, the shorter-term picture looks different: we find marginally significant evidence of price declines on the order of 2–4% between the launch year and $T = 2$. Taking seriously the idea that impacts in existing stock are likely to reflect demand-side pressures, especially in the near term when a filtering process has not yet had time to play out, this looks like removals under SC are associated with a negative demand shock that temporarily reduces prices. Our findings in the new construction segment suggest that it takes a couple of years for reduced output to translate into sharply higher prices, and here the uptick by $T = 3$ seems consistent with the short-term demand side impact being more than counterbalanced by the spillover impact of reduced construction output within several years.

If this interpretation is accurate, one would expect to see sharper price declines in regions where undocumented residents are more prevalent. Therefore, as an indirect test of the mechanism we outline, we split our sample by census tract. Using ACS 5-year estimates from the 2006-2010 vintage, we compare the path of prices in tracts with above median reported LEFB share (∼1.0%) to tracts with no reported LEFB residents (roughly 1/4 of all tracts). The estimation in Panel B of Figure 10 confirms the prediction: we find no evidence of price declines in tracts least likely to contain LEFB residents (blue line), and larger declines in tracts with above-average LEFB share. In both instances, the evidence is again consistent with supply reductions filtering through to existing stock prices within three years.

Online Appendix Figure S6 repeats this estimation in the ACS set of counties. As always, our estimation horizon is one year shorter in this sample, and this may make patterns somewhat less clear to interpret. We still see evidence largely consistent with temporary declines in above-median LEFB tracts, and in these tracts we also see a patterns that is suggestive of a reversal. In contrast to the national sample, we also find price declines in non-LEFB tracts. Over this limited horizon, we do not see a price reversal beginning in these tracts. In general, the ACS sample is comprised of larger counties, and more often contains cities. The smaller differences between no-LEFB and above-average-LEFB tracts may reflect more neighborhood-to-neighborhood mobility in urban environments than in rural ones, leading to stronger price linkages and spillovers between the two sets of tracts. We do not further explore this possibility, and several obvious avenues for future research remain.

5.3.3 Do SC-Associated Population Increases Appear to Explain Higher Prices?

As noted in section 5.2.3, SC appears to be causally linked to population increases in the national sample. In our quantity analysis, this suggests that reduced-form sensitivity between construction labor supply and homebuilding may be understated without taking population changes into account. For a price analysis, however, population increases would tend to push in the other direction, raising the possibility that this channel may affect prices in both new construction and existing-stock segments. We adopt the following approach to assess the extent to which population changes seem to be driving changes in price. We construct an endogenous control variable by conditioning on outcomes: we look at the population growth that actually occurs in each county during the years of SC variation (2009-2012), and then assign each county one of three tercile indicators: high-, medium-, and low-growth. We use this variable in two steps. First, we validate that including this variable (as a set of FE indicators interacted with year) will entirely explain the estimated treatment effect on population in Figure A5(a). This is anticipated and unsurprising: it is the desired consequence of the endogenous variable construction. Then, having shown that including this variable will absorb population changes correlated with treatment, we repeat our price regressions with the same inclusion, and look at the extent to which estimated price impacts change. Large changes would imply a meaningful role for a total population channel to affect price. Small changes imply that population changes cannot be driving price impacts.

Figure A5 shows the results which, in general, suggest that population changes arising from SC have quite limited scope for explaining the price patterns we document. Panel (b) repeats the estimation within the new construction segment for the national sample while including tercile controls for realized growth. This generates almost no change in our estimated impact. Panel (c) includes these controls in the estimation for existing housing stock by tract-LEFB-share, and again finds results that are extremely similar. Figure S7 in our Online Appendix repeats this analysis in the ACS subsample, and shows that controlling for population growth has a similar lack of impact on price patterns.

5.4 Wages

The declines that we document in both construction employment and homebuilding are extremely persistent, which represents a potential puzzle. If we assume that homebuilders optimize over intensity of construction before the SC shock, then what shifts should be expected given an exogenous shock to immigration enforcement? If increased immigration enforcement itself does not meaningfully change the optimal number of homes to build in a given region, then we would assume that builders would attempt to attract workers into the construction sector in sufficient quantity to return to the prior level of employment. Absent a pool of unemployed workers easily enticed into the construction sector (which our ex-post results showing net declines in employment seem to rule out), basic economic theory suggests that static demand paired with reduced labor supply would place upwards pressure on wages. In this section, we test for evidence of such increases.

There are, however, several possible stories that would predict a lack of increases. Perhaps builders currently earn zero economic profits and therefore cannot profitably increase wages. This seems a difficult story to square with patterns of sharply increasing home prices during this time – not to mention evidence in this paper showing that SC increases new construction transaction prices. Another possibility is that increased immigration enforcement changes the builder's optimization. Although we do not have any direct evidence that speaks to this, we note that other literature has tended to find small effects of SC on factors like crime (Miles and Cox 2014, Hines and Peri 2019), and moderate effects on labor markets (East et al. 2018, East and Velásquez 2022) that would not seem to motivate a large shift in optimal homebuilding – again given the overall backdrop of scarcity in US housing markets. A third possibility is that homebuilders have monopsony power in local markets and therefore choose not to raise wages. This could intersect with another potential explanation: that workers' elasticity of substitution for switching into the construction sector might be very small. That would mean that in order to attract large numbers of additional construction workers into the sector, builders would have to raise wages by such a significant amount that it wouldn't be worth it. This would be a profit-based explanation that doesn't require zero economic profit, and indeed is a claim that one can certainly hear large homebuilders make.

To test the response of wages, we use detailed industry data on construction costs as well as standard measures from ACS. Our best measure of wages in the construction sector comes from RSMeans, which provides a cost-estimation platform to the homebuilding industry. As described in Section 4.5, RSMeans extensively surveys builders in local markets about actual construction labor costs each year. They use this information to produce an index of labor cost relative to the national average by location-year. A limitation of this data is that the series begins in 2007. In our estimations, we assign the 2007 figures to 2005 and 2006. This weakens our ability to rigorously test for parallel pretrends with RSMeans data; however, our analyses using ACS data suggest that these hold well.

Other scholars have documented that Secure Communities is associated with absolute reductions in wages (East et al. 2018). For purposes of attracting additional labor to a given sector, the relative level of wages is most relevant. Therefore, we normalize the RSMeans series with a county-level all-industry wage index produced from ACS microdata:

$$RSM\ normalized_{it} = log\left[RSM_{it}\left(\frac{ACS\ aggregate\ county\ pay_{it}}{ACS\ aggregate\ county\ hours_{it}}\right)^{-1}\right] \tag{3}$$

We also explore wage response by subpopulation using ACS data on average weekly wages by LEFB and US-born status. We compute average hourly wage for construction workers in the subpopulation of interest, and again normalize this index with ACS county-level average wages in the same way as equation 3.

Figure 11 shows results. The top-left panel shows that SC does not appear to affect the level of construction wages either at launch or during the subsequent year. At $T = 2$, we see an increase of 2.3% which is marginally significant. The top-right panel shows that relatively flat construction wages do, in fact, translate to relative increases as the overall wage index declines. We find steadily increasing construction wages in relative terms, from 1.3% in the launch year, to 2.0% at $T = 1$ and 4.6% at $T = 2$. The two panels in the bottom row show that these increases are primarily driven by increases for domestic workers. Using ACS-reported data on wages, the bottom-left panel shows no discernible impact on LEFB wages. The bottom-right panel shows increases for US-born workers, which also peak at 4.6% for $T = 2$. In the ACS data, estimated responses are less statistically significant, however the pattern shown is quite similar to the top panel using RSMeans.

This evidence suggests that as workforce shortages mount in the years following SC launch, builders do increase relative wages in the construction sector. However, the majority of the response appears to arise from keeping construction wages level while average wages are driven down through the channels described in East et al. (2018). And further, as our evidence shows, the magnitude of these increases is clearly not sufficient to attract a full complement of replacement workers, nor is it sufficient to forestall sharp reductions in construction output. Against a longstanding national backdrop of housing shortages and given that SC does not decrease overall local population, this suggests that there is some friction within the construction industry that generates wage rigidity. Although it is hard to believe that builders are not making any economic profits – and indeed, our results showing large increases in new construction sales prices suggest that the SC itself may provide builders some extra room to raise wages – the lack of observed wage increases may reflect an equilibrium outcome based on builders' beliefs. Builders may perceive that the market-clearing wage required to attain 100% replacement by domestic labor would be too high to be profitable. That belief (regardless of its veracity) could lead builders to forgo any attempt to raise wages and to reduce activity instead. We believe that exploration of this potential mechanism is a fruitful area for future research.

6 Conclusion

We show that: (i) when undocumented workers are deported, domestic labor only partially fills vacated construction jobs; and an apparent complementarity between immigrant and domestic labor leads to net job losses for US-born workers, (ii) residential construction output is highly sensitive to these declines in labor supply, (iii) the resultant reduction in homebuilding leads to higher home prices, with demand-driven declines in prices being small and highly transient.

We exploit the staggered rollout of additional immigration enforcement under the Secure Communities program to identify shocks to the labor force that are plausibly exogenous to local housing market conditions. We empirically document a first stage: using several proxies for undocumented residents, SC does lead to a reduction in the amount of labor supplied to the construction sector. This effect is heterogeneous by occupation: we show that declines in immigrant labor supplied to low-skilled occupations are partially offset by increases in domestic labor supplied. Within higher-skilled occupations, we find that (smallish) declines in immigrant labor supplied are matched with even larger declines for US workers. Our interpretation is that within residential construction, low-skilled labor is a complement to high-skilled labor. Because domestic labor only partially replaces lost immigrant labor, SC leads to a net decline of low-skilled labor, and that in turn leads to a reduction in total labor demanded.

We show that negative labor shocks are associated with reductions in homebuilding, using two measures: both planned future construction (permits) and realized construction (observed sales of new homes). In both measures, SC leads to an economically and statistically large slowdown in construction activity. We also show how this reduction in housing supply affects the prices of new homes. The quality-adjusted price of newly constructed homes increases following SC. We also show evidence of endogenous adjustment on home size: builders appear to build slightly smaller homes. These shifts lead to muted evidence of increases in price when considering raw transaction prices that do not control for housing stock features. Prices of existing housing stock also increase over a three-year horizon. We find some evidence of temporary declines during the initial two years following SC rollout; these appear to be concentrated in tracts with above-median minority share.

We find some evidence of upwards pressure on wages for US-born workers. The literature has documented that SC leads to reductions in wages. Using an index of construction labor cost sourced directly from local unions and construction companies, we find largely static level wages for construction workers. In conjunction with declines in the overall wage index, this means that relative construction wages increase. These wage gains are fairly small, with a peak impact of slightly less than 5% by the 2nd year after SC launch. Assuming that the immigration shock does not change the optimal amount of construction for firms – which seems broadly reasonable given longstanding housing stock shortages and no total population declines in treated counties – the lack of more robust wage adjustment to attract more labor seems surprising and worth further study.

Housing supply in the United States has been starkly lower than average for most of the past two decades. This paper provides novel evidence on a new channel that has substantial impact on housing supply: shortages in labor supplied to the residential construction sector. One immediate implication of this is that zoning reform or streamlining, often touted as a necessary step for increasing housing supply, may not be sufficient: even if localities stand ready to build more homes, this paper suggests that important frictions may exist in finding the individuals to do the building. This paper also shows that immigration policy, along with other interventions that directly affect domestic labor supply, are important levers for policymakers interested in overall home affordability. In addition, this paper suggests that the overall impact of increased immigration enforcement and deportation programs would be a reduction in housing supply and an increase in house prices, as supply-side impacts appear to dominate any reduction in prices coming from decreased demand.

References

Albouy, David, Gabriel Ehrlich, and Yingyi Liu, 2016, Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis, Technical Report w22816, National Bureau of Economic Research, Cambridge, MA.

Alsan, Marcella, and Crystal S Yang, 2022, Fear and the safety net: Evidence from secure communities, *Review of Economics and Statistics* 1–45.

Altonji, Joseph G., and David Card, 1991, The Effects of Immigration on the Labor Market Outcomes of Less-Skilled Natives, in John M. Abowd, and Richard B. Freeman, eds., *Immigration, trade, and the labor market*, A National Bureau of Economic Research project report (University of Chicago Press, Chicago).

Borjas, George J., 2017, The Wage Impact of the *Marielitos* : A Reappraisal, *ILR Review* 70, 1077–1110.

Borjas, George J., and Joan Monras, 2017, The labour market consequences of refugee supply shocks, *Economic Policy* 32, 361–413.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, 2021, Revisiting event study designs: Robust and efficient estimation, *arXiv preprint arXiv:2108.12419* .

Callaway, Brantly, and Pedro HC Sant'Anna, 2021, Difference-in-differences with multiple time periods, *Journal of Econometrics* 225, 200–230.

Card, David, 1990, The Impact of the Mariel Boatlift on the Miami Labor Market, *Industrial And Labor Relations Review* .

Card, David, 2001, Immigrant Inflows, Native Outflows, and the Local Labor Market Impacts of Higher Immigration, *Journal of Labor Economics* 19, 22–64.

Cohen-Goldner, Sarit, and M. Daniele Paserman, 2011, The dynamic impact of immigration on natives' labor market outcomes: Evidence from Israel, *European Economic Review* 55, 1027–1045.

Cortes, Patricia, 2008, The Effect of Low-Skilled Immigration on U.S. Prices: Evidence from CPI Data, *Journal of Political Economy* 116, 381–422.

Cox, Adam B, and Thomas J Miles, 2013, Policing immigration, *The University of Chicago Law Review* 80, 87–136.

De Chaisemartin, Clément, and Xavier D'Haultfoeuille, 2022, Two-way fixed effects and differences-in-differences with heterogeneous treatment effects: A survey, Technical report, National Bureau of Economic Research.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille, 2020, Two-way fixed effects estimators with heterogeneous treatment effects, *American Economic Review* 110, 2964–2996.

East, Chloe N, Annie Laurie Hines, Philip Luck, Hani Mansour, and Andrea Velasquez, 2018, The labor market effects of immigration enforcement .

East, Chloe N, and Andrea Velásquez, 2022, Unintended consequences of immigration enforcement: Household services and high-educated mothers' work, *Journal of Human Resources* 0920–11197R1.

Friedberg, Rachel M., 2001, The Impact of Mass Migration on the Israeli Labor Market, *The Quarterly Journal of Economics* 116, 1373–1408, Publisher: Oxford University Press.

Gardner, John, 2022, Two-stage differences in differences, *arXiv preprint arXiv:2207.05943* .

Glaeser, Edward, and Joseph Gyourko, 2018, The Economic Implications of Housing Supply, *Journal of Economic Perspectives* 32, 3–30.

Glaeser, Edward L, and Joseph Gyourko, 2003, The Impact of Building Restrictions on Housing Affordability, *FRBNY Economic Policy Review* 21–39.

Goodman-Bacon, Andrew, 2021, Difference-in-differences with variation in treatment timing, *Journal of Econometrics* 225, 254–277.

Grittner, Amanda M, and Matthew S Johnson, 2021, Complaint-driven regulation and working conditions: Evidence from immigration enforcement .

Helsley, Robert W., and William C. Strange, 1995, Strategic growth controls, *Regional Science and Urban Economics* 25, 435–460.

Hines, Annie Laurie, and Giovanni Peri, 2019, Immigrants' deportations, local crime and police effectiveness .

Hunt, Jennifer, 1992, The Impact of the 1962 Repatriates from Algeria on the French Labor Market, *Industrial and Labor Relations Review* 45, 556–572, Publisher: Sage Publications, Inc.

Ihlanfeldt, Keith R., 2007, The effect of land use regulation on housing and land prices, *Journal of Urban Economics* 61, 420–435.

Jackson, Kristoffer (Kip), 2018, Regulation, land constraints, and california's boom and bust, *Regional Science and Urban Economics* 68, 130–147.

Khater, Sam, Len Kiefer, Ajita Atreya, and Venkataramana Yanamandra, 2018, The Major Challenge of Inadequate U.S. Housing Supply, *Freddie Mac Insight* .

Lach, Saul, 2007, Immigration and Prices, *Journal of Political Economy* 115, 548–587.

Malpezzi, Stephen, and Richard K. Green, 1996, What Has Happened to the Bottom of the US Housing Market?, *Urban Studies* 33, 1807–1820.

Miles, Thomas J, and Adam B Cox, 2014, Does immigration enforcement reduce crime? evidence from secure communities, *The Journal of Law and Economics* 57, 937–973.

Molloy, Raven, 2020, The effect of housing supply regulation on housing affordability: A review, *Regional Science and Urban Economics* 80, 103350.

Monras, Joan, 2020, Immigration and Wage Dynamics: Evidence from the Mexican Peso Crisis, *Journal of Political Economy* .

Ohls, James C, 1975, Public policy toward low income housing and filtering in housing markets, *Journal of Urban Economics* 2, 144–171.

Ortalo-Magné, François, and Andrea Prat, 2014, On the Political Economy of Urban Growth: Homeownership versus Affordability, *American Economic Journal: Microeconomics* 6, 154–181.

Rosenthal, Stuart S, 2014, Are private markets and filtering a viable source of low-income housing? estimates from a "repeat income" model, *American Economic Review* 104, 687–706.

Saiz, Albert, 2003, Room in the Kitchen for the Melting Pot: Immigration and Rental Prices, *Review of Economics and Statistics* 85, 502–521.

Saiz, Albert, 2007, Immigration and housing rents in American cities, *Journal of Urban Economics* 61, 345–371.

Saiz, Albert, and Susan Wachter, 2011, Immigration and the Neighborhood, *American Economic Journal: Economic Policy* 3, 169–188.

Sun, Liyang, and Sarah Abraham, 2021, Estimating dynamic treatment effects in event studies with heterogeneous treatment effects, *Journal of Econometrics* 225, 175–199.

Svajlenka, Nicole, 2021, Undocumented immigrints in construction, *Center for American Progress White Paper* .

Sá, Filipa, 2015, Immigration and House Prices in the UK, *The Economic Journal* 125, 1393–1424.

Van Hook, Jennifer, and James Bachmeier, 2013, Citizenship reporting in the american community survey, *Demographic Research* 29, 1–32.

Zabel, Jeffrey, and Maurice Dalton, 2011, The impact of minimum lot size regulations on house prices in Eastern Massachusetts, *Regional Science and Urban Economics* 41, 571–583.

Figure 1: New Construction in the US (population-adjusted)

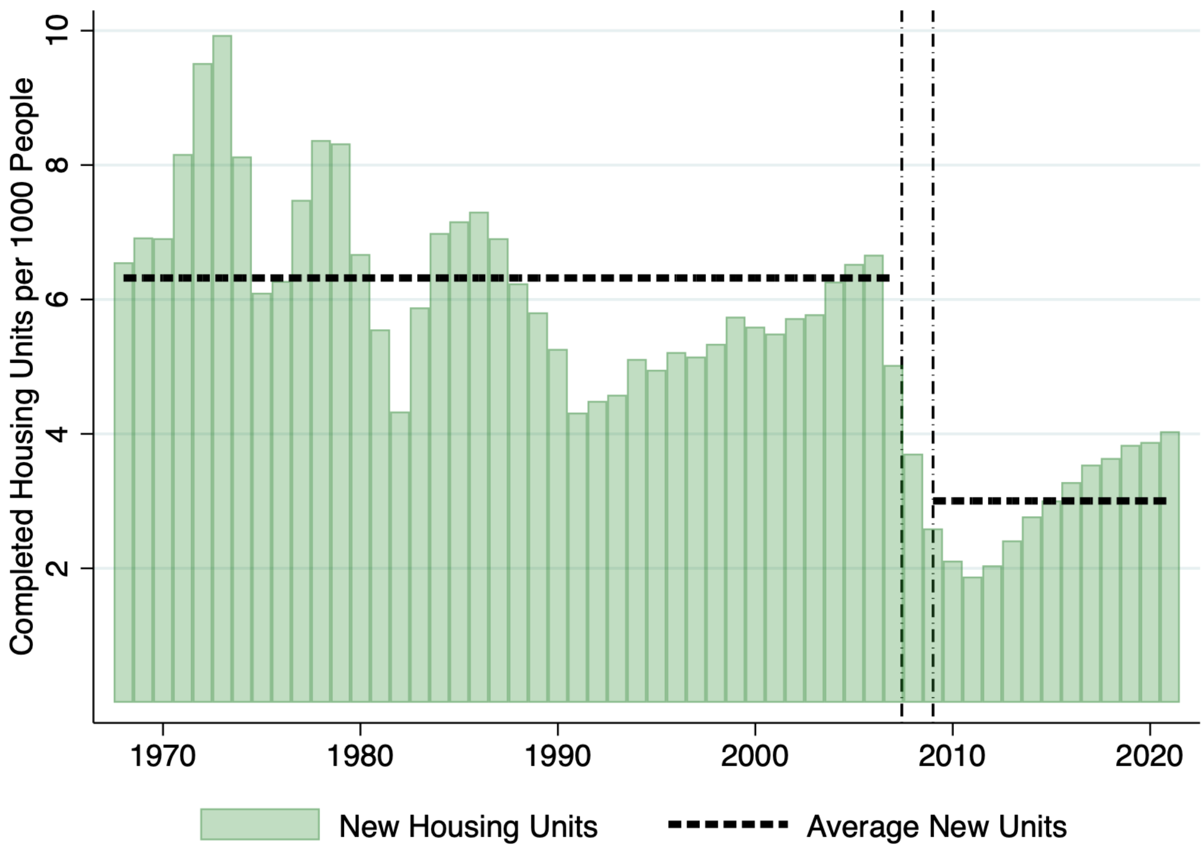

Note: This figure plots the time trends of constructed new housing units. The green bars are the annual new housing units per 1000 population (left axis) in the US from Census Bureau and HUD. The two dashed lines indicates the average levels of new housing units per 1000 population pre-GFC (1968–2007) and post-GFC (2009–2021).

Figure 2: Staggered Rollout of Secure Communities

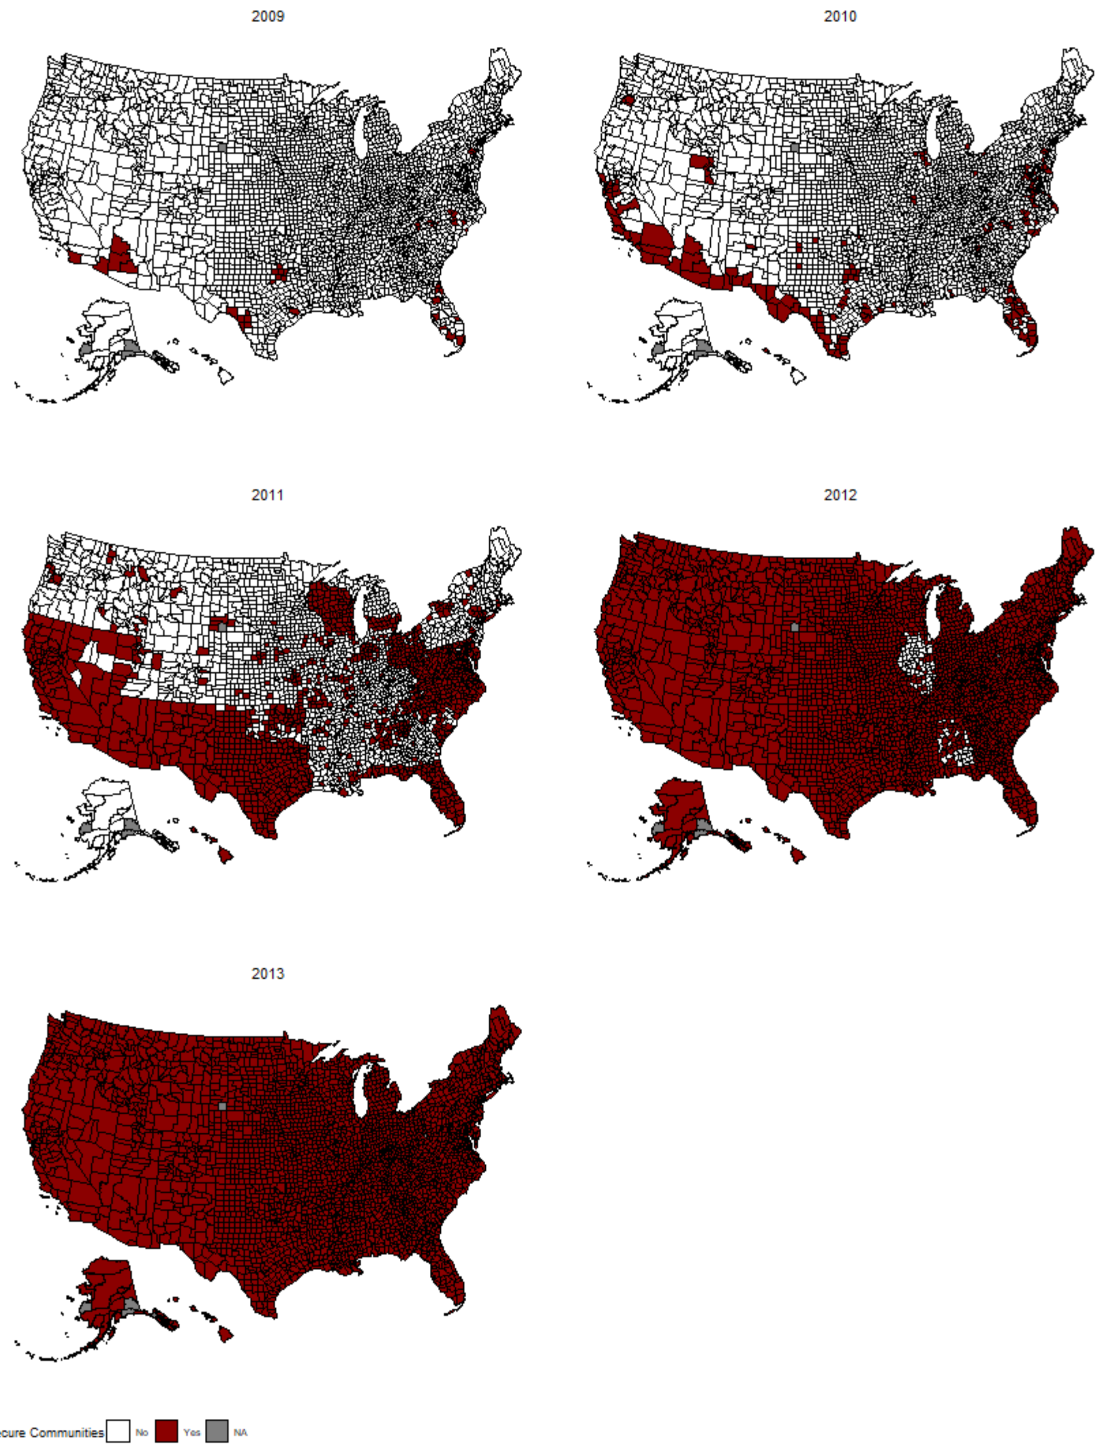

Note: Each panel of this figure shows the counties that implement SC within each year. This map reflects the treatment indicator used in our regressions, which assigns binary treatment status to any county operationalizing SC for at least half the year. Counties launching SC in, for instance, December of year t would therefore be coded as untreated in year t and treated in year $t+1$. Online Appendix Figure S3 shows treatment status by county-year using only the year of adoption without any consideration of how late in the year implementation started.

Figure 3: Cumulative Deportations under Secure Communities through 2013

(a) Along Calendar Years

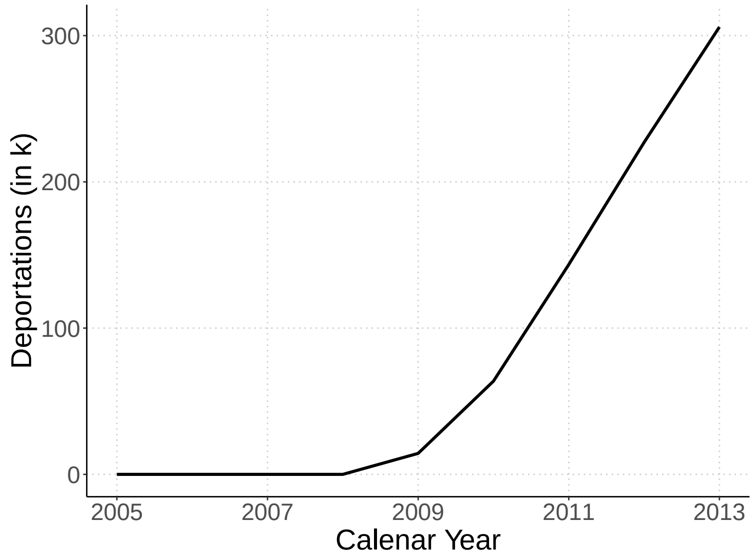

(b) Along Event Years

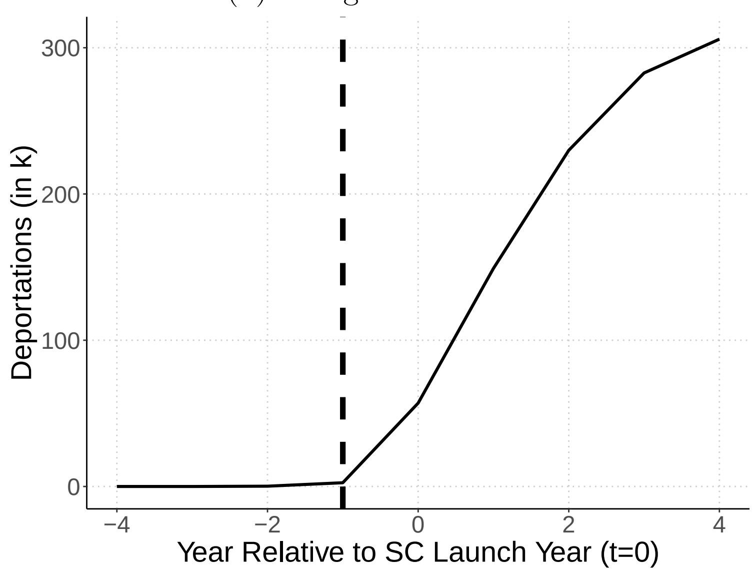

Note: This figure shows the cumulative number of individuals removed under Secure Communities. Data includes all removals through the end of fiscal year 2013 (Sept 30th, 2013), but does not include removals attributed to Secure Communities from 2014 onward. Removal counts by county-year are publicly available from TRAC at Syracuse University, and the figures are compiled from individual-level Department of Homeland Security administrative records. The top panel shows cumulative removals under SC by calendar year (top panel), and the bottom panel shows cumulative removals relative to the launch year.

Figure 4: Construction Workforce Impact of Secure Communities

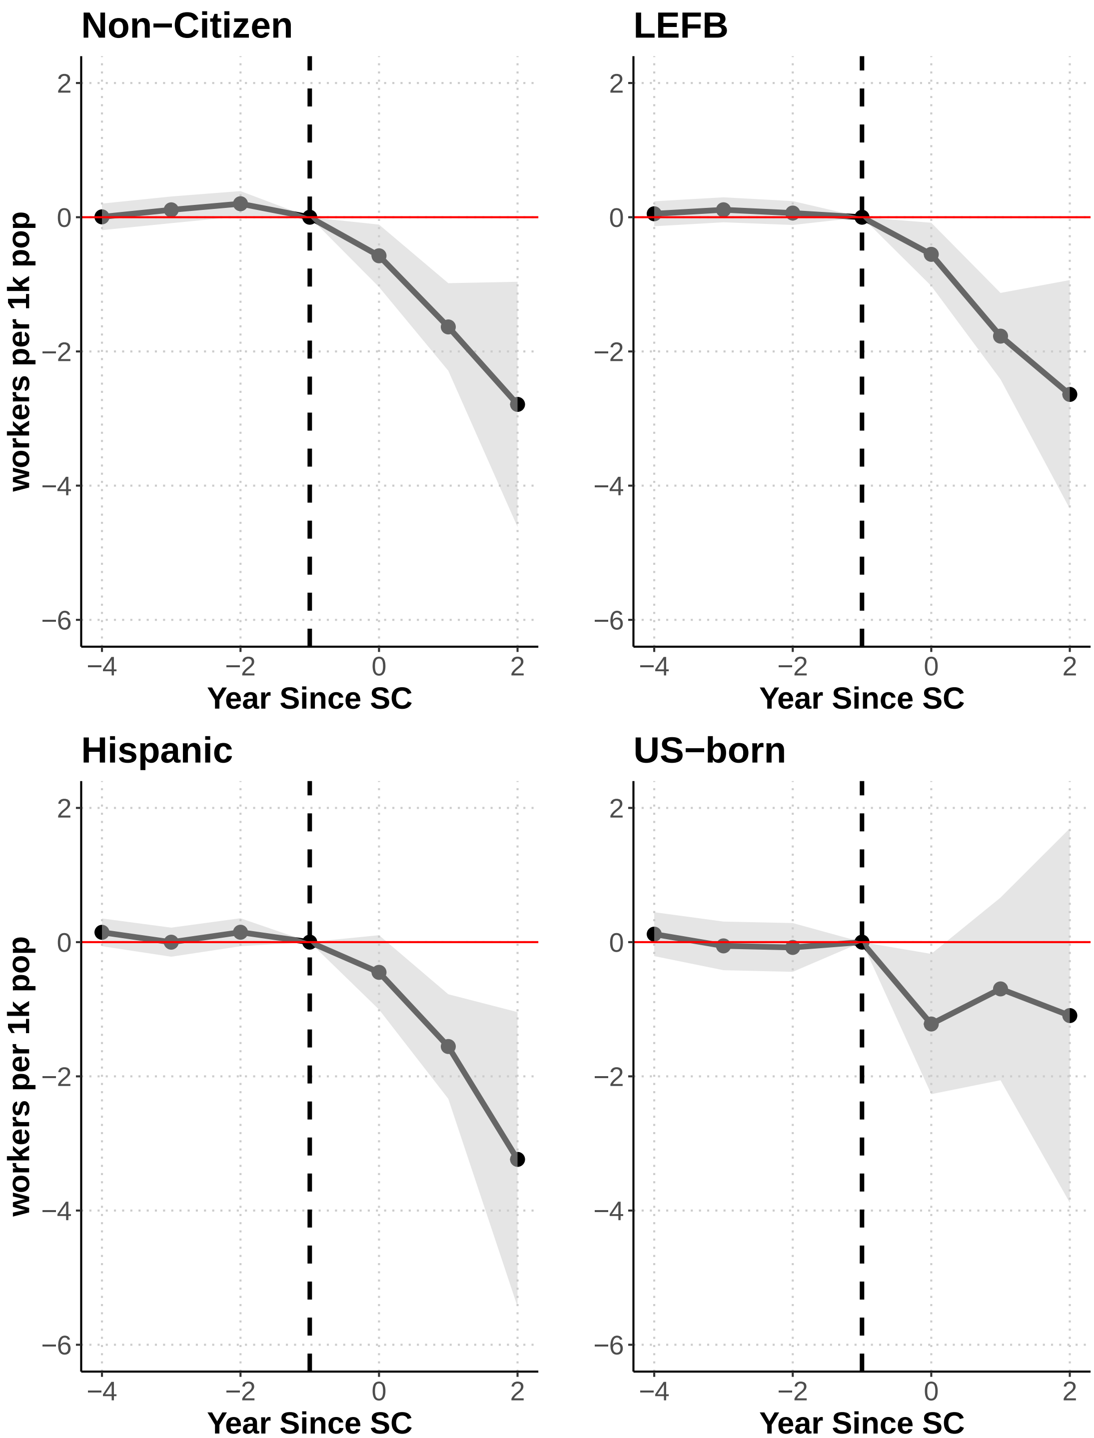

Note: This figure plots the impact of SC on construction employment, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on noncitizen, low-education and foreign-born (LEFB), Hispanic, and US-born workers in the construction sector. We normalize the number of workers by 2005 county-wide population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure 5: Construction-LEFB Impact by Cohort

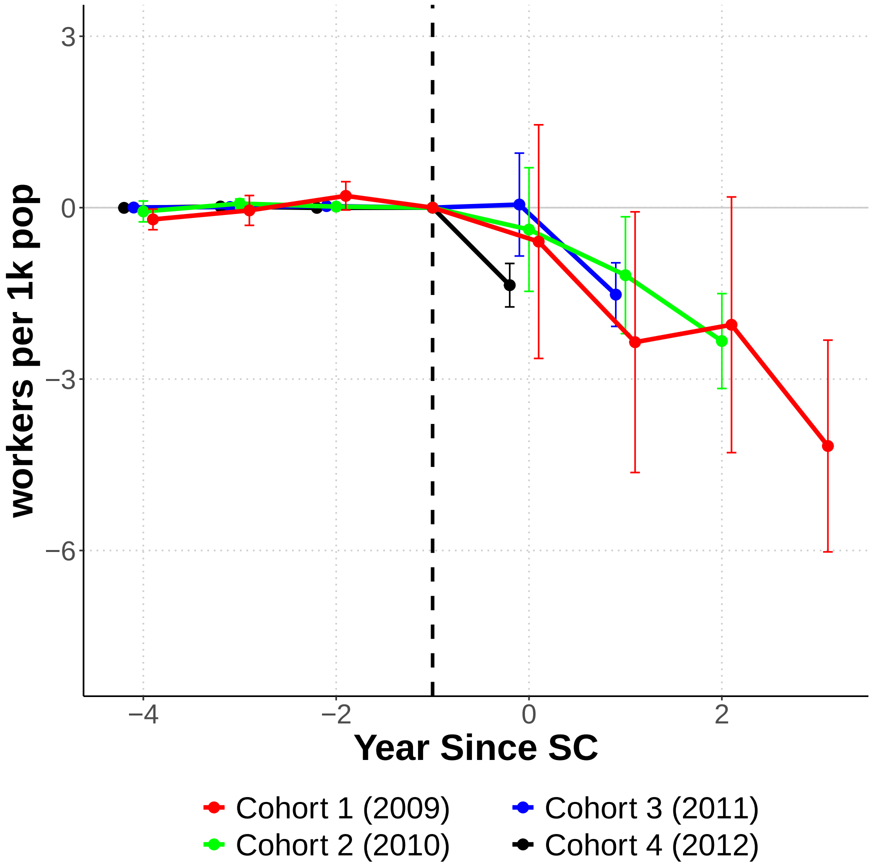

Note: This figure plots the impact of SC on LEFB construction employment by treatment-cohort, with the approach of Gardner (2022) and specification (2). We use a balanced (and constant) control group of the last counties to be treated. We plot the estimates for cohorts treated in 2009, 2010, 2011, and 2012, respectively, using the set of counties still untreated in 2012 as the control group for all regressions. We normalize the number of workers by 2005 county-wide population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure 6: Workforce Impact By Skill

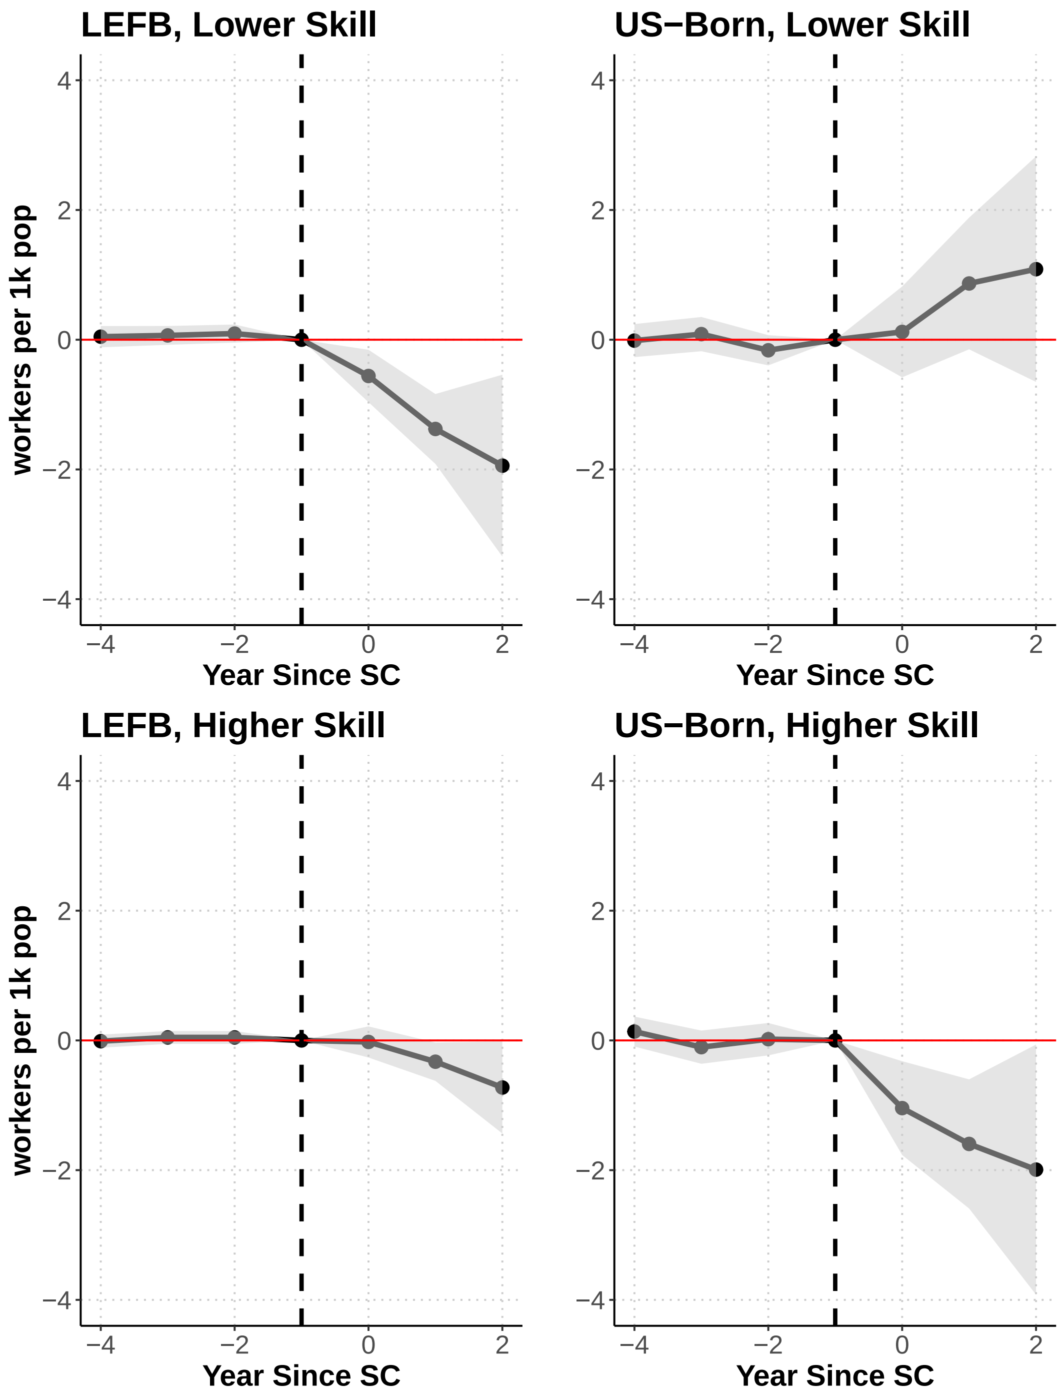

Note: This figure plots the impact of SC on LEFB and US-born construction employment by skill classification, with the approach of Gardner (2022) and specification (2). We define the following occupations as lower skill, construction laborers, helpers in construction trades, painters and maintenance workers, drywall installers, carpenters, and roofers. All remaining occupations are regarded as higher skill. The four panels plot the estimate impact on LEFB-lower skill, US-born-lower skill, LEFB-higher skill, and US-born-higher skill workers in the construction sector. We normalize the number of workers by 2005 county-wide population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure 7: Total Permits (Intended Construction)

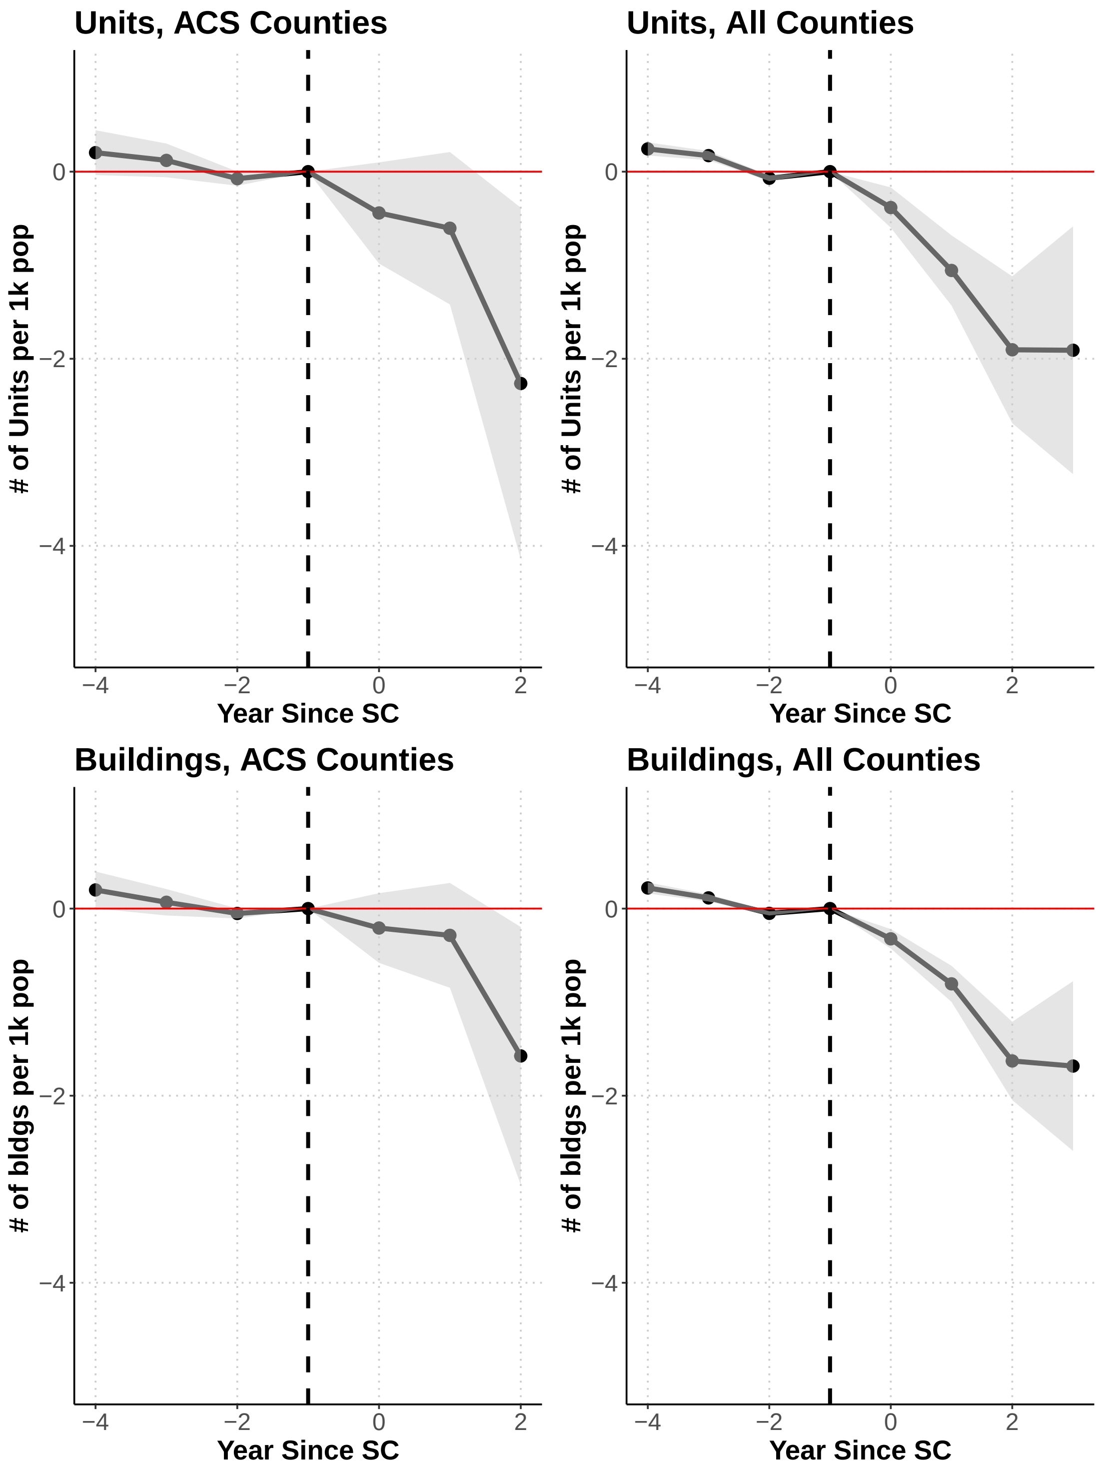

Note: This figure plots the impact of SC on residential construction activity measured by residential permits (intended construction), with the approach of Gardner (2022) and specification (2). The left columns examines the impact within ACS subsample counties, and right within the national sample. The top panels examine the impact on total permitted units, and the bottom examines permitted buildings. The outcome variable is total permits normalized by 2005 county-wide population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure 8: New Construction Entering Market

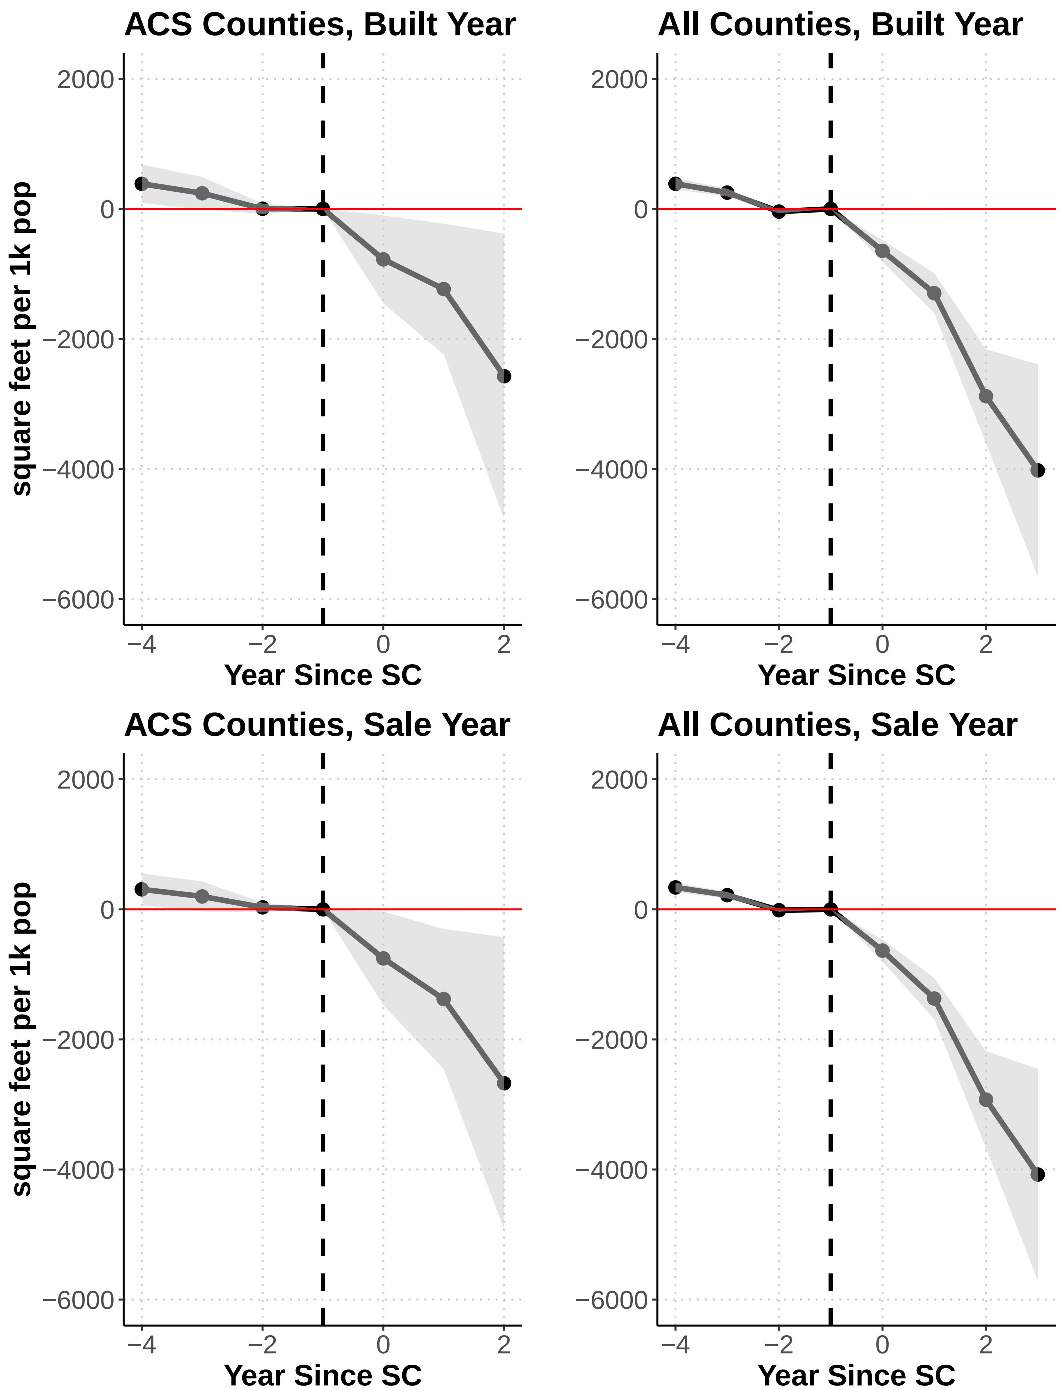

Note: This figure plots the impact of SC on residential construction activity measured by observed new construction (completed new construction), with the approach of Gardner (2022) and specification (2). Completed new construction is aggregated to county-year level using administrative tax-roll data from CoreLogic. The top row examines the impact on new construction measured based on build-year (preferred measure), and the bottom row uses sale-year. Results on the left are based on the national sample, and results on the right are based on the subset of counties separately identifiable in the ACS microdata. The outcome variable is total square footage normalized by 2005 county-wide population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of period is a year.

Figure 9: Price Response, Single-Family New Construction

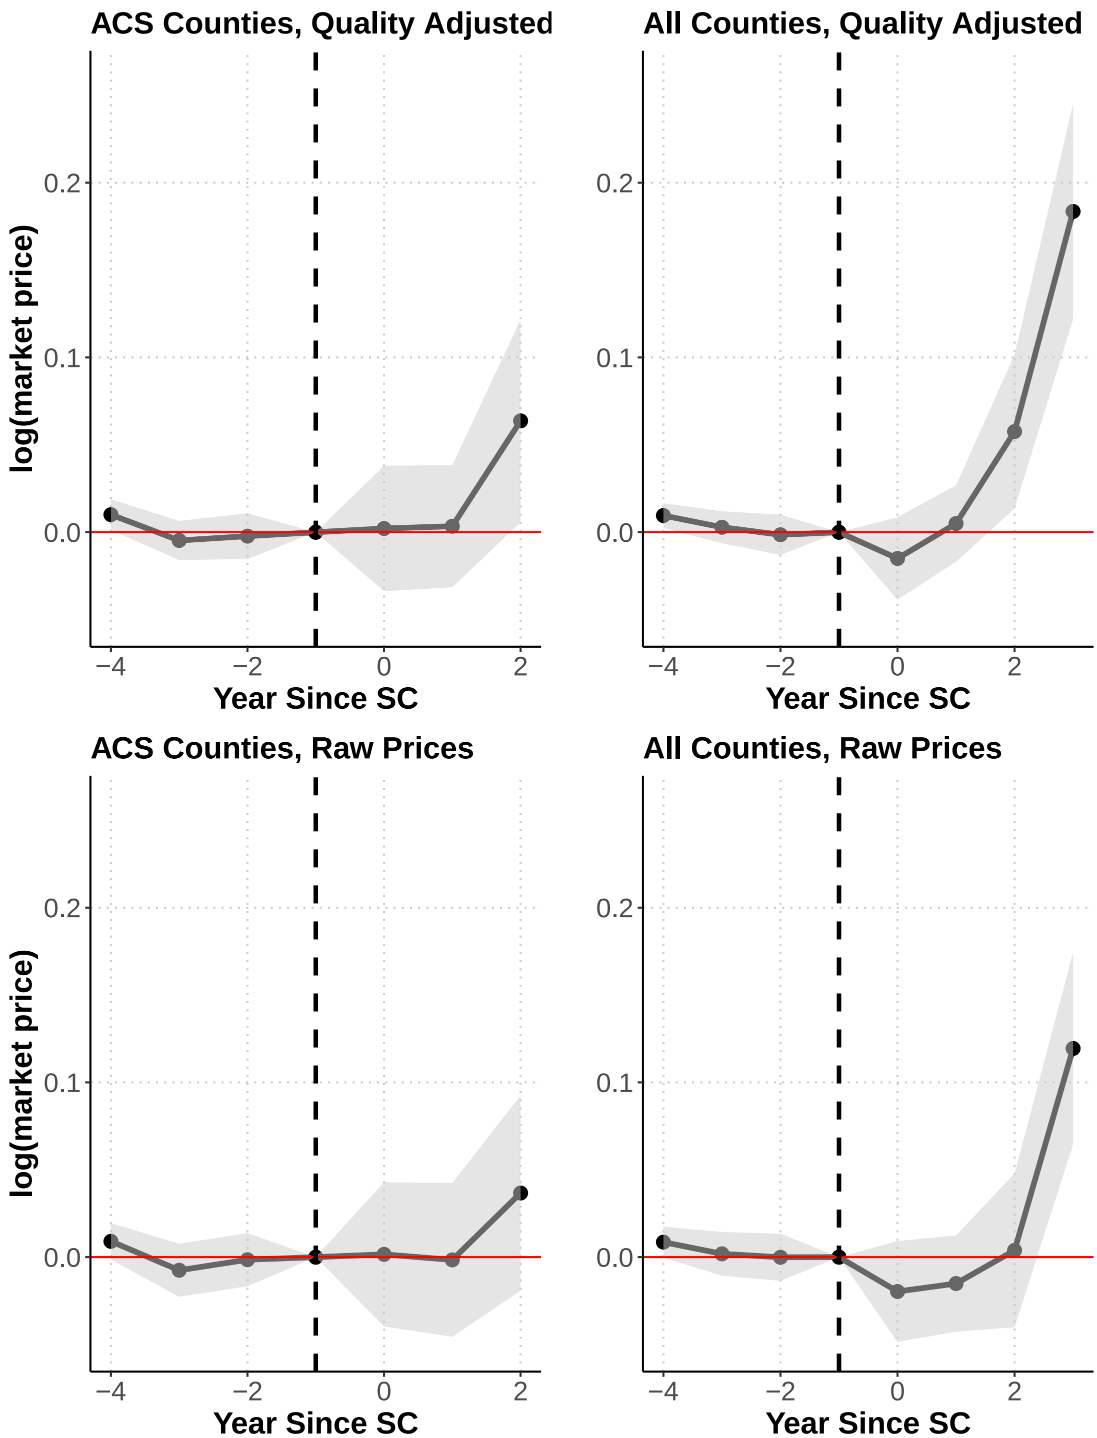

Note: This figure plots the impact of SC on single-family new construction prices, with the approach of Gardner (2022) and specification (2). The top row shows the effect on single-family homes without attribute controls. The bottom row includes hedonic controls to show the impact on quality-adjusted prices. Results on the left are based on the subset of counties separately identifiable in the ACS microdata, and results on the right are based on the national sample. The outcome variable is the natural log of recorded market price. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure 10: Resale Price Response, Single-Family Homes

(a) Average Price Response

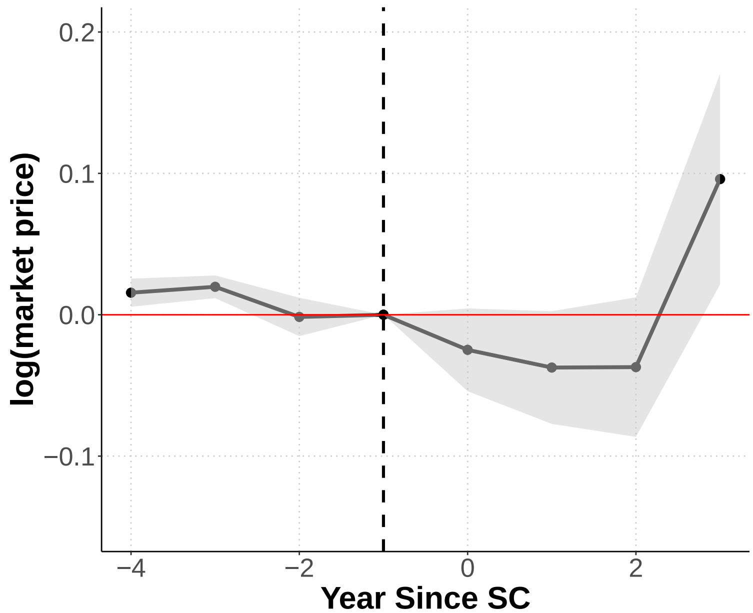

(b) By Tract LEFB Share

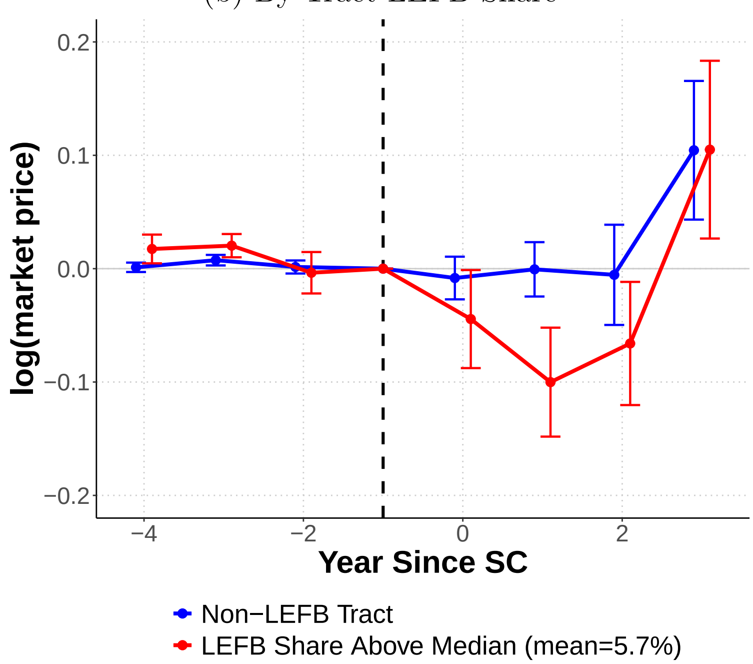

Note: This figure plots the impact of SC on the resale prices of the existing single-family home stock, with the approach of Gardner (2022) and specification (2). Results are based on the national sample. The top panel shows the average effect. The bottom panel shows the price responses for tracts with 0% LEFB share (blue), and for tracts with above median LEFB share (red). The outcome variable is the natural log of recorded market price. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure 11: SC Effect on Wages

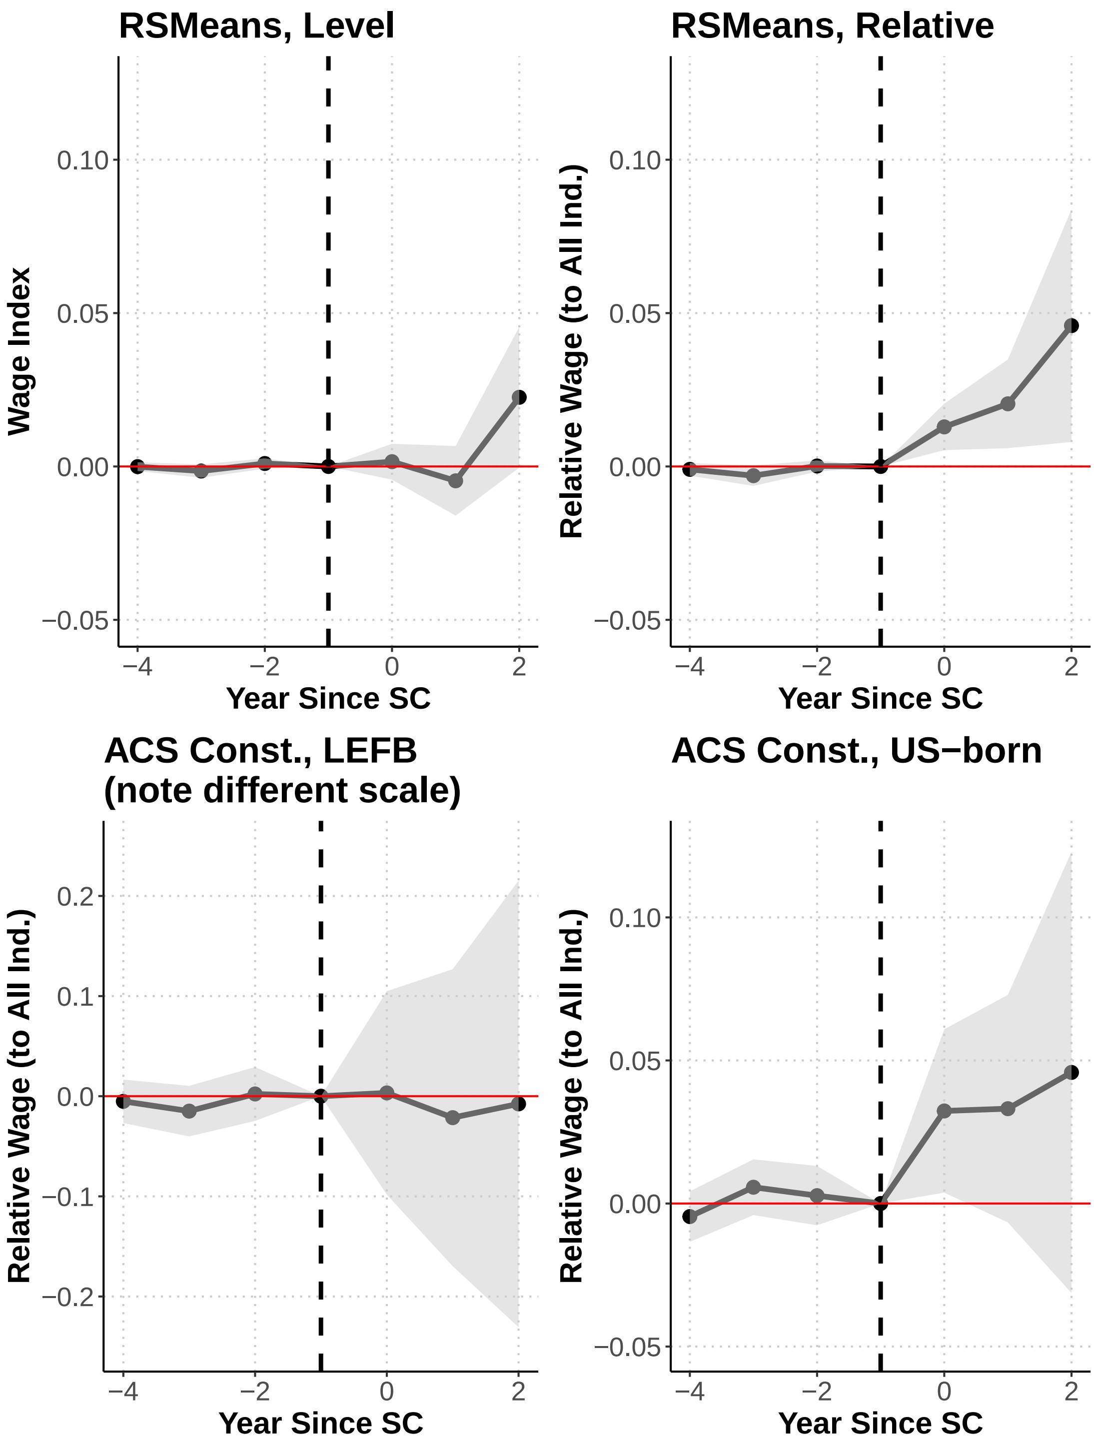

Note: This figure plots the impact of SC on average county-level construction wages. We use the approach of Gardner (2022) and specification (2). The top panel shows the level impact on RSMeans' regional wage index, normalized by county-level all-industry wages from ACS. The middle and bottom panels show the effect on ACS construction wages (again relative to all-industry wages) for LEFB and US-born workers respectively. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Table 1: Summary Statistics

	ACS Counties				All Counties			
	Mean	SD	Median	Obs.	Mean	SD	Median	Obs.
Construction Labor as Share of 2005 Total Population (%)								
Non-citizen	0.50	0.56	0.30	5,621				
LEFB	0.56	0.60	0.36	5,621				
Hispanic	0.71	0.83	0.40	5,621				
US-born	2.49	0.94	2.41	5,621				
Permits per 1k 2005 Total Population								
Buildings, Total	3.12	3.49	2.01	5,621	2.16	3.36	1.10	48,347
Units, Total	4.23	4.38	2.84	5,621	2.59	4.25	1.31	48,347
Construction per 1k 2005 Total Population								
Square Feet, Built Year	3959	6258	1843	2,648	1418	4399	42	23,886
Square Feet, Sale Year	4341	6310	2233	2,648	1548	4447	56	23,886
New Construction Transactions – Micro Data								
Market Price ($k)	316	205	264	1,345,620	302	200	250	2,071,526
Square Feet	2402	878	2238	1,345,620	2369	894	2198	2,071,526
# Bedrooms	3.58	0.81	4.00	1,345,620	3.53	0.80	3.00	2,071,526
# Bathrooms	2.89	0.82	3.00	1,345,620	2.84	0.83	3.00	2,071,526
Age at Sale	0.44	1.01	0.00	1,345,620	0.47	1.03	0.00	2,071,526
House Resale Transactions – Micro Data								
Market Price ($k)	336	331	240	5,541,935	310	306	222	8,675,762
Square Feet	1957	895	1750	5,541,935	1958	894	1753	8,675,762
# Bedrooms	3.28	0.83	3.00	5,541,935	3.26	0.82	3.00	8,675,762
# Bathrooms	2.38	0.94	2.00	5,541,935	2.38	0.94	2.00	8,675,762
Age at Sale	32.46	26.46	26.00	5,541,935	31.42	26.86	25.00	8,675,762
Hourly Wages ($)								
ACS Construction	19.67	4.68	19.22	5,621				
ACS Construction, LEFB	16.41	7.40	15.07	4,516				
ACS Construction, US-born	20.82	5.14	20.20	5,621				

Note: This table summarizes the county-level characteristics. Column "ACS Counties" uses only the subset of counties separately identifiable in the ACS microdata, while Column "All Counties" includes all counties. We report the mean, standard deviation, median, and number of observations, of the normalized population and construction employment, for each group of non-citizen, LEFB, Hispanic, and US-born, normalized number of permitting buildings and units, normalized square feet of new construction sold, transaction values and characteristics of both new construction and house sales, in addition to various wage measures.

Table 2: Is SC Rollout Predictable?

	Binary for Rollout						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Hispanic Share	0.6143***	0.5678***	0.6143***	0.5935***	0.6138***	0.6232***	0.5904***
	(0.1503)	(0.1509)	(0.1512)	(0.1509)	(0.1504)	(0.1506)	(0.1512)
Total Pop (M)		0.0325					
		(0.0312)					
3Yr Pop Growth, US			0.0011				
			(0.2840)				
3Yr Pop Growth, Hisp				−0.0464*			
				(0.0238)			
3Yr Pop Growth, LEFB					−0.0033		
					(0.0320)		
3Yr NC Growth						−0.0028	
						(0.0026)	
01-07 Price Runup							0.1035
							(0.0710)
Observations	962	962	962	962	962	934	962
R^2	0.4971	0.4993	0.4971	0.4981	0.4972	0.4900	0.4987

*p<0.1; **p<0.05; ***p<0.01

Note: This table explores predictors of rollout. The dependent variable is a binary indicator for whether rollout occurs in a given county. The dataset is stacked over rollout years 2009-2012. Each stack codes counties launching SC in that year as 1 and counties that have not yet launched SC as 0. Counties that have already launched are excluded from a given stack. The regression includes Census-region-by-stack fixed effects, and standard errors are clustered at the county-level.

A Appendix

Figure A1: Population Growth by Census Region

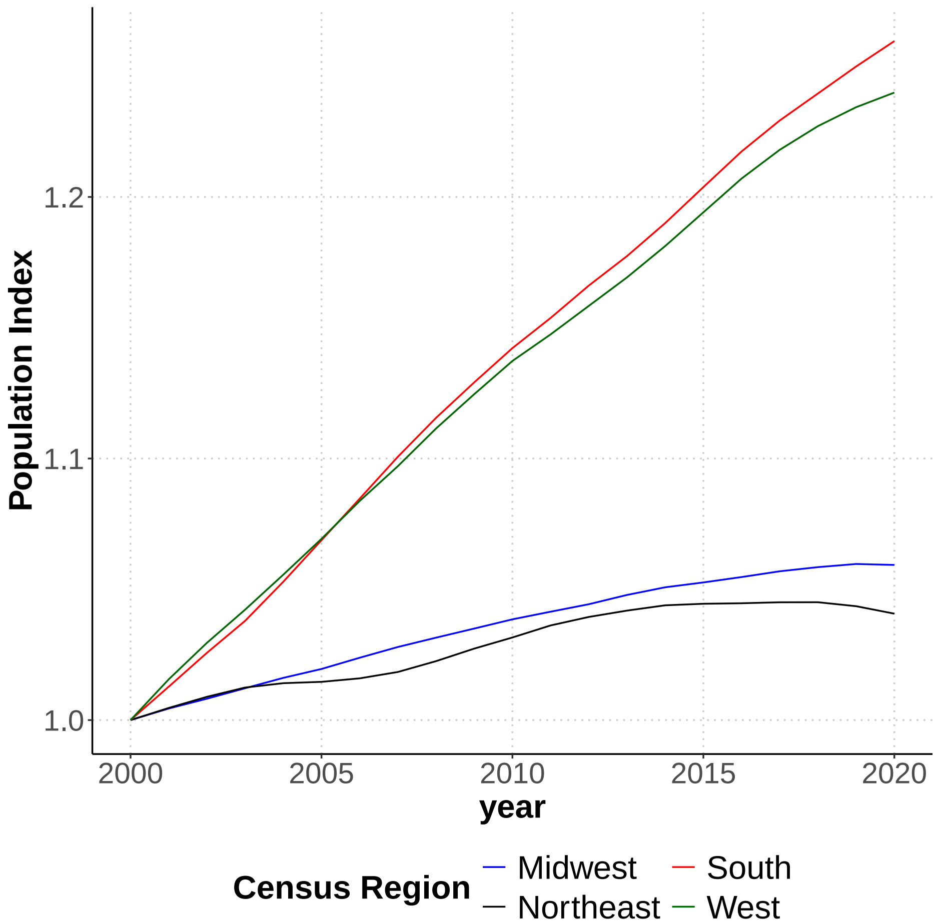

Note: This figure shows population by Census Region, indexed to 2000. Data comes from US Census American Community Survey.

Figure A2: Total Labor Supplied to Construction Industry

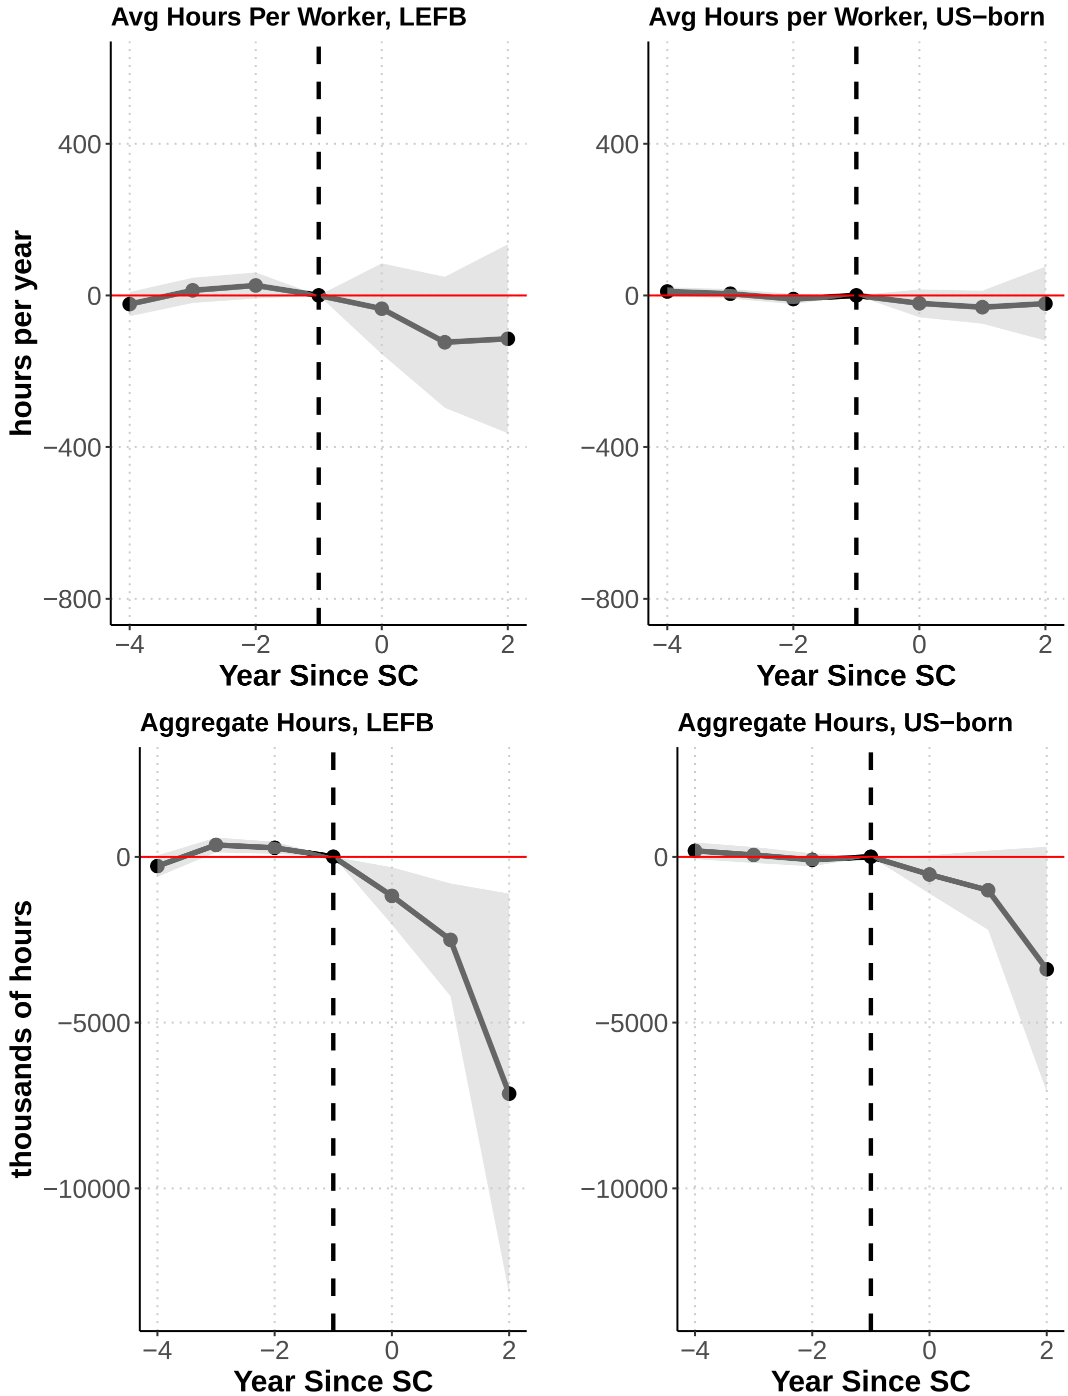

Note: This figure plots the impact of SC on the intensive margin of construction employment, with the approach of Gardner (2022) and specification (2). Using ACS data, we compute the average hours worked per worker in each subpopulation. The top panels shows average working hours per LEFB worker, and the bottom shows US-born workers. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure A3: New Construction Attribute Shifts, All Counties

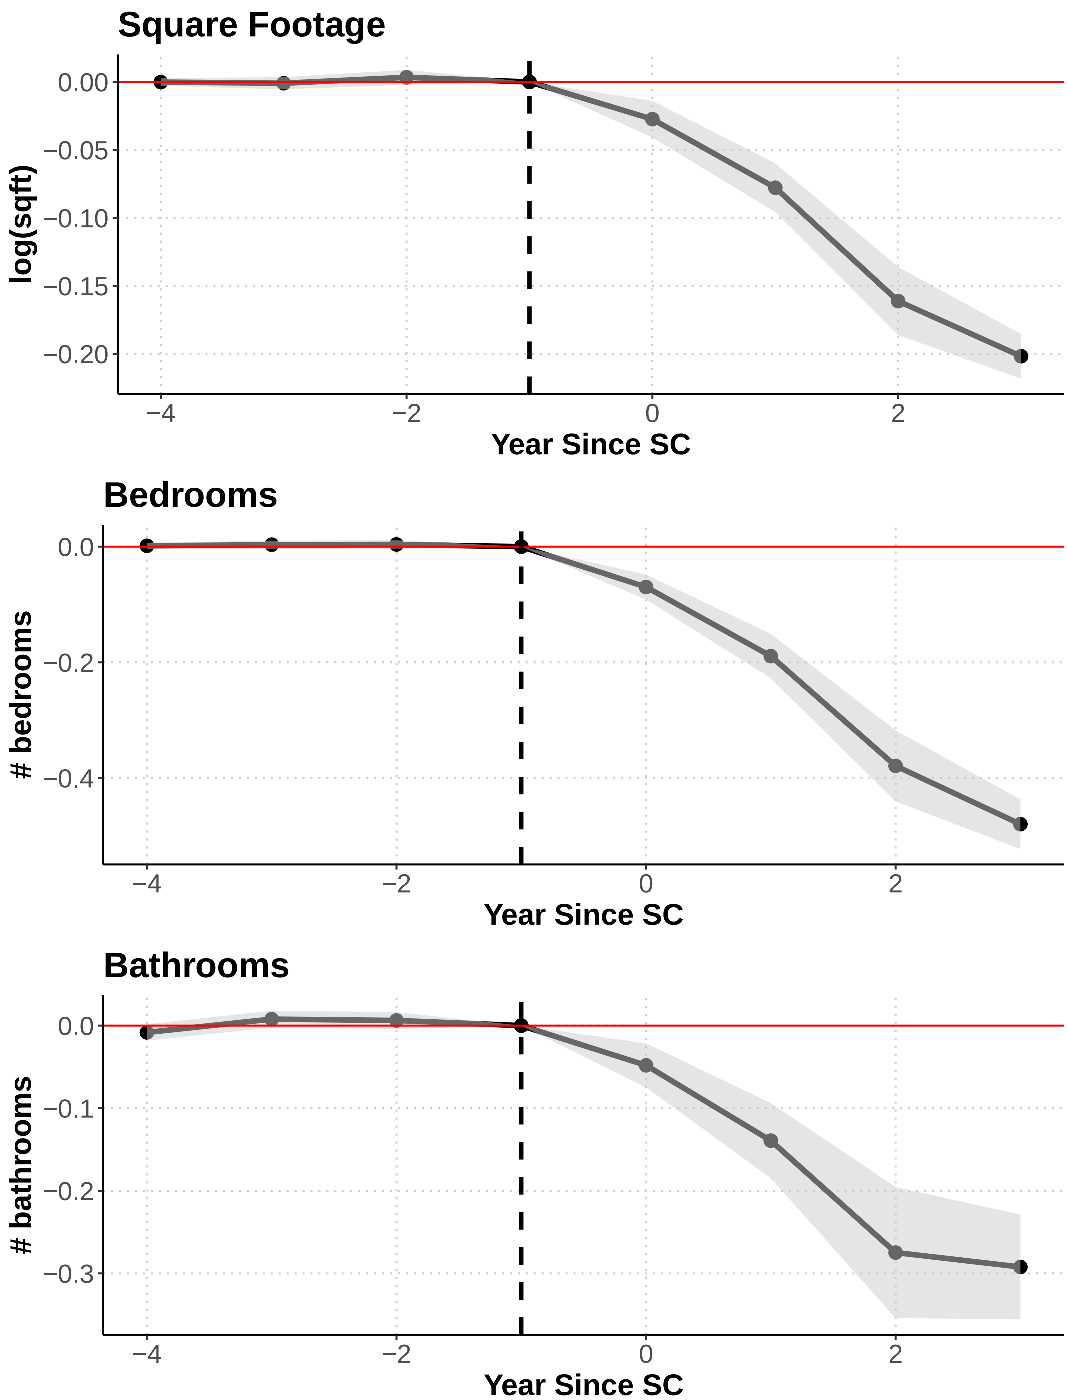

Note: This figure plots the impact of SC on various hedonic attributes of newly constructed properties, with the approach of Gardner (2022) and specification (2). Estimations in this figure are based on the full national sample. In each panel, 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure A4: Population Impact of Secure Communities

(a) ACS Counties

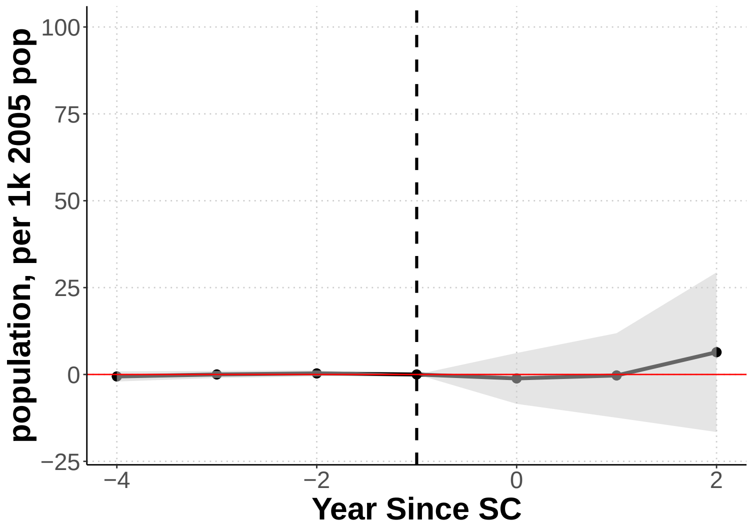

(b) All Counties

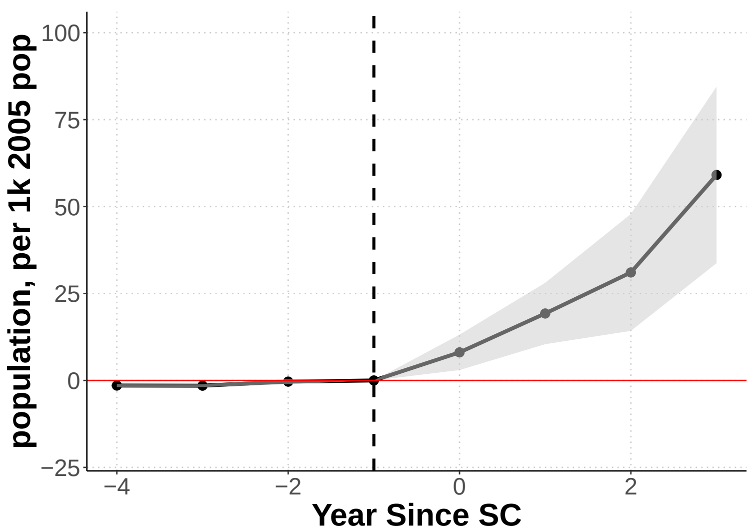

Note: This figure plots the impact of SC on county population, with the approach of Gardner (2022) and specification (2). The top panel is based on the subset of counties separately identifiable in the ACS microdata, and the bottom panel is based on the national sample. We normalize the population by 2005 level population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.

Figure A5: House Price Reponses After Controlling Population Growth

(a) Population (Control Population Growth)

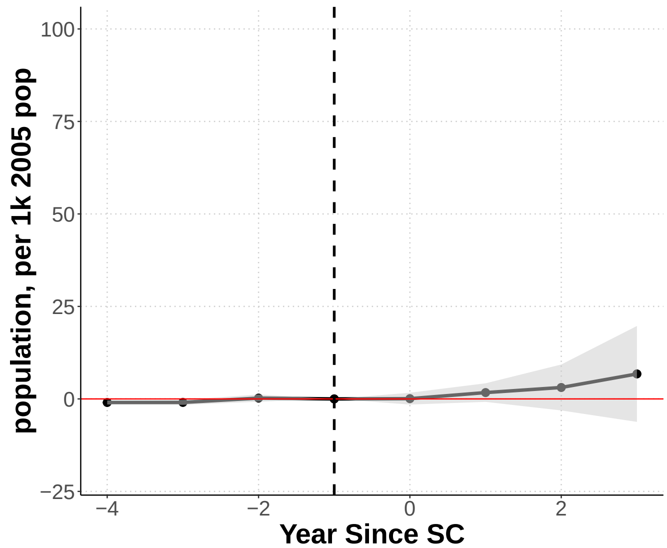

(b) New Construction Price (Control Population Growth)

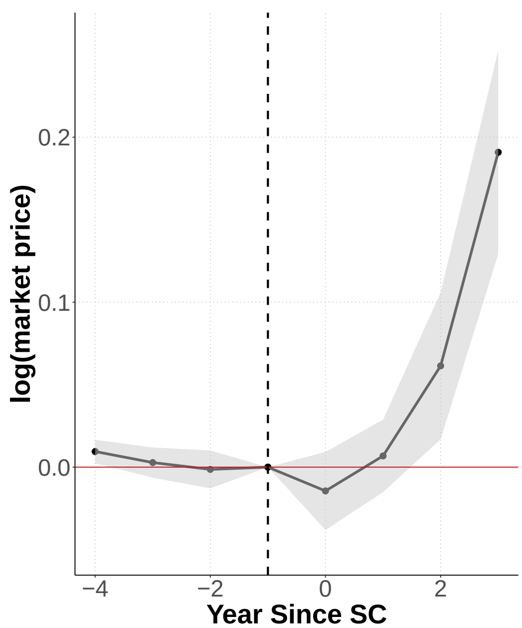

(c) Resale House Price (Control Population Growth)

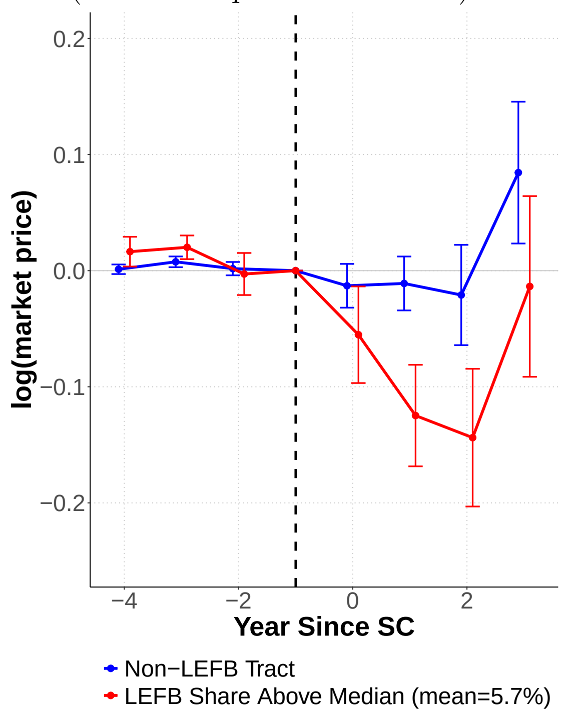

Note: This figure repeats several estimations with the inclusion of growth-tercile-by-year fixed effects. Results are based on the national sample. Panel (a) estimates the SC effect on population (paralleling the bottom of Figure A4). Panel (b) shows the SC effect on new construction prices (paralleling the top right of Figure 9). Panel (c) shows the house resale price responses for tracts with 0% LEFB share (blue), and for tracts with above median share (red) (paralleling the bottom of Figure 10). 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.