How Do Labor Shortages Affect Residential Construction and Housing Affordability?*

Troup Howard† Mengqi Wang‡ Dayin Zhang§

February 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 a less explored channel affecting housing supply: shortages of construction labor. 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: deporting undocumented construction workers reduces labor supplied by domestic construction workers on both extensive and intensive margins.

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, Albert Saiz, and Christopher Timmins; 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), the 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.

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

‡University of Wisconsin-Madison. Email: mengqi.wang@wisc.edu.

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

1 Introduction

The United States has faced a secular shortage of housing supply in the past decade. Since 2011, the US added an average of 1.1 million new housing units every year, which is 30% lower than the long-run equilibrium before the Great Recession and 34% lower than the annual new construction demand estimated by Freddie Mac (Khater et al. 2018). As Figure 1 shows, despite a gradual recovery in homebuilding since 2011, the level of annually completed housing units reached by 2021 corresponds only to the lowest point over the preceding five decades. While housing underproduction is widely perceived as a key contributor to rapid and sustained growth of home prices over the prior decade, the extensive academic literature exploring the drivers of low homebuilding has tended to focus on the distortionary effects of excessive housing regulations (i.e. zoning and building codes) as the central factor in limiting housing supply (Glaeser and Gyourko 2018, Molloy 2020). This paper provides novel evidence documenting that labor supply is also an important channel affecting housing supply.

To recover the causal relationship between labor supply and homebuilding, we need a shock to regional construction workforces that is otherwise unrelated to the set of local economic conditions that determine housing supply in equilibrium. Our setting 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.[1] 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. Other scholars have documented large impacts on local population and employment. As the residential construction sector is well known to be a large source of employment for undocumented immigrants (Svajlenka 2021), we use the population shock of Secure Communities as a laboratory for exploring the impact of reductions in regional construction workforce. We present findings along four main dimensions.

We begin by documenting the first-stage impact on construction workforces. 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

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

other industries or from outside the labor force) fail to fully offset employment losses, but in fact, immigration enforcement leads to reductions in total construction workforce and in average hours worked for US-born workers also. We show that this effect is heterogeneous by skill. For lower-skilled occupations, reductions in undocumented workers are partially offset by increases in domestic labor supply. However, employment in higher-skilled occupations is reduced for both undocumented immigrants and US citizens.

Our second set of results focuses on new construction quantity. We show that 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 has foregone the equivalent of an entire year's worth of additional residential construction: 2,423 fewer buildings are permitted, and 1,997 fewer newly constructed homes enter the market.

Third, we empirically document the anticipated link between reduced quantities and increased prices. Several factors complicate this analysis. First, it is clear that increased immigration enforcement will have a demand-side impact as well as a supply-side effect: at the most mechanical level, an increase in deportations may mean fewer residents demanding housing services. However, by focusing on new construction, we restrict focus to a market segment that is relatively less likely to include undocumented buyers. We also find that SC is not associated with declines in total county-level population, ruling out a channel whereby a shrinking population directly reduces demand for homes. Second, the workforce impact of SC may be associated with endogenous shifts in housing characteristics that in turn affect prices. We control semiparametrically for hedonic characteristics of housing and show a large increase in the quality-adjusted price of new residential construction. Three years after SC implementation, the average new construction parcel is 17% more expensive relative to the baseline (after controlling for home attributes).

Finally, 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 three different wage measures to explore this relationship, including a region-by-year index of journeyman construction wages provided directly by local unions. We find a striking lack of any wage increases in the construction sector, either in level terms or in relative terms, across all sub-populations of construction workers. This evidence strongly suggests the existence of some friction or market failure in the construction sector that leads homebuilders to reduce output rather than increase wages.

Overall, our results show that housing supply 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 appear to offset workforce shocks by raising wages to attract additional workers. 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.

Furthermore, this paper shows that US homebuilding is highly sensitive to labor supplied by undocumented immigrants in particular. When construction jobs held by undocumented workers are vacated for exogenous reasons, domestic labor supplied to the construction sector is highly inelastic. Net reductions in labor supply induce a slowdown in building, and as a consequence negative shocks to undocumented labor supply translate into employment declines for US-born workers on both the extensive and intensive margin. This complementarity between undocumented workers and domestic labor suggests that immigration policy has first-order impact on both housing supply and home prices.

Related literature. This paper contributes to three important lines of literature. First, we shed new light on housing affordability trends in the US. 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. This paper complements the literature showing that lack of labor is another critical factor that can exacerbate housing underproduction.

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 (Molloy 2020), and Jewish immigration into Israel after the collapse of the Soviet Union (Friedberg 2001, Cohen-Goldner and Paserman 2011). The crucial issue is that immigrants are likely to be correlated with the economic trend of the local economies, so the cross-sectional estimation tends to be biased upward. Many of these 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.

Compared with studies focusing on labor market outcomes, very few papers have looked at the impact of immigration policy on the product market. Two exceptions are Lach (2007) and Cortes (2008), which show immigration flows are associated with lower prices for nontradable goods and services, using previous immigration labor share as instrumental variables. Our paper focuses on housing, the largest component of durable consumption for households. We provide a detailed picture of the effects along both quantity and price dimensions.

Third, this paper also contributes to the line of literature that studies the role of immigration in 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) and Saiz (2007) show more immigrants are associated with inflated housing rents. Our paper provides novel evidence that immigration flows have a first-order impact on housing supply. And we also document a link between this supply shock and home prices.

The rest of this paper proceeds as follows. Section 2 describes the institutional details of the immigration shock that we employ as a laboratory for studying how labor impacts housing supply. 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).[2] DHS would then match those fingerprints against an internal database of foreign-born individuals. A subset of individuals appearing in this database are 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 for 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 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. Appendix Figure A1 maps treatment cohorts using the actual date of SC launch, without consideration for when in the year that initial date falls.

[2] 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.

The last set of untreated counties adopted SC in January 2013.[3] 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 occasion 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.[4]

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. Allowing i to denote cross-sectional units and t time periods, the analysis is commonly implemented with two-way fixed-effects (TWFE) OLS:

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

In recent years, several papers have shown a potential for bias in the estimated causal

[3] 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.

[4] 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.

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 TWFE model necessarily places greater weight on potentially problematic pairings that use already-treated units for the comparison 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 strongly 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.[5]

Several papers have presented estimators that can address these concerns with the standard 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 (cross-sectional and yearly) from untreated units. Practically, this means that increasing treatment effects over time will not erroneously shift cross-sectional averages, nor will time fixed effects late in our sample rest primarily on outcomes in treated regions. Purged of problematic variation, these estimated fixed effects are then used to produce fitted values which residualize

[5] 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.

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?

A glance at the rollout maps shown in Figure 2 quickly suggests 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 shows that county size does not meaningfully predict rollout. Column 2 confirms the existing finding from the literature that 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 3 onward, we retain Hispanic share as a control variable, and test for the marginal effect of other predictors. Column 3 shows that SC rollout is associated with a three-year growth trend in US-born population. The standard deviation of growth in the sample is 4.4%, meaning that an increase of one standard deviation in US population growth correlates with a 3.7pp increase in SC launch in a given year. In Section 5, we do ultimately show that SC corresponds with an increase in US (and total) population which may be related to this margin of selection. In our setting, an increased propensity to implement SC in faster growing counties would tend to bias us against finding any decreases in construction as more people (all else equal) implies a need for more homes. Columns 4 and 5 show that growth in Hispanic and low-education foreign-born (LEFB) populations do 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) prior to launch (year $t-4$ to year $t-1$). We do not find any evidence that prior building predicts SC rollout either economically or statistically.[6] This means that our results are very 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. In addition, we find that initial rollout is more likely in places with a growing population – which would tend to bias us away from finding declines in homebuilding. 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.

4 Data

This section introduces datasets used in this study.

[6] 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.

4.1 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.[7] 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.[8] SC was implemented in all counties by 2013.[9] Among 3,126 counties, 2% adopted SC in 2009, 11% in 2010, 31% in 2011, 53% in 2012, and 3% in 2013.[10]

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.

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

[8] 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.

[9] 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).

[10] 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 three different measures of wages: (i) county-year-by-industry wage data produced by the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW), (ii) wage indexes constructed directly from ACS microdata, and (iii) 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. The QCEW measure draws upon administrative employer reporting to state unemployment insurance programs, while the ACS measure relies on employee survey response. While employing potentially richer data, the QCEW measures are not disaggregated by subpopulation. The ACS measure can be produced for each of our populations of interest. However, 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.[11] The RSMeans data is also obtained via survey: for each locality, 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).

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

5 Results

5.1 First-Stage Impact of Secure Communities

In this section we show that the first-stage impact of SC generates a relevant setting in which to explore the relationship between labor supply and homebuilding. We begin by estimating the impact of SC on total populations. With this initial focus on population, we demonstrate that our empirical strategy picks up the effect of SC. There is independent empirical support for this: administrative records of individuals deported under the program. Therefore, while we estimate the impact of SC on population, we already know that the null of no impact should be rejected because ICE reports statistics on individuals removed under SC. However, replicating this analysis in our setting allows us to demonstrate that our empirical proxies for

[11] Following the literature, we compute average wage as total income divided by total hours worked for those between 20 and 64 years old who worked at least half-time in the prior year. The concern is a precisely reported downward shift in the numerator (arising from less working overall) that is not matched by a downward shift in the denominator because reductions are not large enough to move an individual between weeks-worked bins.

likely-undocumented populations in the ACS data are reasonable, and also to evaluate the accuracy of our estimates by comparing implied magnitudes from our analysis with aggregate data from administrative records.

After validating our empirical approach by exploring population impact, we turn attention to the central proposition of our setting: that SC had a meaningful impact on construction workforce. This analysis cannot be paralleled with administrative records because no information on occupation exists for those deported under SC.[12] The ACS data contain 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 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

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

[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.

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. We demonstrate a first-stage impact on population and employment in this subsample. Then, when we move on to focus on second-stage construction outcomes, we will estimate effects both within the ACS-covered subset of counties, as well as the entire national sample. We find extremely similar patterns in both samples.

Our preferred specifications are event study versions of equation 1:

$$y_{it} = \alpha_i + \gamma_t + \sum_k \beta^k \mathbb{1}(time\_since\_treatment_{it} = k) + \epsilon_{it}. \tag{2}$$

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 setting, SC is implemented between 2009 and 2013, which means that we can estimate an impact for four periods: the contemporaneous effect plus three subsequent years.

5.1.1 Population

Figure 3 shows how SC affected overall population. Our independent variable is the log of the group share: log(group population / 2005 total county population). We use this log normalization because of meaningful heterogeneity in the baseline share of immigrant populations across counties. This specification implies that SC has a proportional rather than an additive impact on group shares. By using predetermined population from the start of our sample period, we ensure that shares are not affected by any total population changes that may arise from SC. To facilitate comparison between our estimates and nationwide figures, we weight these regressions with 2005 population. There are non-trivial econometric nuances both to normalization of the dependent variable and to the use of regression weights. We discuss these in detail in Section 2 of our Online Appendix.

The top two figures show the impact of SC on noncitizen and LEFB populations respectively. For both populations, we find no differential pretrends and clear declines pursuant to SC. The statistical significance of these declines is marginal, but the trend is clear. Focusing on the LEFB population, after three years, the point estimate suggests a peak decline in LEFB population of 4.67 percent. For the average county, this represents a decline of approximately 2,570 people. It is important to realize that such reductions will be driven not only by direct deportations through SC but also by several related channels. Facing increased immigration enforcement, some individuals may elect to voluntarily depart the country, and some who would otherwise have immigrated may elect not to. Additionally, when any individual is officially deported under SC, family members or close associates may also voluntarily leave the country alongside the deported individual.

How reasonable is the correspondence between our estimates and administrative estimates of removals under SC? The Transactional Records Access Clearinghouse (TRAC) center at Syracuse University has obtained records on SC under the Freedom of Information Act. Using aggregate data at the month-year level which TRAC makes public, we can compute that slightly over 250,000 people were removed under SC during the period from 2008 to 2012. This is the rollout period, so many counties were treated less than four full years. In 2013 – the only full year where all counties were treated before the policy transition from SC to PEP – 79,000 individuals were removed. The implied four-year total would be 315,000, though this may be understated if SC impact is declining in treatment duration. Our estimate of a 4.67 percent decline, applied to the total LEFB population for the entire county in ACS, implies a reduction of 1.12M individuals – or, in other words, implies that for each person removed under SC, another 2.5 people either leave the country voluntarily or refrain from emigrating in the first place.

The bottom right panel of Figure 3 shows that SC has little discernible impact on the overall Hispanic population. Of course this grouping includes large numbers of citizens, and also excludes many non-Hispanic undocumented immigrants. This grouping is of interest simply because a large share of undocumented workers in the construction sector has historically been Hispanic. While we subsequently show an unambiguous impact on Hispanic construction workers, Figure 3's lower right panel suggests that SC's impact on total Hispanic population is difficult to observe in aggregate. The bottom right figure estimates the impact on US-born population. This panel shows that SC is associated with population increases. As noted in Section 3.1, this is likely, at least partially, due to a correlation between SC rollout and medium-frequency population growth trends. It also may be the case that immigration enforcement has a causal impact on total population. If jobs vacated by deported individuals are filled by US citizens, and citizens have on average larger families than undocumented workers (many of whom have families that remain in their country-of-origin), this would also lead to population increases.

Because our central interest is in identifying a shock to construction workforce, we do not explore these population impacts further. Instead, we take the overall evidence of Figure 3 as clear evidence confirming that SC did affect regional populations and that implied reductions align order-of-magnitudewise with administrative records. We also note that because US-born population is much larger than immigrant populations, Figure 3 also provides strong evidence that total population does not decrease as a result of immigration enforcement. The LEFB population is definitionally disjoint from the US-born population (a distinction that does not apply to the population indicating Hispanic heritage) and increases in the number of US-born residents outweigh estimated declines for LEFB residents. This means that net population losses are not a potential explanation for any declines in residential construction that we find.

5.1.2 Workforce

In Figure 4 we estimate the impact on construction workers within the same subpopulations as before. Again, we normalize the number of workers by 2005 county-wide population. Here, as there is less heterogeneity of population shares conditional on the construction sector, we estimate impact in simple shares rather than log-shares to facilitate comparison across groups. 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 an increase during 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 8.4bps of the county population. For the median county in our data, this is equivalent to 412 fewer workers. The net impact of SC increases over time. At $T = 3$, we estimate the total reduction to be nearly half a percent of county population—2,185 workers for the median county. We highlight that our measure of workforce is not conditioned on employment. We are looking at the total number of workers in a given county-year who report any construction-related occupation. We return to intensive margin response later, but the precise interpretation of Figure 4 is 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, for the average county, estimated reductions in construction (2,185) are large relative to total population reductions (2,570). In the ACS data, approximately 9-12% of LEFB respondents report construction occupations. 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 occupations or workers leaving the labor force.

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 only counties untreated through 2012; and likewise for cohorts 3 and 4. 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 which, especially in light of the population results showing increases in US-born residents, is surprising. 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 a clear decline. Magnitudes are smaller: the peak effect, two years after treatment, represents a decline of approximately 725 workers for the average county. This is a strongly statistically significant result.[16]

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.[17]

[16] There is possibly some weak evidence to support a trend towards recovery at $t = 3$; however, this is not statistically significant. The relatively larger standard error for the latest event-study coefficient is common across most of our results and is likely due to having very few untreated units four years after rollout begins. As a consequence of the Gardner (2022) procedure, this means that late-sample γ_t fixed effects are estimated with less power, leading to larger standard errors in the second stage (event-time) estimate.

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

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 test this hypothesis by using ACS occupation codes to sort workers by skill.

We assign the following occupations 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.

Figure 6 shows event-study results by skill classification. Within lower-skill occupations, domestic labor appears to be a substitute for immigrant labor. The direct impact of SC is a reduction in the more-likely-undocumented grouping of LEFB construction workers. This is 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 17% and 45% in the periods following SC adoption. After three years, the net effect suggests a loss of approximately 1,050 workers in these lower-skilled occupations for the average county.

An opposite pattern holds within higher-skilled occupations. SC also causes reductions in high-skill LEFB employment; however, the impact is substantially smaller within this population. (The estimates are, in fact, statistical zeros; though a downward trend in the point estimates is still apparent.) Rather than experiencing any offsetting increase, higher-skilled US-born workers also see employment declines. 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. And 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 $T = 3$, our estimates imply a reduction of just under 1,500 higher-skilled workers for the average county; and 80% of the reduction comes from US workers.

Results on the intensive margin are consistent with domestic labor being a complement to undocumented labor. From the ACS data, we form a measure of average hours worked in each county. We compute total hours 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. Dividing by the number of people in a given workforce subpopulation (LEFB or US-born) gives us a measure of employment intensity for a subpopulation. Figure 7 shows the results. We find little impact on the LEFB population (although standard errors are large, so these are not precise zeros). We've already shown that there are fewer LEFB construction workers overall, but these results show that the remaining workers in this group are working the same amount, on average. By contrast, the bottom panel shows that US-born workers are working less. The reduced amount of work available for US-born construction workers appears to outstrip the decline in workforce, leading to fewer hours worked for the average US-born worker. This is strikingly consistent with a complementarity story. SC directly reduces LEFB construction workers, but there is an additional intensive margin effect in addition to the extensive margin reductions, suggesting an overall slowdown in the construction sector. In the next section, we document this slowdown in output directly.

5.2 Direct Evidence on Homebuilding

The prior section provides indirect evidence of reduced construction activity through reductions in overall construction employment. In this section, we explore the impact on residential homebuilding more directly. We focus on two measures of residential construction activity: intended construction (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 5.1.2 (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 8 shows the results for permitted buildings (top) and total permitted units (bottom). These two measures are very highly correlated ($\rho = .91$), and so results are quite similar between the two specifications.[18] Focusing on the ACS subsample at left: SC leads to a sharp reduction in both permitted building and total planned units. This effect is quite large. Focusing on the top panel of Figure 8, SC leads to .55 fewer buildings per 1,000 residents in the launch year, approximately

[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.

1 building per 1,000 residents in the next two, and 2.40 fewer buildings at $T = 3$. For the average county, this implies a total reduction of 2,423 buildings over three years. While the Great Recession complicates measurement of baseline activity, the average number of buildings permitted across counties in 2005 was 2,658, and the average number permitted across all county-years prior to SC implementation (this includes a number of post-housing-boom observations) is 1,333. This means that the three-year effect of SC corresponds to a year's worth of boom-time residential construction, or more than two-years of post-boom construction activity. This is a very large reduction relative to either baseline. Patterns in the national sample (right) are extremely similar to the ACS subsample.

The event studies in Figure 8 show some evidence of a statistically significant downward pre-trend, especially in the ACS subsample. This raises the concern that SC was initially rolled out in regions that were already experiencing declines in construction. Although nothing in the public discourse of SC implies this, that does not alleviate all concerns, as rollout being associated (deliberately or not) with any correlate of construction activity would be problematic. As we have discussed in Section 3.1, the inability of housing-boom-era construction or home price growth to predict rollout means that our results are unlikely to be driven by a cyclical pattern of overbuilding before SC and subsequent contraction. Another possibility is that statistically significant estimates in the preperiod represent anticipation effects: a public awareness of future immigration enforcement may induce contemporaneous response. Obtaining a permit is a forward-looking action, by definition, and so perhaps builders pull forward permitting activity (and construction plans) in anticipation of future labor shortfalls. The declines pursuant to SC treatment are larger than any preexisting downward trend would predict, but we acknowledge that the reader should be mindful of the possibility of confounding influences.

Figure 9 disaggregates the effect by building size for the ACS subsample. Single-family homes are the chief driver of the decline, which is perhaps unsurprising as single-unit buildings also represent the majority of housing stock in most places. It also becomes clear that a possible downward trend in the preperiod is a characteristic only of single-family homes. The parallel trends assumption seems to hold much more clearly for other classes of buildings. While we consider it very unlikely that SC rollout was explicitly a function of trends in one subclass of residential housing, the concern that rollout may have been influenced by some factor that correlates with single-family home construction remains.

No large treatment effect is evident for medium-size buildings. If anything, the event study suggests that SC is associated with increases in two-unit buildings, though magnitudes are small. This may point towards some margin of endogenous substitution by builders when labor becomes scarce. In Section 5.3 we show evidence of endogenous changes in housing stock attributes within property class, but we do not have evidence that speaks more directly to substitution between single family homes and multi-unit construction. No statistically significant response is evident for three- or four-unit buildings. For the largest buildings, declines are also evident. While the number of buildings is small – two to four fewer buildings per year – this may nonetheless reflect a meaningful impact on housing supply, since these buildings contain large numbers of units. Figure A2 repeats the estimation by building size but focuses on total units. For the largest buildings (bottom figure), the peak reduction occurs at T=1 and is approximately 25 units. Also of note is that the long-horizon estimate for total number of units is both economically and statistically zero. While the analysis shows a statistically meaningful reduction in the number of large buildings at T=3 (bottom graph of Figure 9), the evidence in Figure A2 suggests only a moderate reduction in total units across all large buildings being built. Again, this may reflect endogenous changes in the planned number of units per building occasioned by labor scarcity, but we do not empirically explore this channel further. Figures A3 and A4 show buildings by size-class and units by size-class, respectively, for the full national sample. Again, patterns are quite similar to those evident in the ACS subsample.

5.2.2 Observed New Construction

For several reasons, changes in permitting activity might not correspond to changes in actual construction. It is perhaps most likely that permitting activity overstates construction: builders may obtain a permit and subsequently decide to abandon the project due to downstream frictions like hiring labor or securing financing. In this case, our permit-based estimates would understate the effect of immigration enforcement. However, if there are changes in unpermitted construction, it could be the case that permitting declines overstate the true impact of labor shortages. 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 10 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. Once again, 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.[19]

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, peak between 4,000 and 4,500 sq. ft. per 1,000 residents. This is a flow measure of building, so to understand total magnitude over our estimation window we cumulate estimates between $T = 0$ and $T = 3$ and multiply by average county population. For the subset of ACS counties, the total implied reduction is just over 4.05M sq. ft. The median new home in ACS counties (prior to SC) is 2,031 sq. ft., which implies a total of 1,997 fewer homes built by $T = 3$. The corresponding figure based on permitting activity was 2,423, which is broadly similar. Higher observed responses in permitting may suggest that some permitting activity is speculative (one could imagine builders filing despite being less than certain a given project will proceed, for reasons of bureaucratic or timing efficiency), leading to a larger response than in the subset of projects that are completed.

Our preferred specifications in Figure 10 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 A5 in our 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

[19] In the next section, we will show direct evidence of shrinking time between new home completion and sale – a natural response to reduced supply. However, the effects suggest changes on the order of months, not years.

any observation missing square footage.

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 focus on new construction prices. The reasoning is twofold. First, this is the segment exhibiting a quantity response. Filtering theories in housing markets would lead us to expect price spillovers to existing housing stock. However, this link is both indirect and potentially realized on a longer timescale. Second, focusing on new construction allows us to consider a segment where demand-side shocks arising from SC are less likely to confound estimates. At the very most fundamental level, an increase in deportations within a given area mechanically means fewer residents demanding housing services.[20] In addition, other work has shown that SC leads to a range of economic spillovers (East et al. 2018, Miles and Cox 2014, Alsan and Yang 2022). However, 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 our price estimates.

We estimate event studies using 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

[20] However, as discussed above, our empirical results suggest that SC is in fact associated with net increases in population.

and bathrooms, along with fixed effects for age and census tract. We conduct our analysis separately between single-family homes and multi-family.

Figure 11 shows the result for single-family homes. The top row shows the impact on raw transaction prices, and the bottom shows the effect of quality-adjusted prices. As usual, the left column focuses on the ACS subsample, and the right uses the national sample. We find that prices drift higher, both unconditionally and in a quality-adjusted measure. This occurs at a delay: there is no meaningful price response until two years following implementation, after which there is a strong trend upwards. Within the ACS subsample counties, the average quality-adjusted new construction property has become 17% more expensive three years after SC rollout: an increase of \$57,300 relative to the average price of new-construction before SC. Price impact is slightly smaller in the national sample: a peak effect of 12.3% after three years. This is exactly consistent with a straightforward supply-and-demand framework. Because SC leads to sharply reduced supply of new construction, if the increase in immigration enforcement does not meaningfully change demand, we would expect to see this pattern of increases. 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 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 11 suggest that it takes about two years for shortages to become salient enough to have a large impact on market prices.

We also find endogenous shifts in the characteristics of homes that are built, which helps explain the finding that raw prices increase by more than quality-adjusted prices do in the ACS counties. Figure 12 shows the quantity response by home size for ACS counties (left) and the national sample (right). Smaller homes are less than 1,880 sq. ft., medium homes are less than 2,616, and larger homes are the remainder (these are tercile cutpoints from the data). In ACS counties, we see reductions in smaller and medium size homes but no change in aggregate building within the segment of larger homes. This means that larger homes represent an increased share of overall building after SC, which will mechanically push up raw transaction price. This is also a pattern that is only evident within the ACS subsample – one of the few instances in which we find different impacts between ACS-identifiable counties and the national sample. In general, the set of counties that can be unambiguously crosswalked to Public Use Microdata Areas tend to be larger and more urban than those which cannot. In the national sample, we see roughly equivalent declines across all size terciles. This suggests that when labor shortages lead builders to reduce output, those in larger cities prioritize construction at the upper end of the market and reduce output of smaller homes. In the sample that includes many more rural areas – where there may be less demand at the higher end of the market – the evidence suggests declines in all segments, as shown in Figure 12.

Figures A6 and A7 show additional evidence of attribute shifts within newly constructed properties. Figure A7 shows that homes in ACS counties become larger on average (a peak increase of 6.9%) and that the average number of bedrooms and bathrooms increases. Figure A7 confirms that this pattern switches in the full national sample: average size declines by 3.6% by year 3, and the average number of bedrooms shrinks as well. It is worth pointing out that these estimates likely conflate both extensive margin across size terciles as discussed in the prior paragraph, as well as intensive margin adjustment within tercile: small homes becoming even smaller in the national sample, for instance.

The bottom right figure in both A6 and A7 uses the age at sale as the dependent variable. In both cases, the evidence shows newly constructed homes selling slightly faster as we would expect with reduced supply. Age at sale is an annual measure, because we don't have the month of completion, so the downward pattern in these figures is coming either from a reduction in properties that sit on the market for at least a year before sale or from an increase in pre-sale. This is also something that we control for in measuring impact on quality-adjusted prices.

Finally, Figure A8 shows the price and square footage adjustment for condos and duplexes. Here, we find different patterns from single family homes. In the top row, we see no statistically significant impact on raw prices. In the second row, once adjusting for attributes, we see marginally statistically significant declines in quality-adjusted prices: on the order of 5% or so, with larger but very imprecise estimates at $T = 3$. The bottom row suggests that unit size is decreasing, which may help explain why raw prices remain even as supply shrinks, although these estimates are statistically imprecise. Quality-adjusted declines may suggest a reduction in quality on unobservable margins. It is also the case, as the permitting evidence suggested, that declines in multi-unit construction are much smaller than in single-family homes. Figure A9 shows the impact on new condo and duplex square footage on the same scale as Figure 10 for comparability. While there are declines, they are substantially smaller, also suggesting limited upward pressure on prices.

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? The first-stage impact, as we've shown, is a reduction in workers but no reduction in total population (indeed, an increase). 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 would certainly seem to rule out), basic economic theory would suggest that static demand paired with reduced labor supply would place upwards pressure on wages. In this section, we test for evidence of such wage increases.

There are, however, several possible stories that would predict a lack of wage 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 (Section 5.3). 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 three data from three sources. Each of these has flaws and limitations. However, taken together, we argue that our findings across all three wage measures strongly suggest a lack of any large wage increases.

5.4.1 Quarterly Census of Employment and Wages Measures

We first use data from QCEW. In addition to extensive establishment-level survey information, QCEW uses wage information reported to state unemployment insurance offices, which means that BLS staff can check and validate the information received via survey. The chief limitation of the QCEW data, however, is that we cannot disaggregate by subpopulation; we can only look at the response for all workers in a county. The top left panel of Figure 13 shows the estimated effect of SC on wages across all industries. Our results parallel a central finding from East et al. (2018): SC appears to depress wages overall. SC is associated with a decline of approximately $1,000-1,500 annually, a reduction of 2-3% from the mean. QCEW average wages are computed as the ratio of total wages to total employment, meaning that this measure also captures intensive margin shifts. The top right panel repeats this estimation for construction workers: we find declines that are slightly smaller than the all-industry benchmark.

In the bottom left panel, we estimate response for the hospitality sector, one of two sectors that employ a greater number of undocumented workers than construction.[21] Here we see evidence of wage increases although statistical significance is marginal and magnitudes are much smaller. This suggests, however, that there isn't a structural feature of the data which necessarily generate negative findings: we are able to observe wage increases in other industries that are likely significantly impacted by SC. Finally, in the lower right, we look at the impact on relative wages, which is the most meaningful measure for thinking about attracting additional workers. Since wages go down overall, the construction sector will become relatively more attractive even with wage declines, as long as those are smaller. We find that relative construction wages do increase – however, the magnitudes are small. A year after SC implementation, relative wages in the construction sector are up by about 2%. That relative increase appears to go away over the next two years. Compared to the magnitude of the declines we document in construction workforce and homebuilding output, this does not seem likely to be a highly meaningful wage response.

5.4.2 ACS Wage Measures

We next use ACS data to produce measures of wage impact by subpopulation. As discussed in Section 4.5, the ACS data structure does have a significant flaw: precise measures of income are reported, but total hours worked are reported imprecisely (via binning of weeks worked). Wages are the ratio of total pay to total hours. Therefore, given our results showing a slowdown in the construction sector, this means the econometrician is likely to observe downwards shifts in pay but fail to observe the downward shifts in hours that drive some of this pay reduction; the result will be a wage index that is lower than it should be. In other words, the ACS data is likely biased towards finding reductions. Figure 14 shows the

[21] Sources vary in ranking agriculture, hospitality/service, and construction as employers of undocumented workers. Most sources seem to place agriculture or hospitality as the top two and construction third.

estimated impact on workers by subgroup. In the top panel, we find a decline in wages for all construction workers. The peak effect is just over $1/hour, which is approximately 5% of the median construction wage. As anticipated, the decline is larger than in the QCEW data. The middle panel shows the impact on LEFB workers. Point estimates show declines, however these estimates are extremely large (peaking at $5 hour) and very statistically imprecise. The bottom panel shows statistically significant declines for US-born workers. The peak decline is $1.27 per hour, which is 6% of the median wage.

In Figure 15 we consider the relative wage response in the ACS data. The top panel suggests that overall construction wages decline more or less in parallel with county-average wages: the (log) percentage difference between construction wages and average wages is statistically zero. The middle panel again has very noisy point estimates suggesting relative declines for LEFB wages up to 25%; however the standard errors are too large to place any confidence in this finding. The bottom panel suggests that construction wages for US-born workers also evolve very similarly to the overall relative wage index: a marginally significant relative increase of 1.8% in period $T = 0$ is followed by insignificant declines of 2-3% in years 2 and 3. Due to the potential for downward bias in the ACS data, at a high level, we interpret Figures 14 and 15 as a meaningful lack of evidence against any substantive *increase* in wages in either absolute or relative terms.

5.4.3 RSMeans Wage Measures

Our third 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 means that we are not able to rigorously test for parallel pretrends with RSMeans; however, the evidence using both QCEW and ACS suggests that this assumption holds quite well. Figure 16 shows results. The top panel suggests that construction wages decline steadily upon SC rollout, with a peak decline of 5%. The bottom panel parallels the relative wage estimation of the prior sections. Because RSMeans does not survey costs outside of the construction sector, we use the QCEW series for all industries to normalize. We find some evidence of relative declines for construction wages; again these are small, with a peak effect of -48bps.

Across all three wage measures, we find no evidence to support meaningful or sustained increases in construction wages in either level or relative terms. In turn, 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 leads builders not to increase wages to attract additional workers. Although it is hard to believe that builders are not making any economic profits – and indeed, our results showing 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 negative shocks to construction labor supply are highly persistent, and have a large effect on the construction of residential housing. 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 reduced population at the county level. SC also leads 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. While the average home sold declines in price, our results suggest that this comes from endogenous adjustment on home characteristics. The quality-adjusted price of newly constructed homes increases following SC.

We also find a striking lack of evidence for builders increasing wages in order to attract replacement workers. 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 wage adjustment to attract more labor seems surprising. We show that SC does appear to induce wage increases in another industry which draws heavily upon undocumented workers, suggesting that industry-specific features may drive the lack of wage increases for construction workers.

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. This paper also suggests that immigration policy, along with other interventions that directly affect domestic labor supply, may be important levers for policymakers interested in overall home affordability.

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.

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* .

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

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)

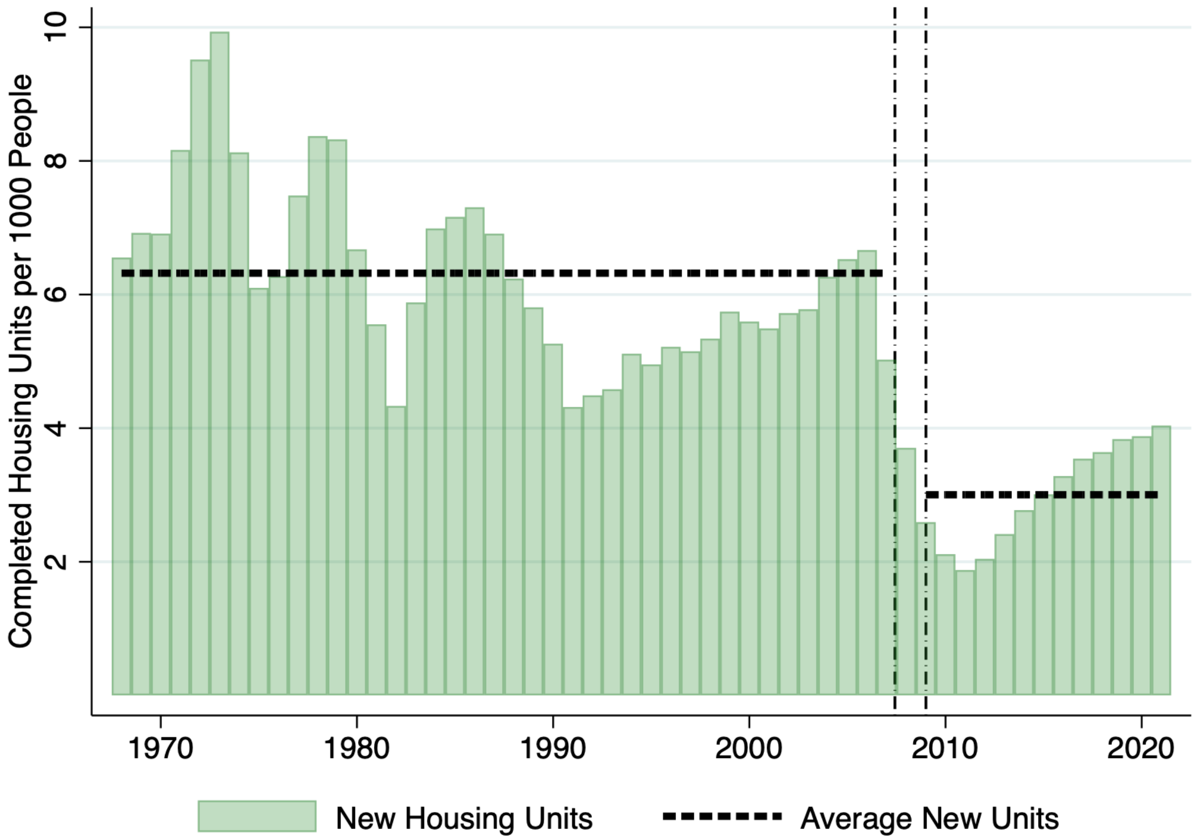

Note: This figure plots the time trends of constructed new housing units. The green bars are the annual new housing units per 1000 population (left axis) in the US from Census Bureau and HUD. The two dashed lines indicates the average levels of new housing units per 1000 population pre-GFC (1968–2007)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

Figure 2: Staggered Rollout of Secure Communities

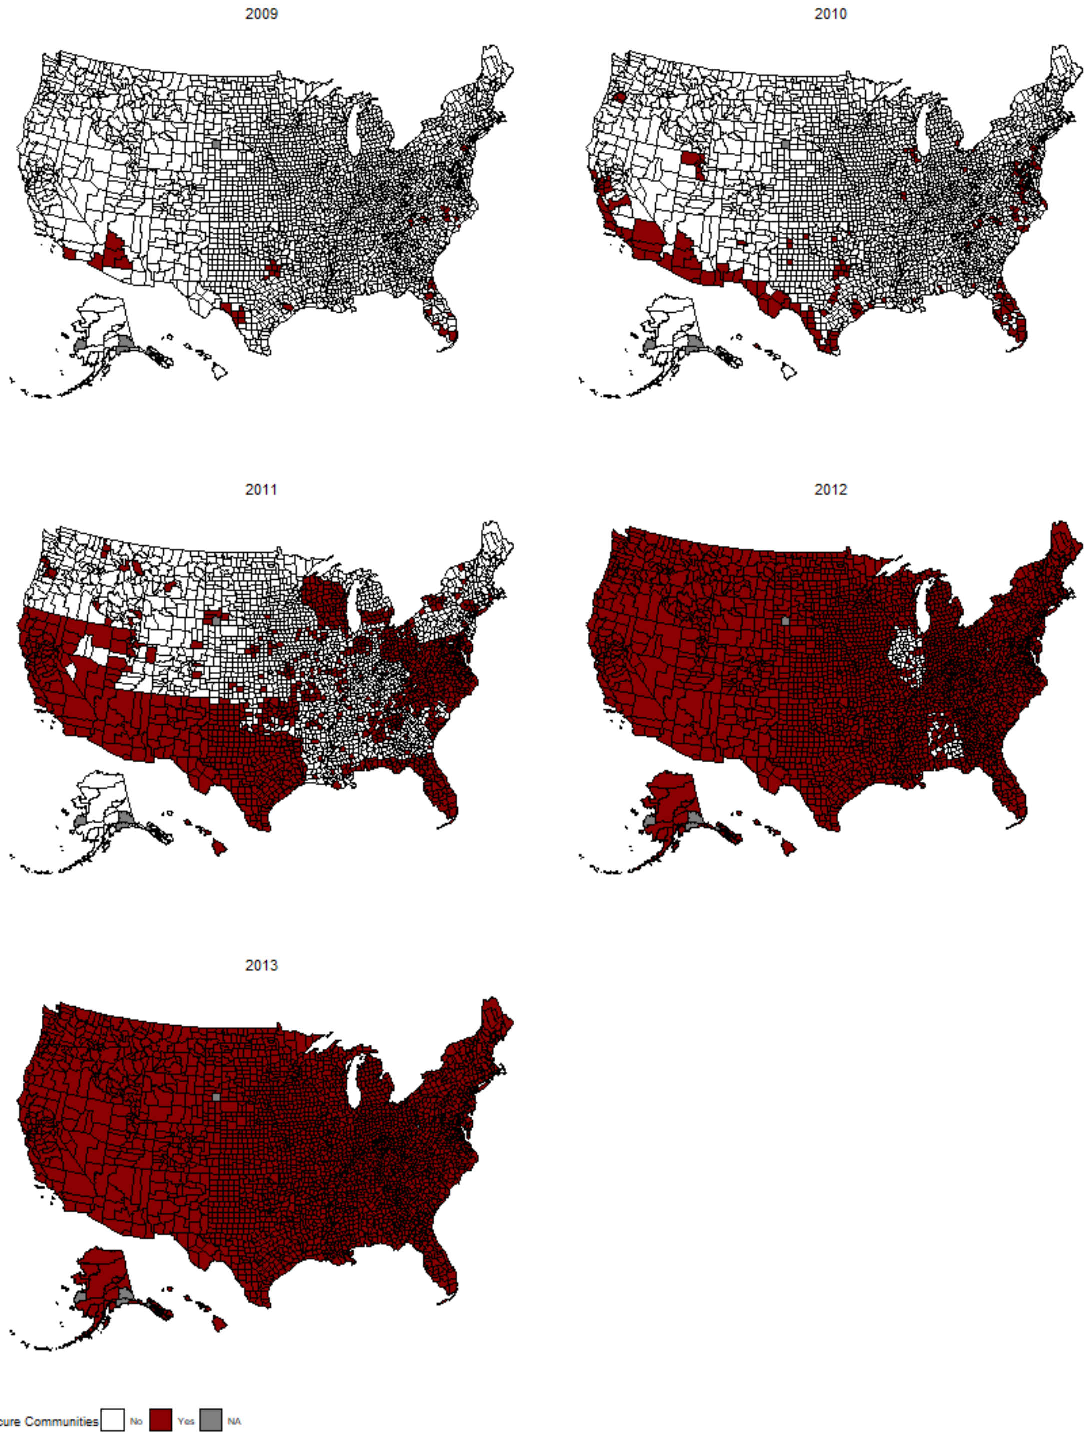

Note: Each panel of this figure shows the counties that implement SC within each year. This map reflects the treatment indicator used in our regressions, which assigns binary treatment status to any county operationalizing SC for at least half the year. Counties launching SC in, for instance, December of year t would therefore be coded as untreated in year t and treated in year $t+1$. Appendix Figure A1 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: Population Impact of Secure Communities

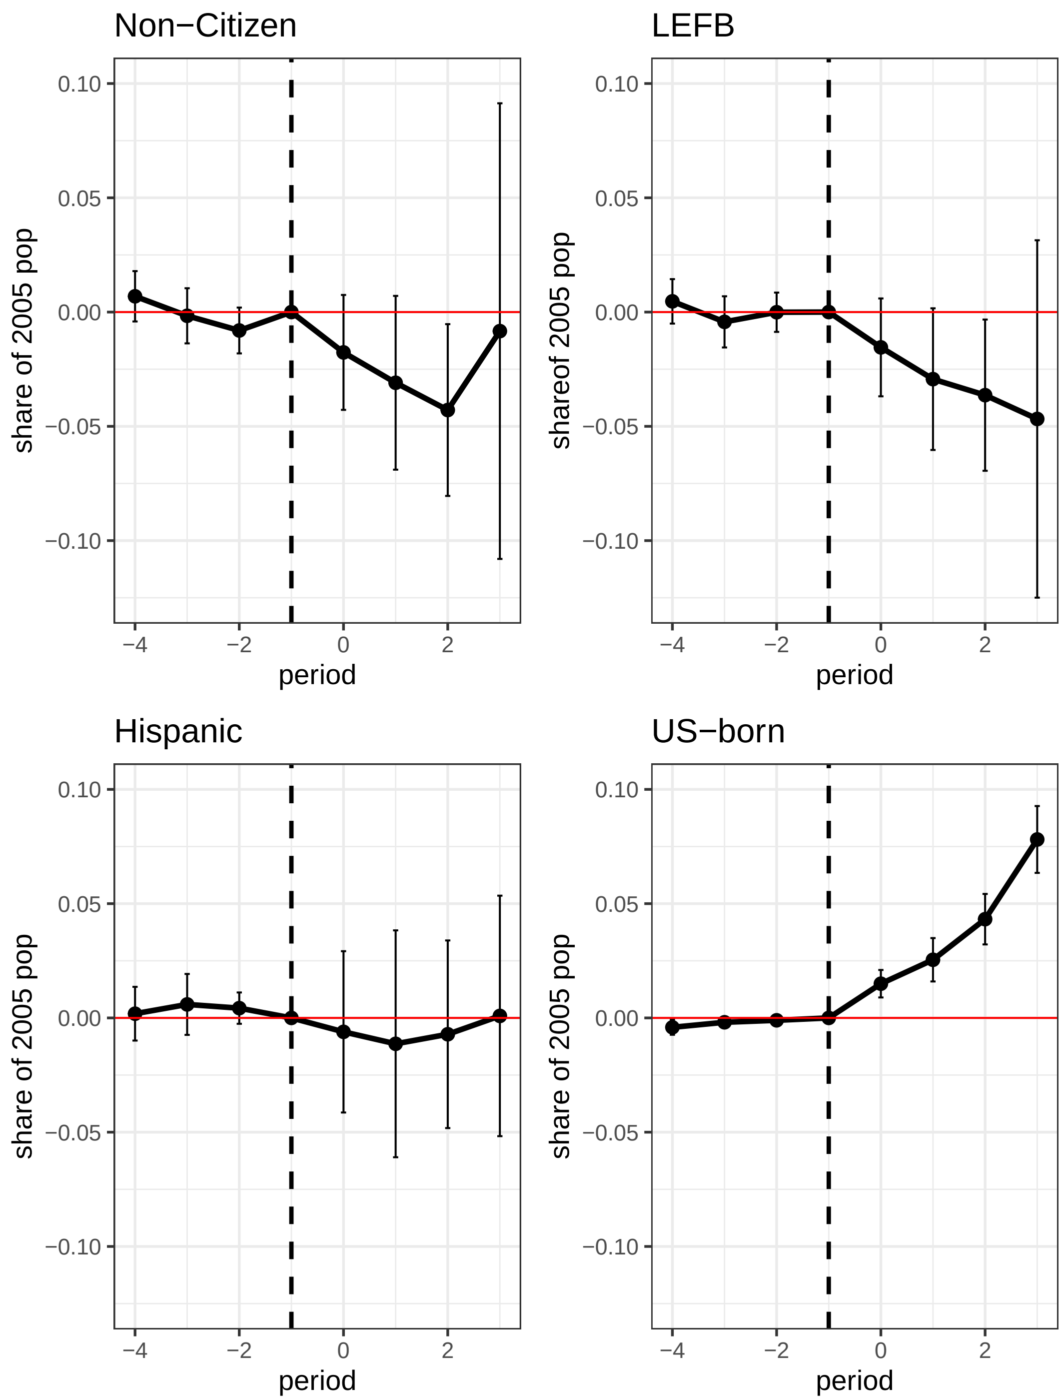

Note: This figure plots the impact of SC on population changes, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on noncitizen, low-education and foreign-born (LEFB), Hispanic, and US-born populations. We use the 2005 population to normalize the dependent variable to population share. 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 4: Construction Workforce Impact of Secure Communities

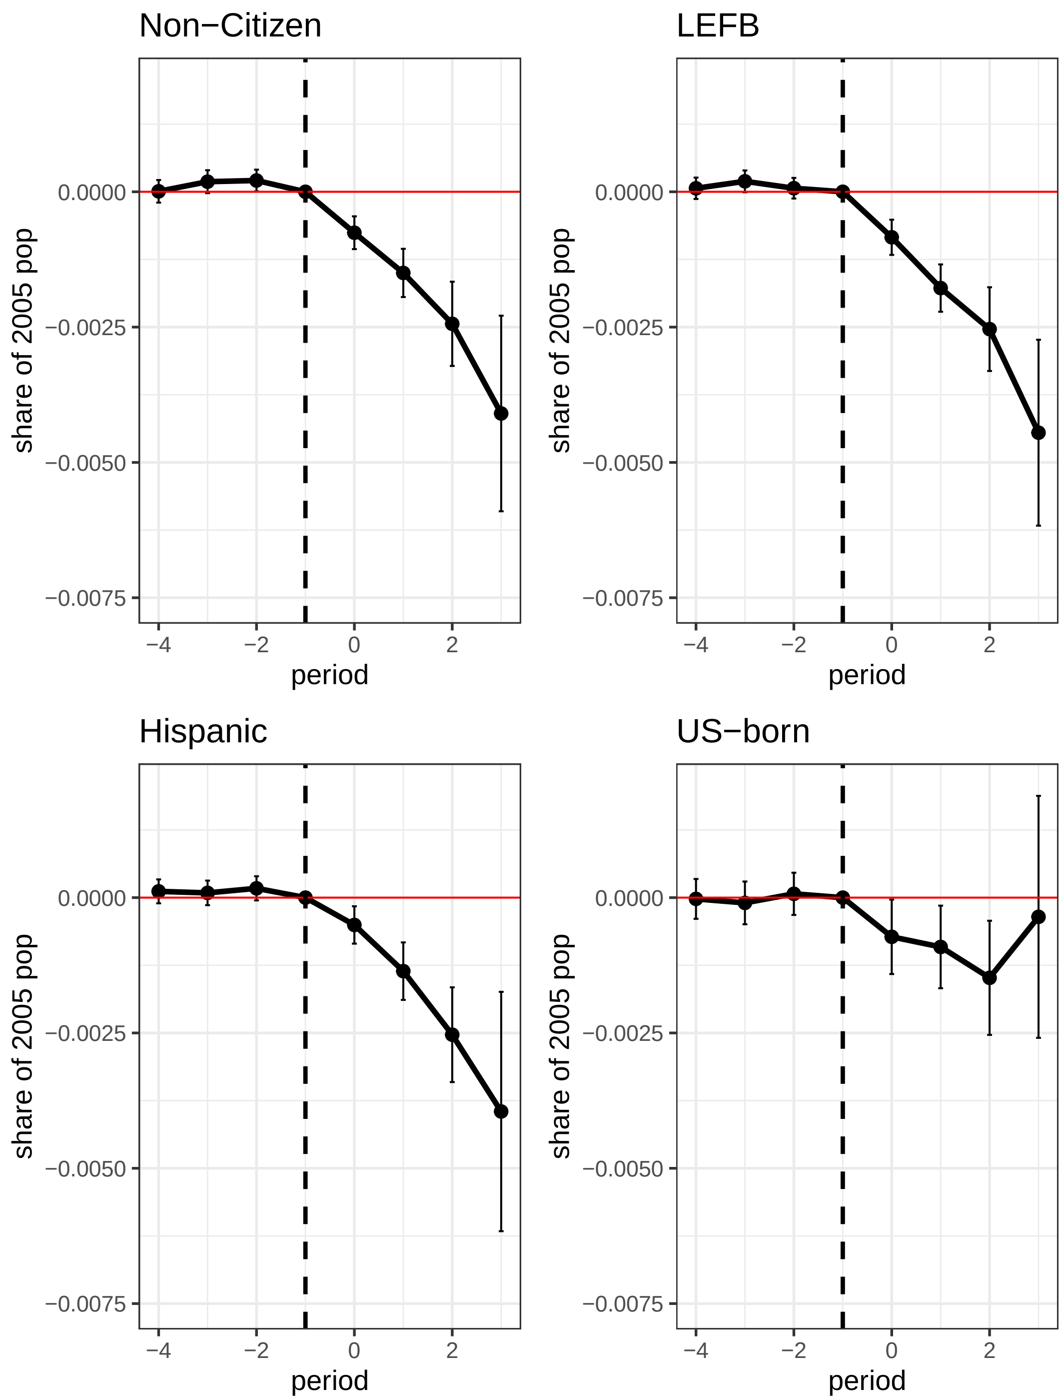

Note: This figure plots the impact of SC on construction employment, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on noncitizen, low-education and foreign-born (LEFB), Hispanic, and US-born workers in the construction sector. We normalize the number of 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

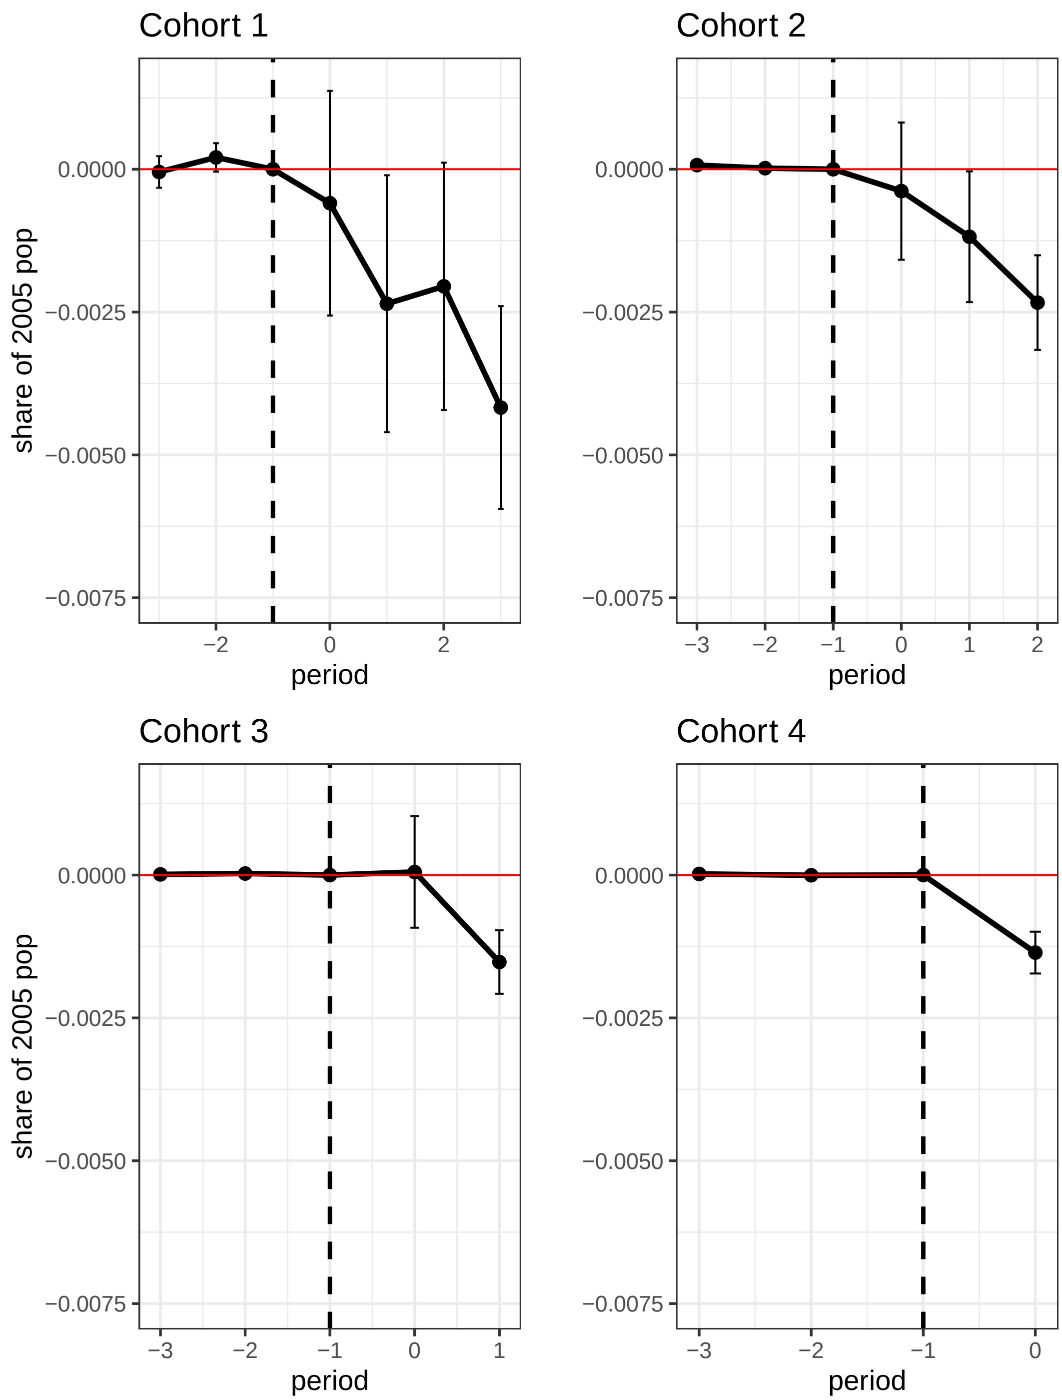

Note: This figure plots the impact of SC on LEFB construction employment by treatment-cohort, with the approach of Gardner (2022) and specification (2). We use a balanced (and constant) control group of the last counties to be treated. The four panels compare counties treated in 2009, 2010, 2011, and 2012, respectively, with only counties untreated through 2012. 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

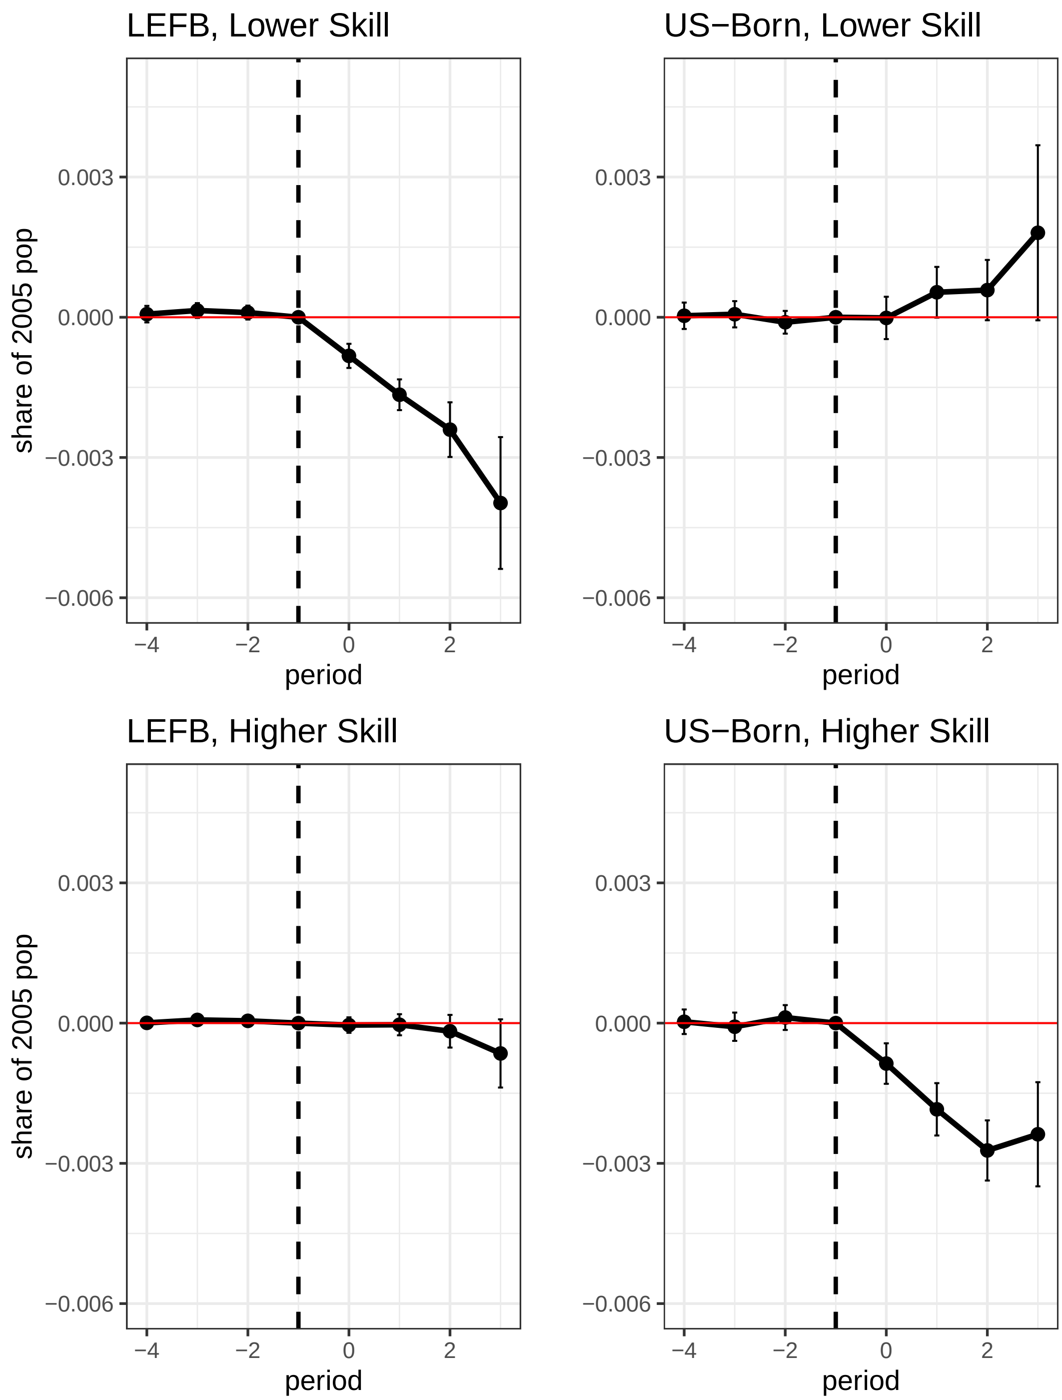

Note: This figure plots the impact of SC on LEFB and US-born construction employment by skill classification, with the approach of Gardner (2022) and specification (2). We define the following occupations as lower skill, construction laborers, helpers in construction trades, painters and maintenance 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: Intensive Margin Impact - Construction Workers

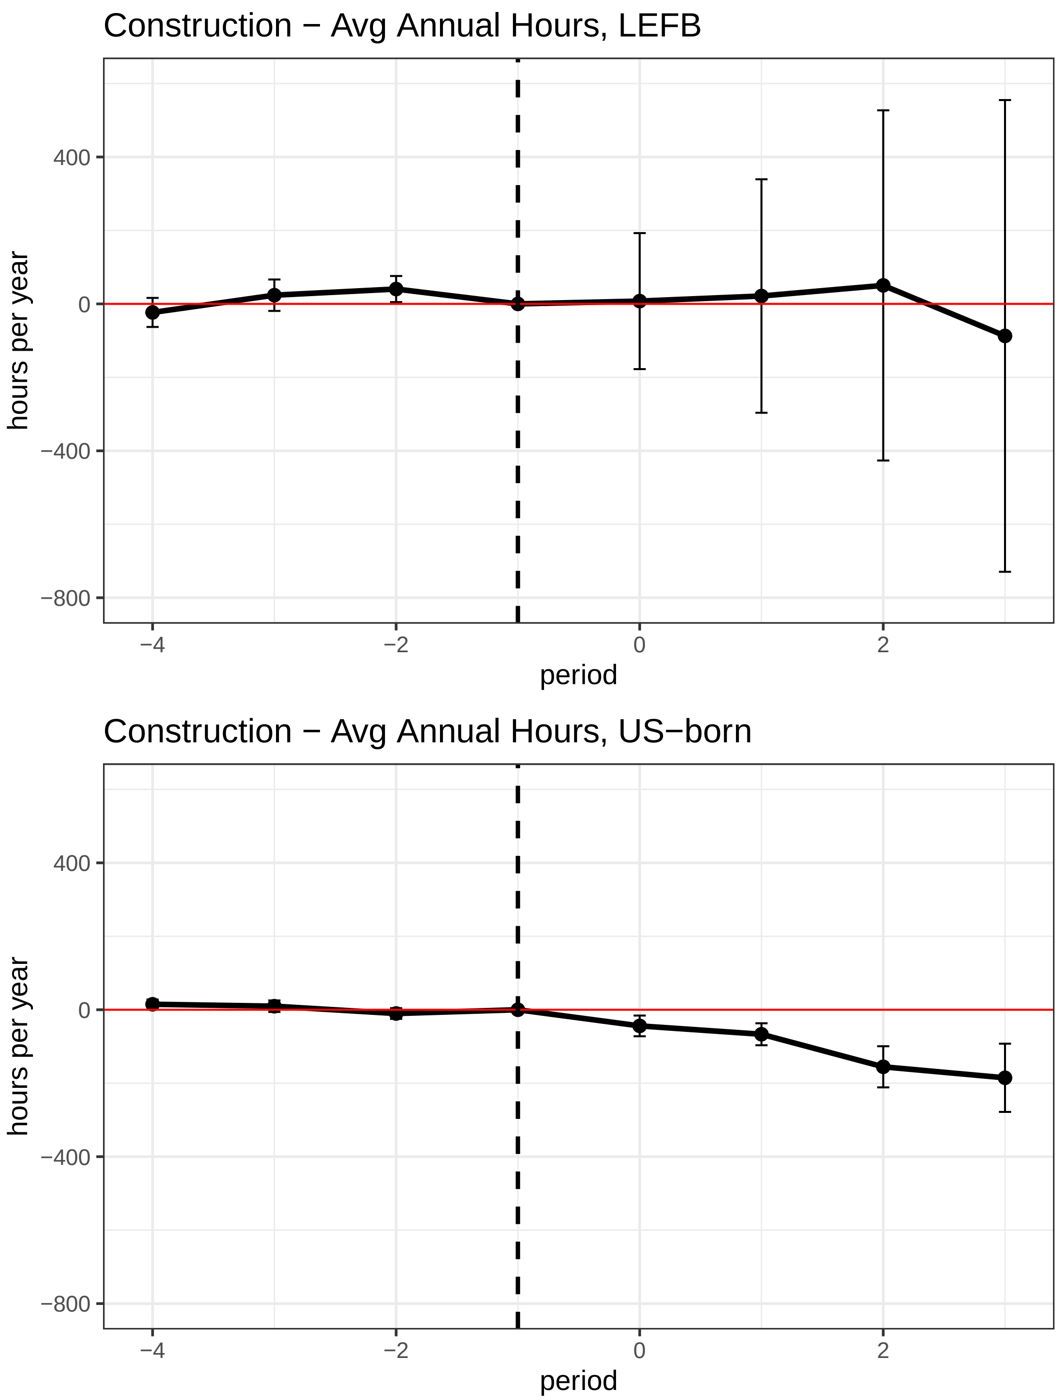

Note: This figure plots the impact of SC on the intensive margin of construction employment, with the approach of Gardner (2022) and specification (2). Using ACS data, we compute the average hours worked per worker in each subpopulation. The top panels shows average working hours per LEFB worker, and 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 8: Total Permits per 1,000 residents

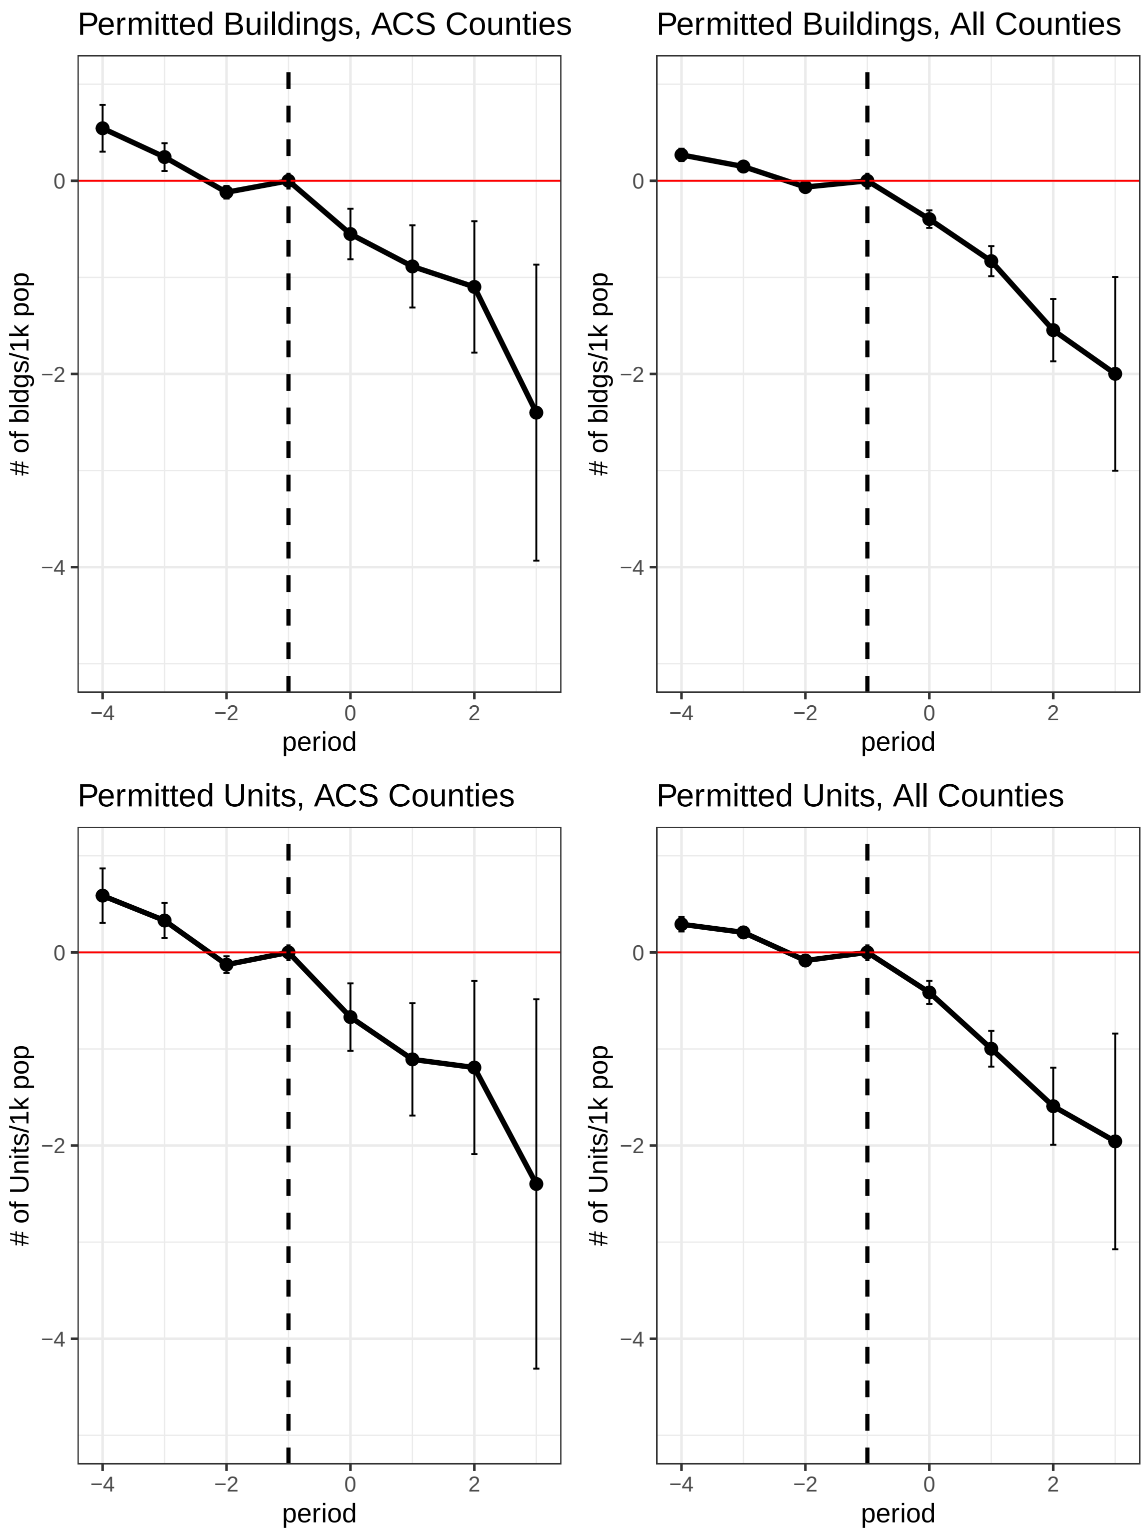

Note: This figure plots the impact of SC on residential construction activity measured by residential permits (intended construction), with the approach of Gardner (2022) and specification (2). The left columns examines the impact within ACS subsample counties, and right within the national sample. The top panels examine the impact on permitted buildings, and the bottom examines total permitted units. We use total permits per 1,000 residents as the outcome variables. 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 9: Permits by Building Class, ACS Counties

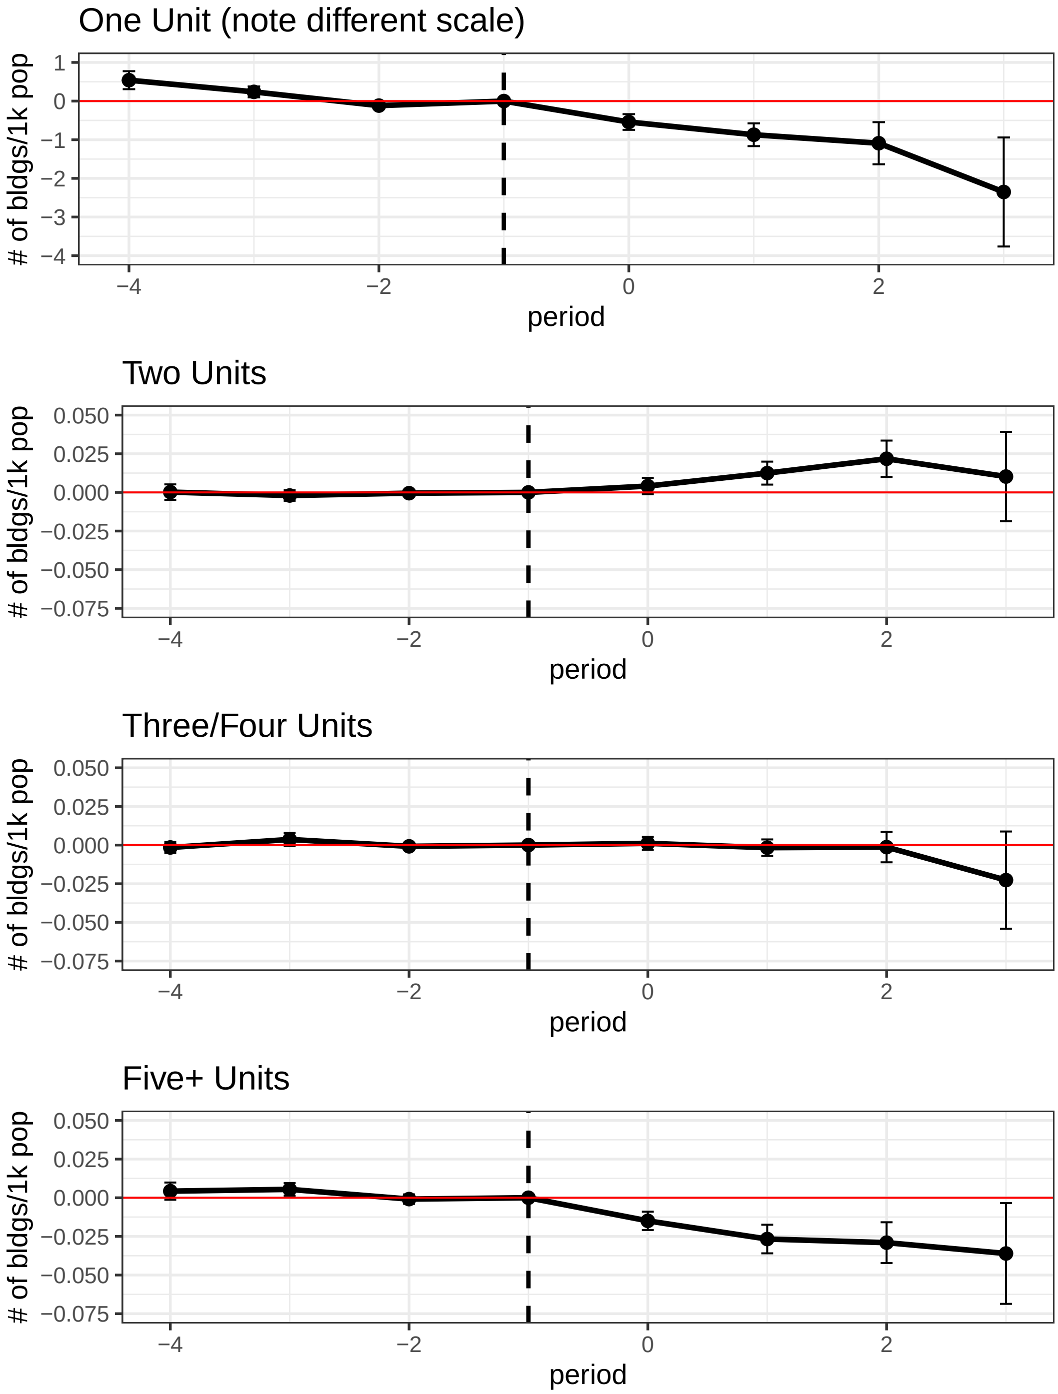

Note: This figure plots the impact of SC on residential permits (intended construction) by building size, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on permitting buildings of one-unit, two-unit, three/four-unit, and five/more-unit. Estimations in this figure are based on the ACS subsample. We use total permitting buildings per 1,000 residents as the outcome variables. 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: New Construction Entering Market

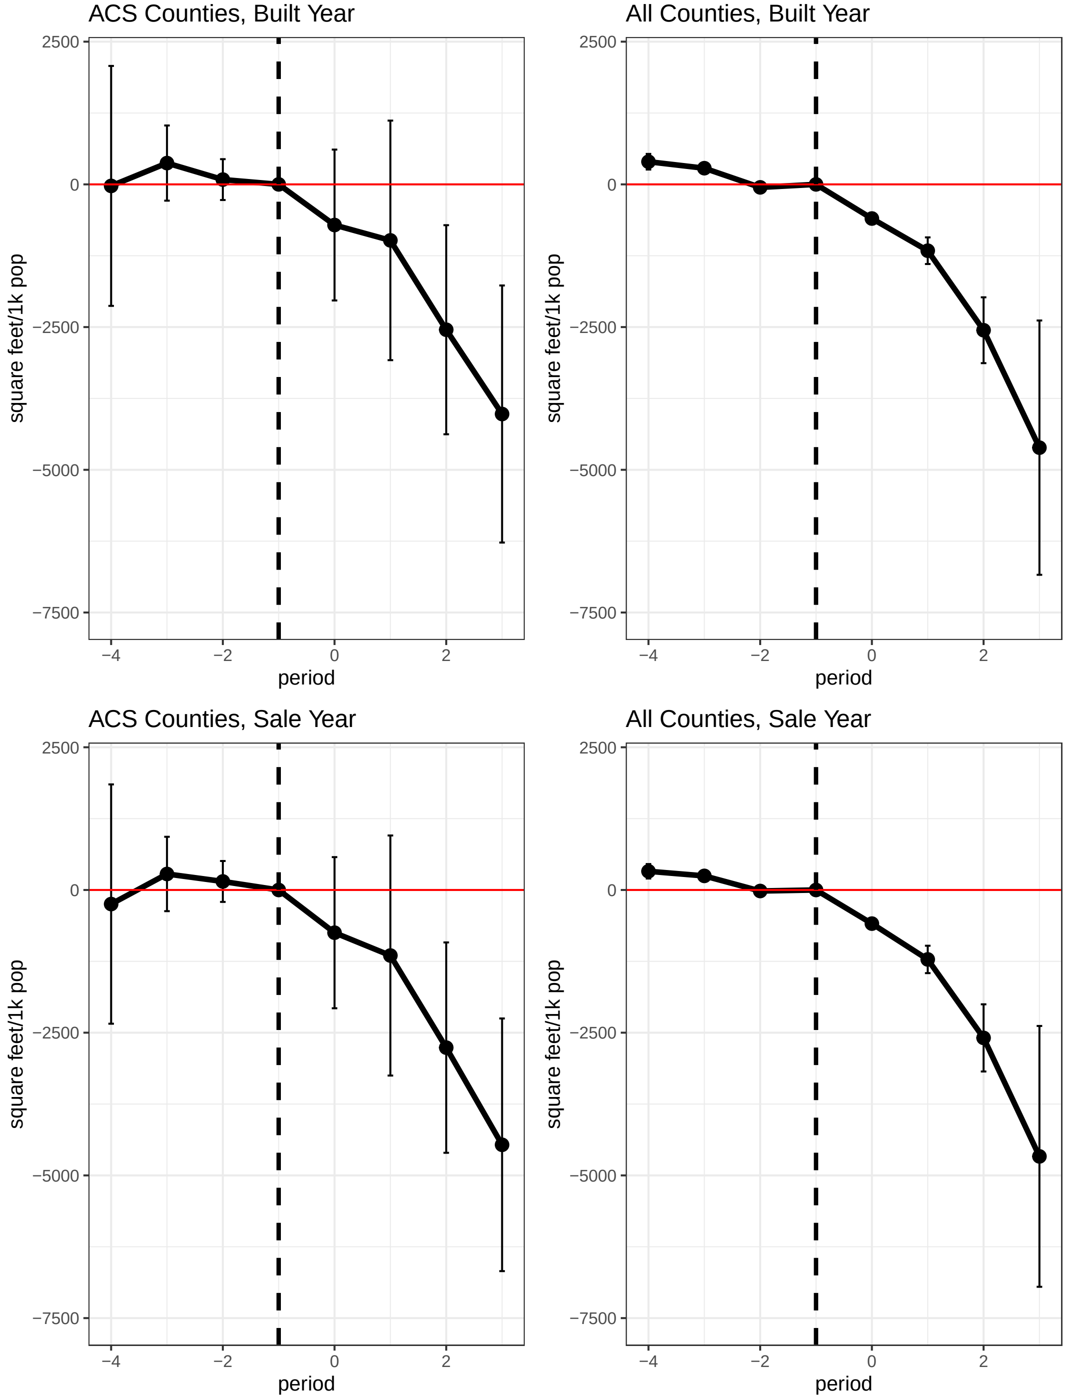

Note: This figure plots the impact of SC on residential construction activity measured by observed new construction (completed new construction), with the approach of Gardner (2022) and specification (2). Completed new construction is aggregated to county-year level using administrative tax-roll data 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 one period is a year.

Figure 11: Price Response, Single-Family Homes

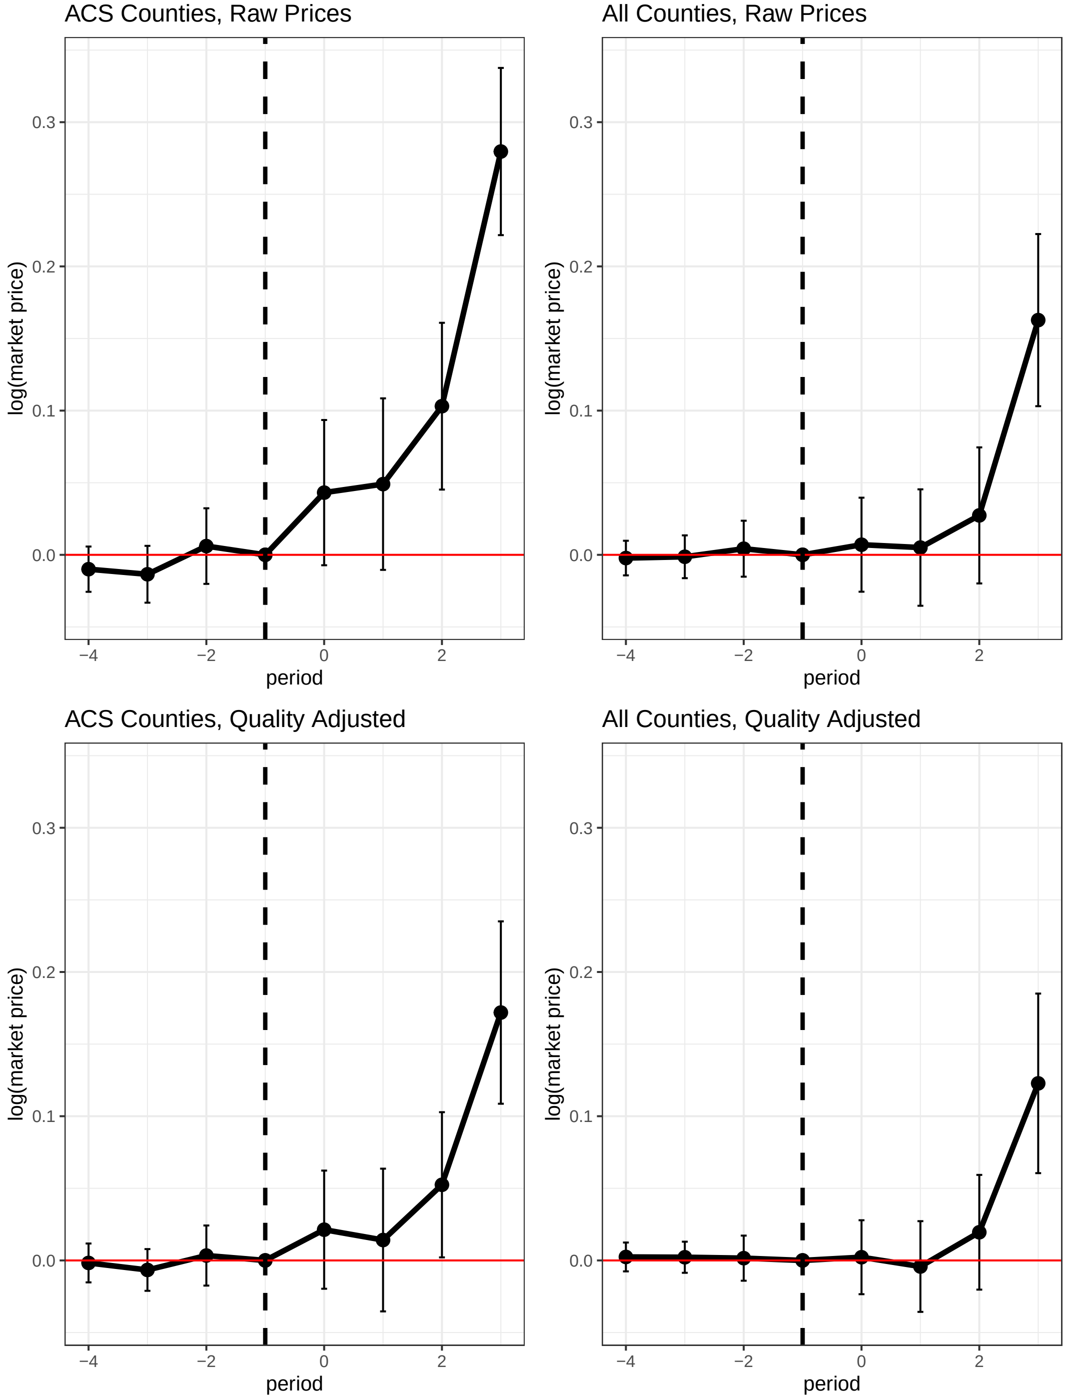

Note: This figure plots the impact of SC on single-family home prices, with the approach of Gardner (2022) and specification (2). The top row shows the effect on single-family homes without attribute controls (but with a tract fixed-effect). The bottom row includes hedonic controls to show the impact 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 12: Portfolio Shifts - Results by Size Tercile

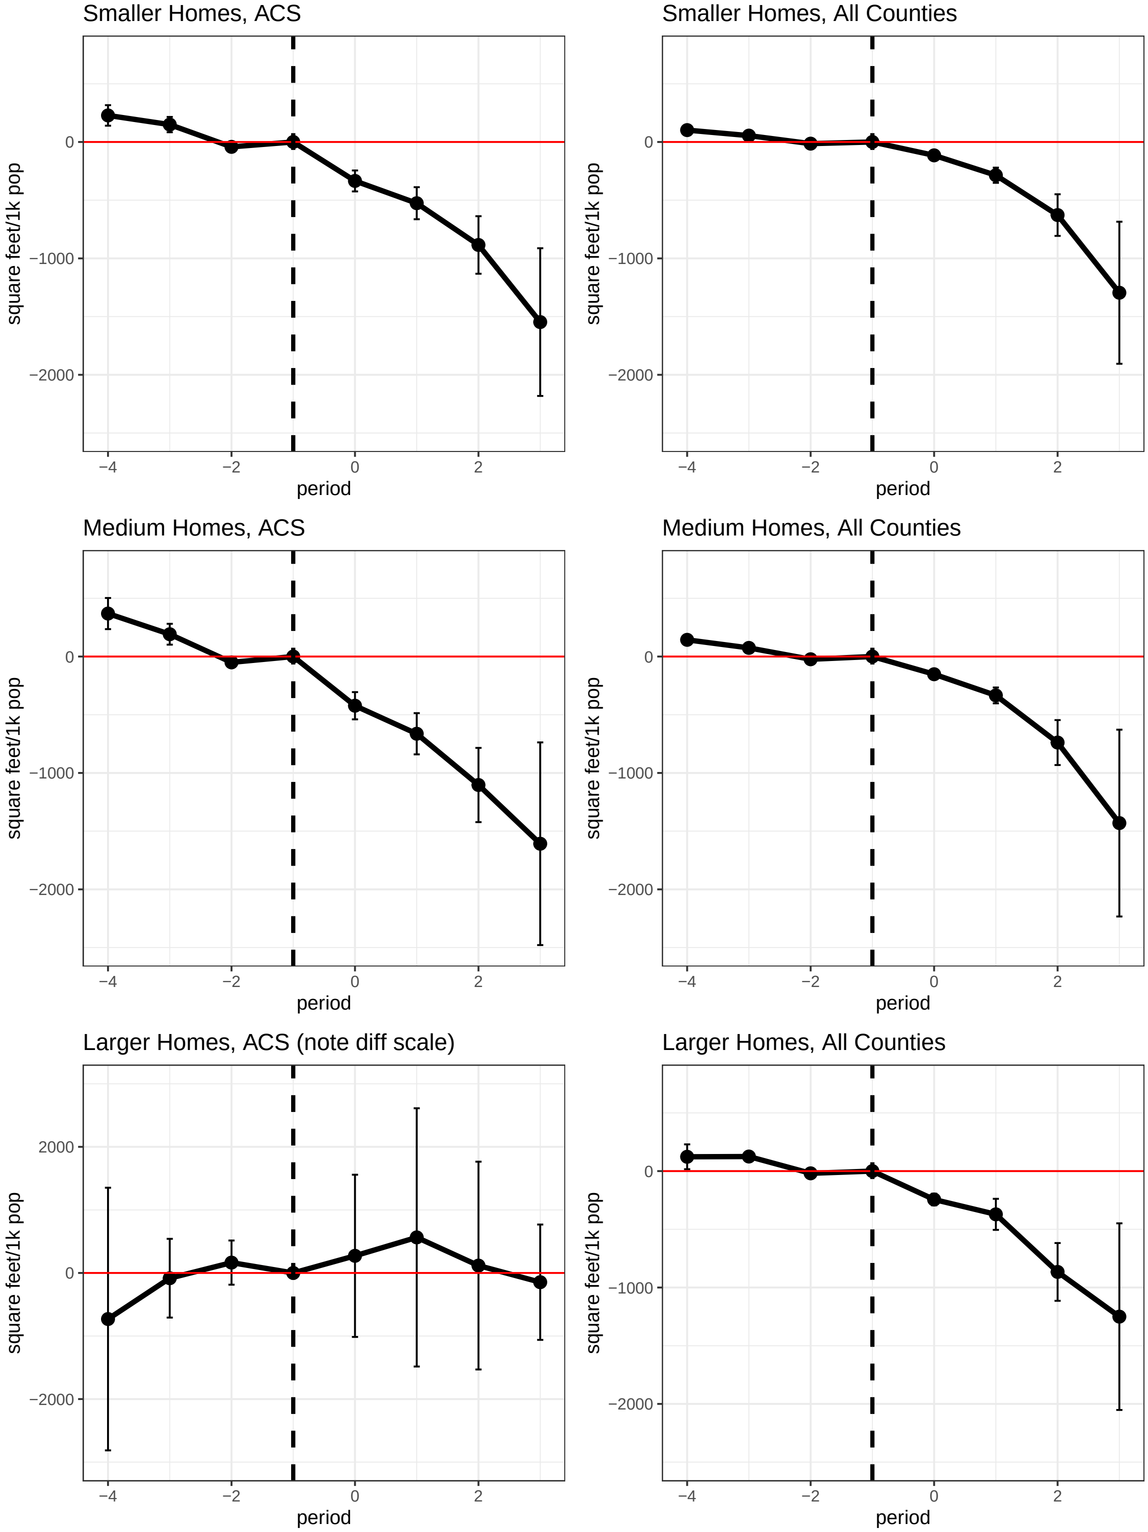

Note: This figure plots the impact of SC on residential construction activity by size tercile, with the approach of Gardner (2022) and specification (2). Results on the left are based on the national sample, and results on the right are based on the subset of counties separately identifiable in the ACS 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 one period is a year.

Figure 13: Effect on Wages, QCEW (Annual)

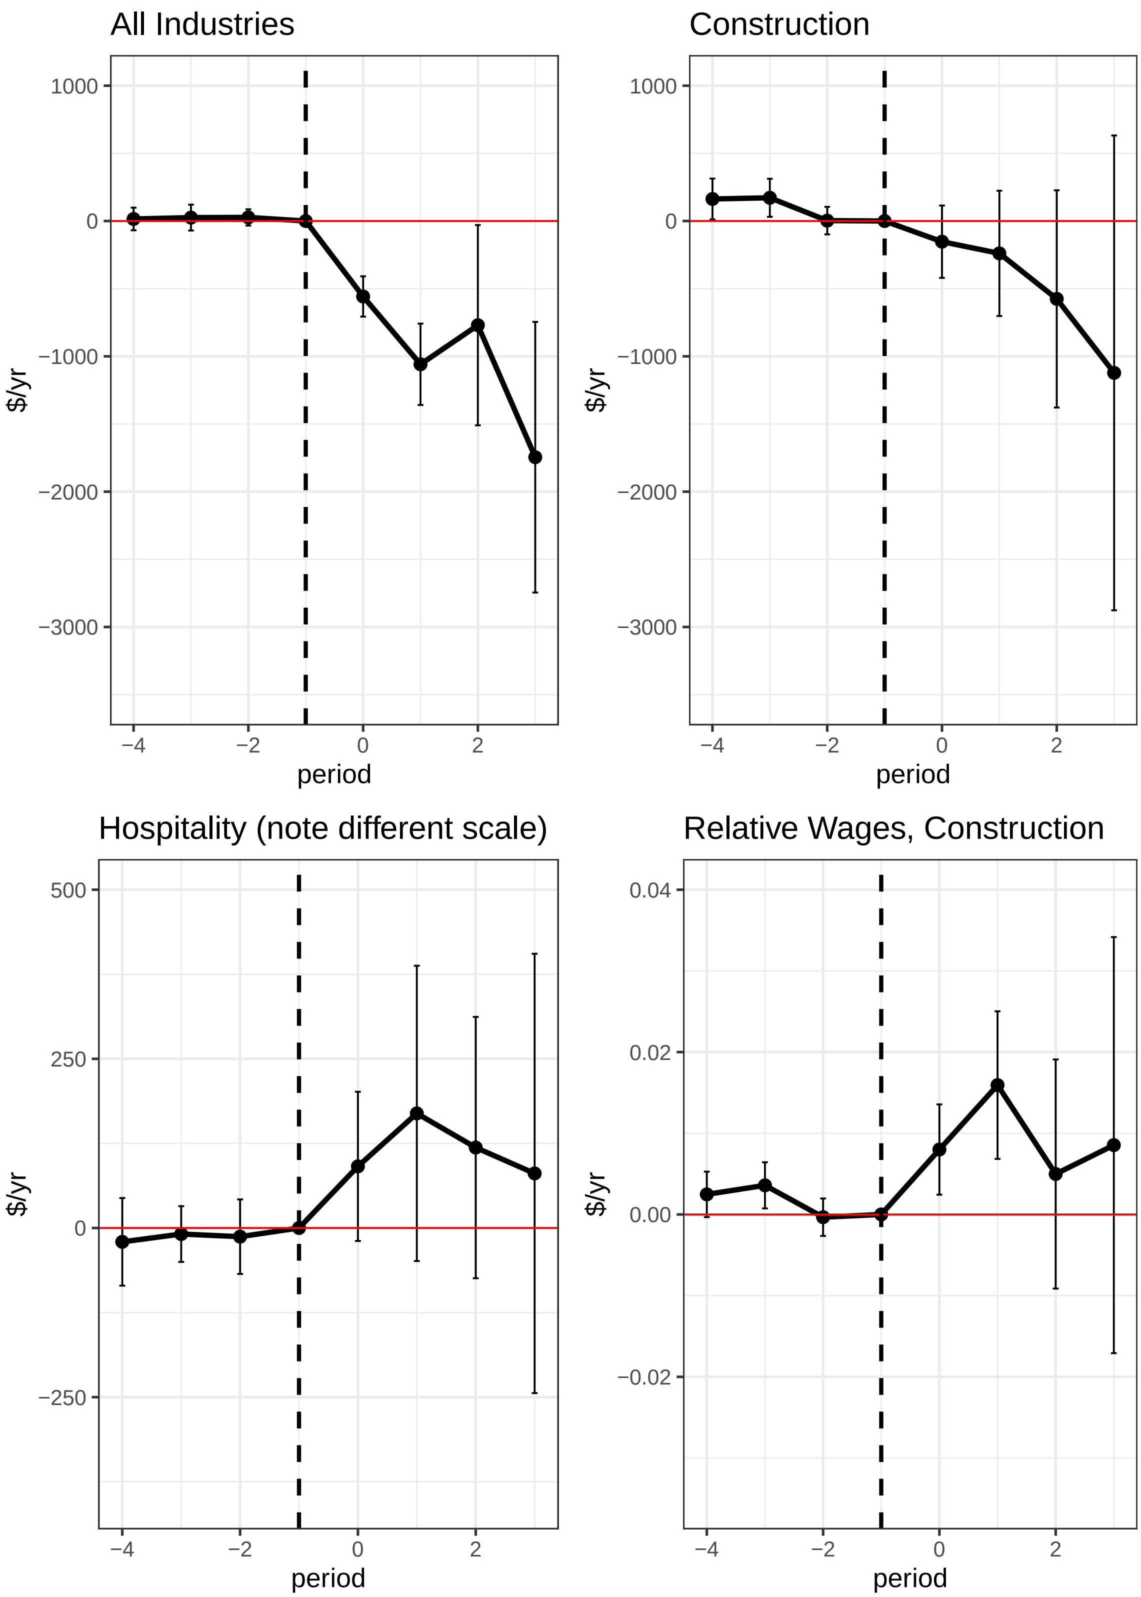

Note: This figure uses QCEW data on annual wages to plots the wage impact of SC with the approach of Gardner (2022) and specification (2). The two top figures and the bottom left show the level impact for all industries, construction, and hospitality respectively. The bottom right figure plots relative wage for the construction industry: the ratio of construction wages to the average across all industries. 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 14: Effect on Construction Wages, ACS

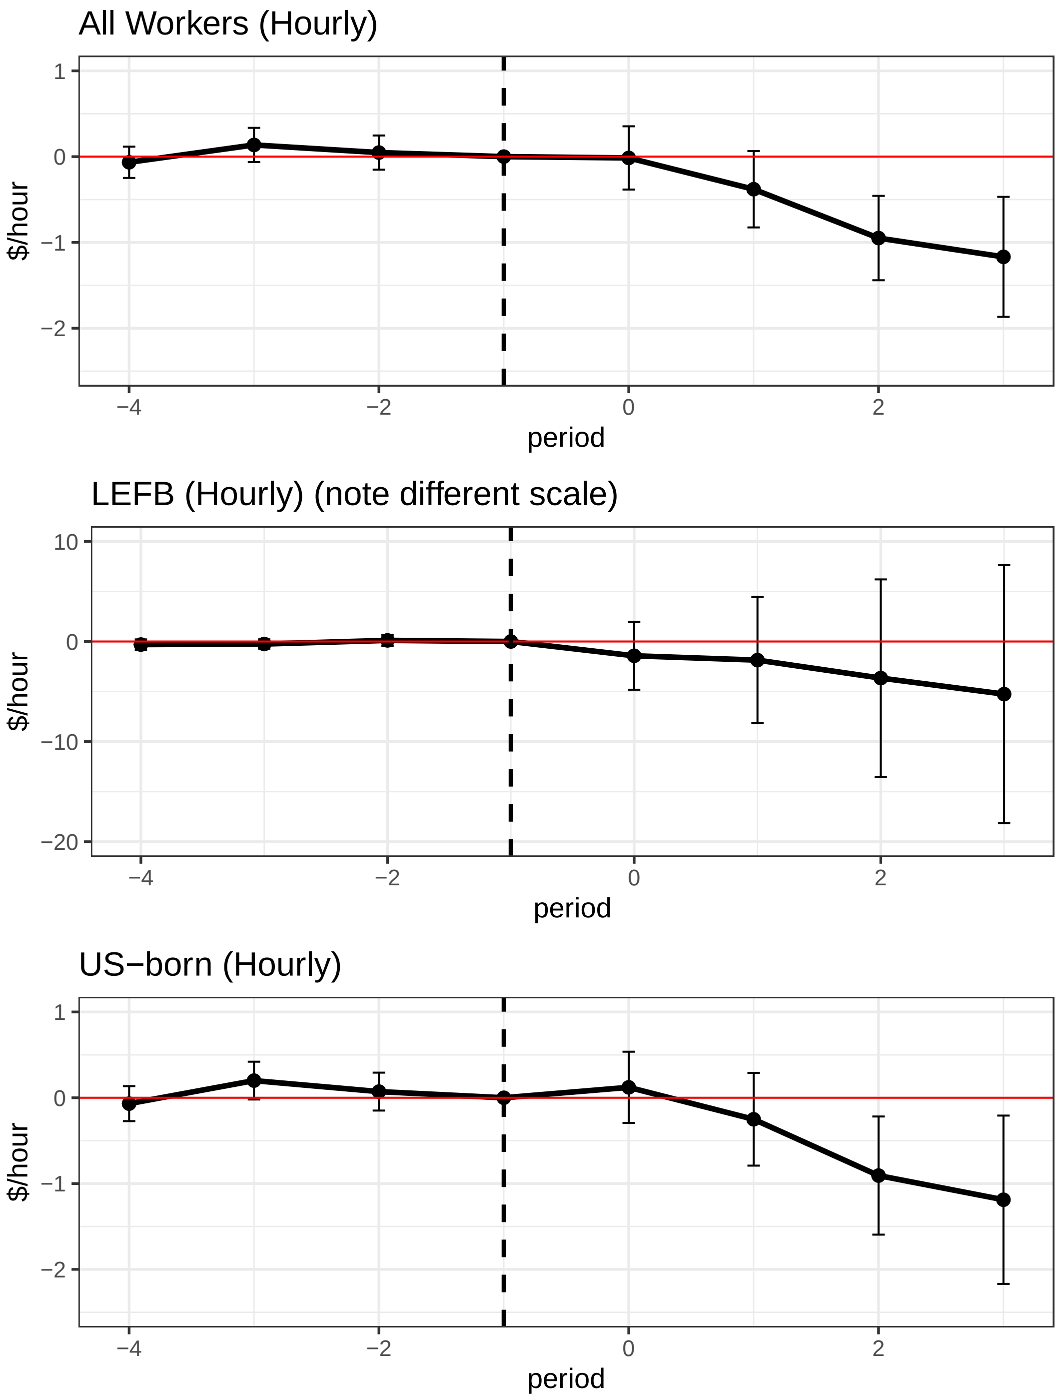

Note: This figure uses ACS data to plot the impact of SC on average county-level wages for construction workers with the approach of Gardner (2022) and specification (2). All panels use a measure of hourly average wage constructed from ACS microdata. The top panel shows impact on all workers, the middle panel shows impact on LEFB workers, and the bottom panel 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 15: Relative Construction Wages, ACS

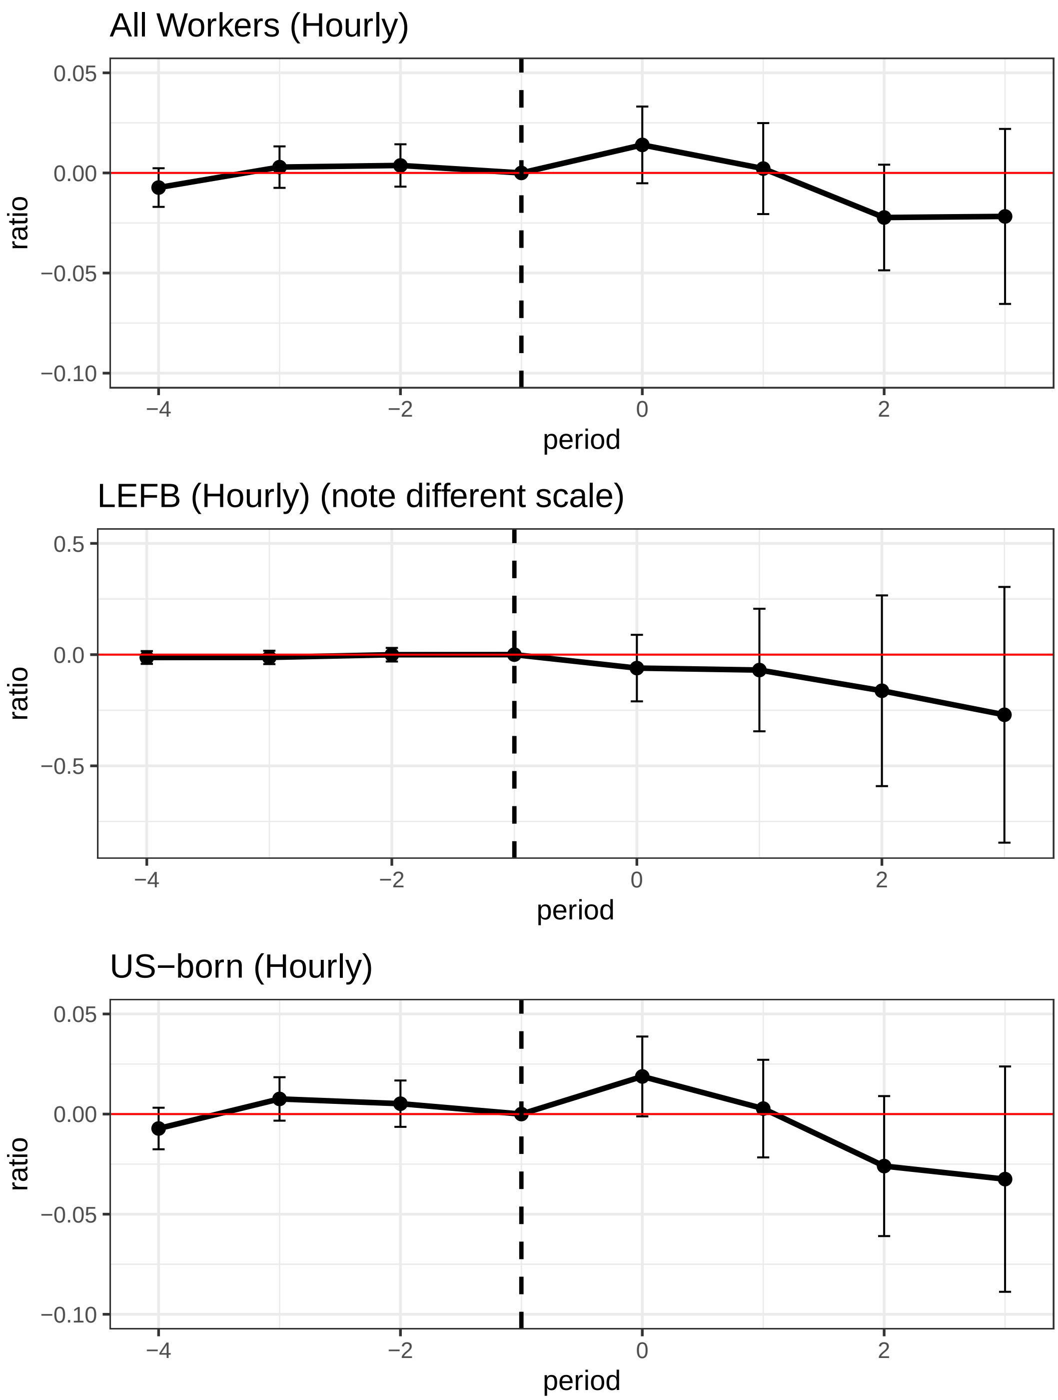

Note: This figure uses ACS data to plot the impact of SC on *relative* county-level wages for construction workers with the approach of Gardner (2022) and specification (2). Relative wages for a given group are defined as the within-group ratio of average construction wages to the average across all other industries excluding construction. The top panel shows impact on all workers, the middle panel shows impact on LEFB workers, and the bottom panel 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 16: Effect on Wages, RSMeans

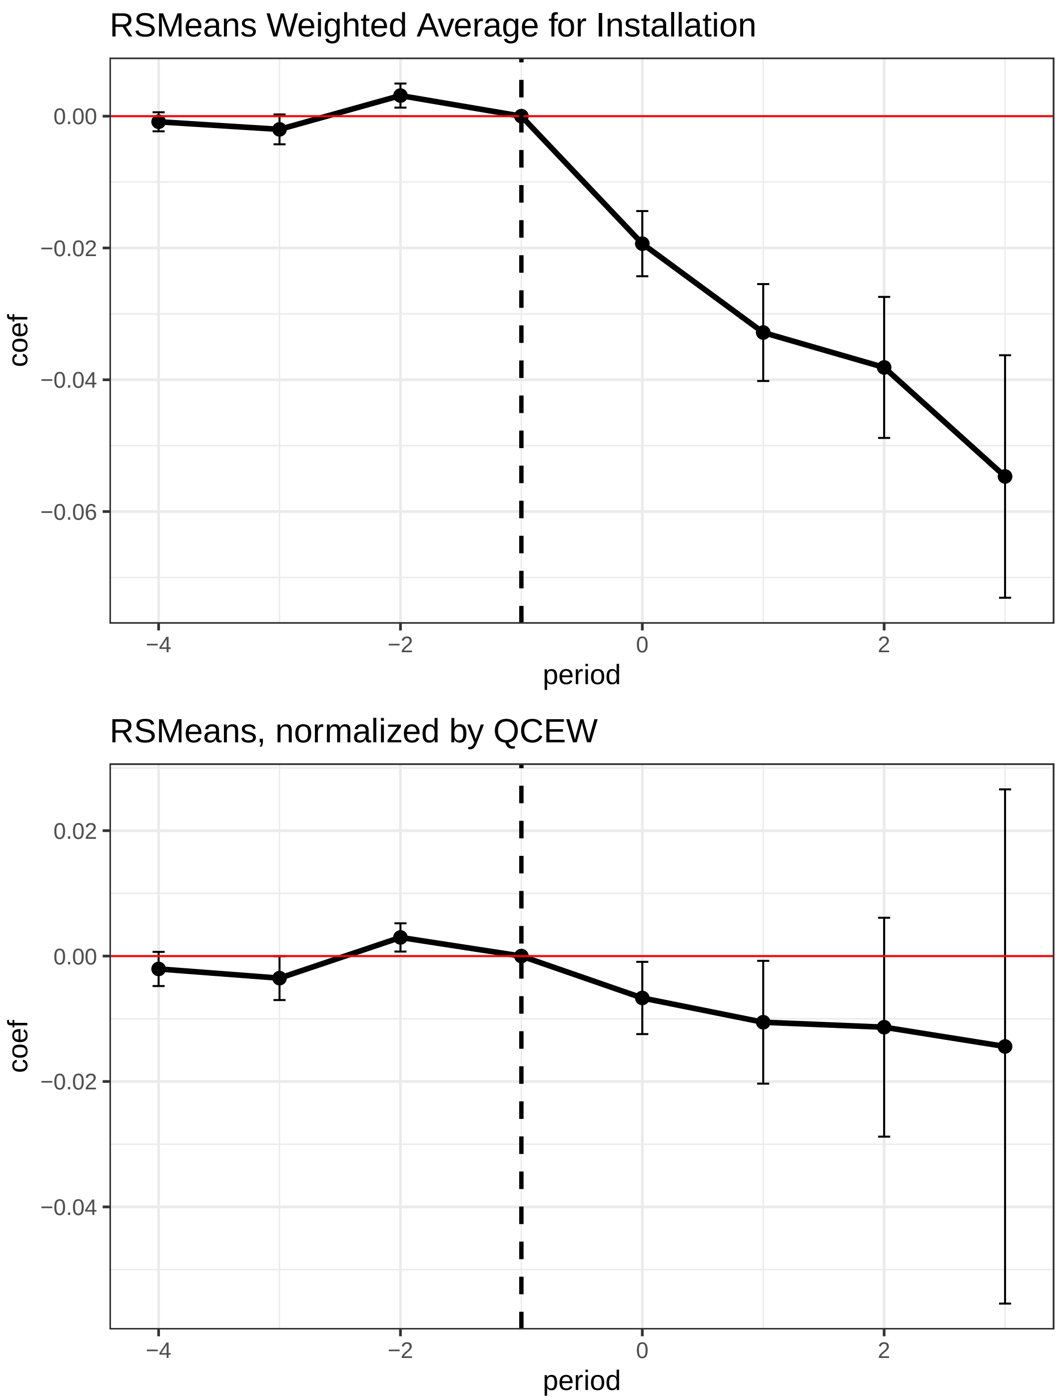

Note: This figure plots the impact of SC on average county-level construction wages according to industry data compiled by RSMeans. We use the approach of Gardner (2022) and specification (2). The top panel shows the level impact on RSMeans' regional wage index. The bottom panel shows the effect on relative wages, where the RSMeans index is normalized by an index of wages for all industries created from QCEW data. 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.
Population as Share of 2005 Total Population (%)								
Non-citizen	6.48	5.26	4.69	5,621				
LEFB	7.48	6.07	5.35	5,621				
Hispanic	14.05	17.79	7.30	5,621				
US-born	97.19	12.38	97.95	5,621				
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	4531	29674	1847	2,648	1416	10387	8	25,040
Square Feet, Sale Year	4918	29673	2254	2,648	1540	10406	15	25,040
Units, Built Year	1.77	2.73	0.83	2,648	0.63	1.97	0.00	25,040
Units, Sale Year	2.01	2.86	1.05	2,648	0.70	2.03	0.00	25,040
New Construction Transactions – Micro Data								
Market Price ($k)	320	222	264	1,666,794	305	215	250	2,500,913
Square Feet	2232	983	2063	1,666,794	2228	1411	2045	2,500,913
# Bedrooms	3.36	1.17	3.00	1,666,794	3.35	1.63	3.00	2,500,913
# Bathrooms	2.78	1.30	3.00	1,666,794	2.75	1.24	3.00	2,500,913
Age at Sale	0.63	1.38	0.00	1,666,794	0.62	1.33	0.00	2,500,913
Annual Wages from QCEW ($)								
All Industries	44370	13547	41691	5,621	35753	10410	34105	50,266
Construction	50872	13331	49637	5,621	37693	17271	38492	50,105
Hospitality	18120	5741	16726	5,621	13726	5197	13051	50,146
Hourly Wages from ACS ($)								
Construction	19.67	4.68	19.22	5,621				
Construction, LEFB	16.41	7.40	15.07	4,516				
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 new construction, in addition to various wage measures.

Table 2: Is SC Rollout Predictable?

	Binary for Rollout						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Total Pop (M)	0.0450						
	(0.0313)						
Hispanic Share		0.7220***	0.6829***	0.7099***	0.7277***	0.7215***	0.7329***
		(0.1171)	(0.1180)	(0.1181)	(0.1172)	(0.1225)	(0.1304)
3Yr Pop Growth, US			0.8312***				
			(0.2814)				
3Yr Pop Growth, Hisp				−0.0295			
				(0.0296)			
3Yr Pop Growth, LEFB					0.0356		
					(0.0340)		
3Yr NC Growth						−0.0001	
						(0.0002)	
01-07 Price Runup							−0.0188
							(0.0580)
Observations	962	962	962	962	962	887	962
R^2	0.3419	0.3774	0.3824	0.3778	0.3780	0.3629	0.3775

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 a stack fixed effect, and standard errors are clustered at the county-level.

A Appendix

Figure A1: Staggered Rollout of Secure Communities (by exact date)

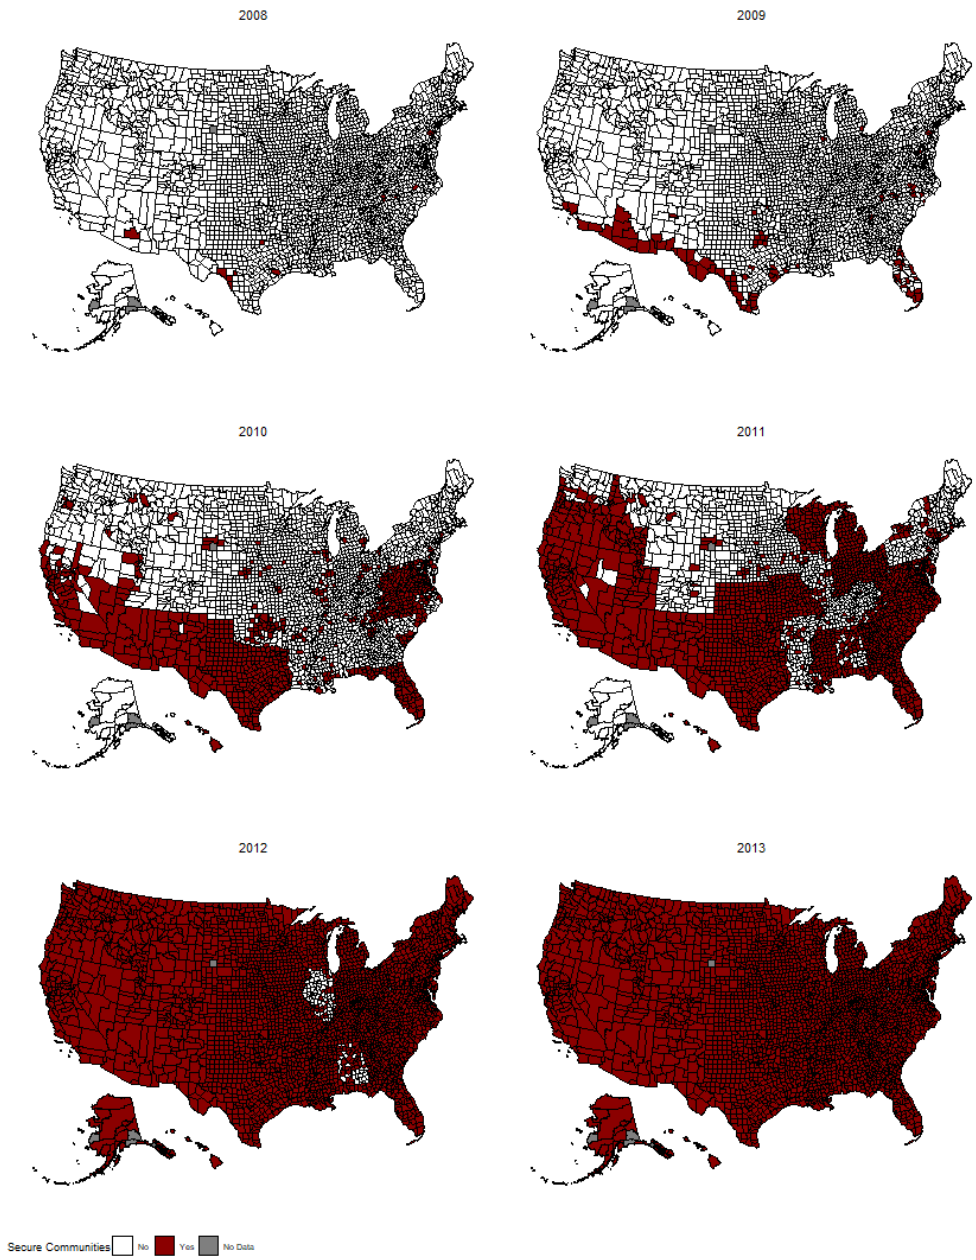

Note: Each panel of this figure shows the counties that implement SC within each year. This map reflects treatment based on exact date of implementation: a county is coded as treated in year t if the launch date falls at any point within year t. All regressions in this paper assign annual treatment status only to counties which have been treated for at least half a year; a corresponding map of this empirical treatment indicator is shown in Figure 2.

Figure A2: Permitted Units by Building Class, ACS Counties

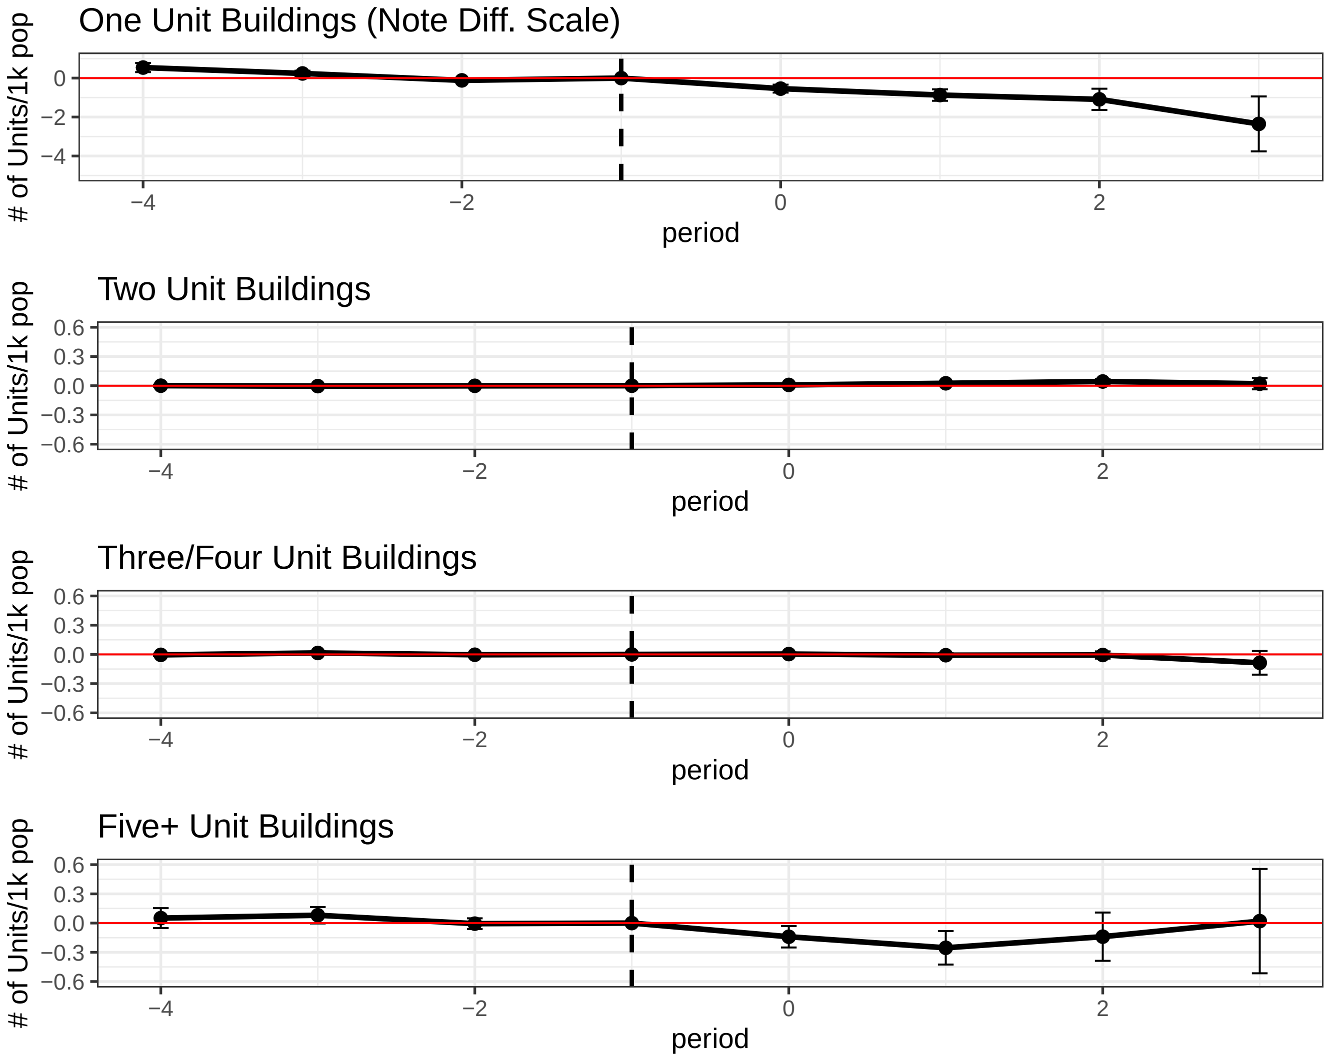

Note: This figure plots the impact of SC on residential permits (intended construction) by building size, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on permitted units of one-unit, two-unit, three/four-unit, and five/more-unit buildings. Estimations in this figure are based on the ACS subsample. We use total permitted units per 1,000 residents as the outcome variables. 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: Permits by Building Class, National Sample

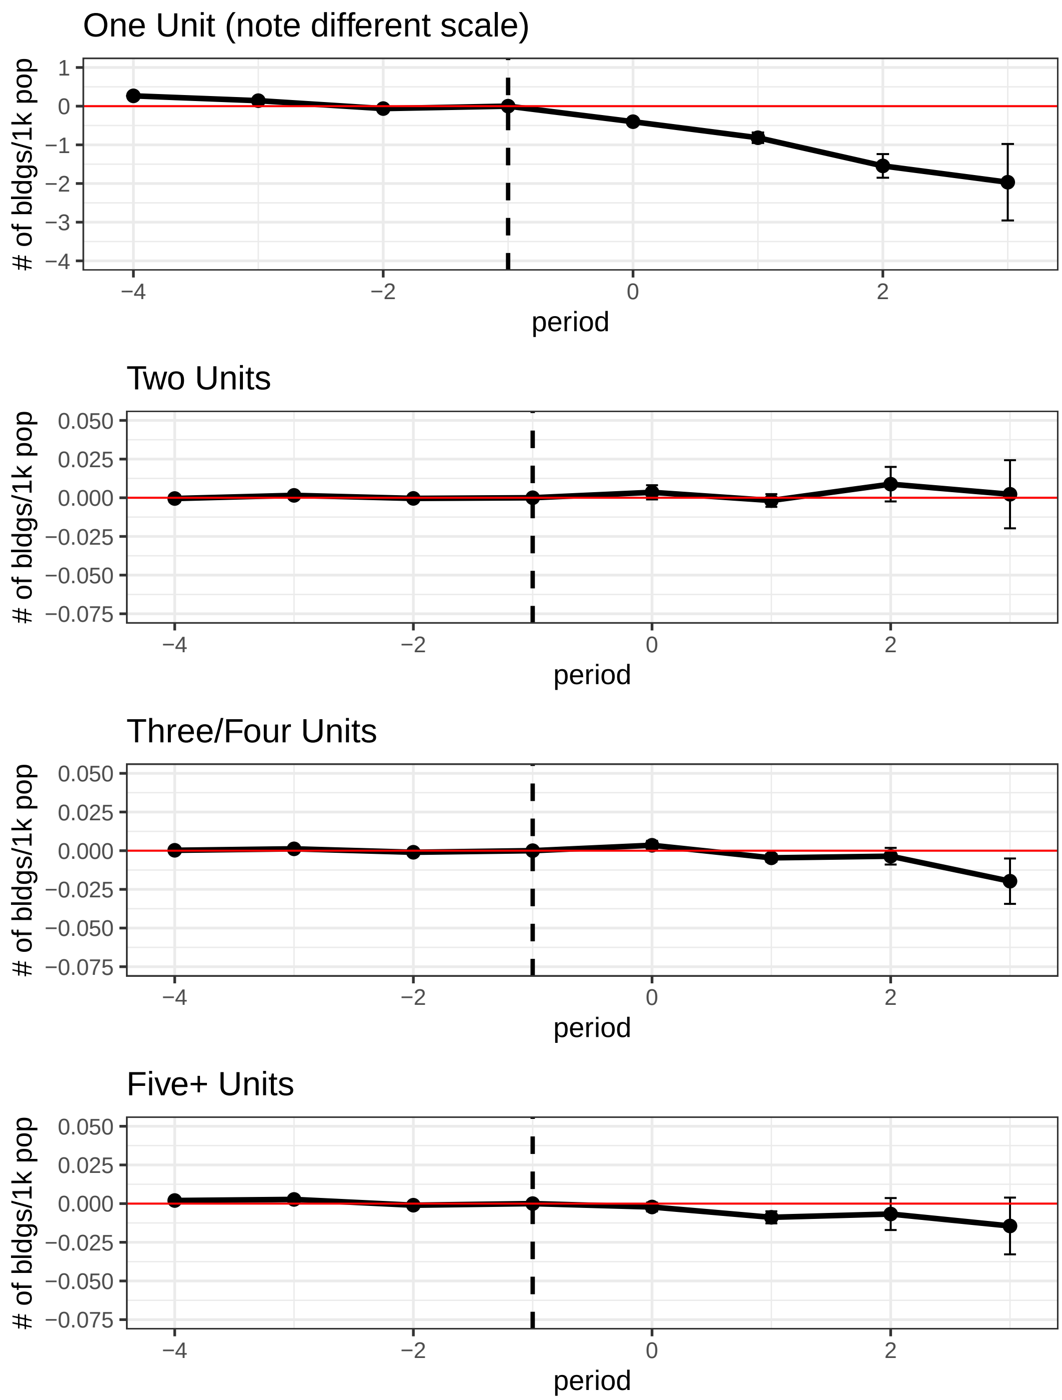

Note: This figure plots the impact of SC on residential permits (intended construction) by building size, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on permitting buildings of one-unit, two-unit, three/four-unit, and five/more-unit. Estimations in this figure are based on the national sample. We use total permitting buildings per 1,000 residents as the outcome variables. 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: Permitted Units by Building Class, National Sample

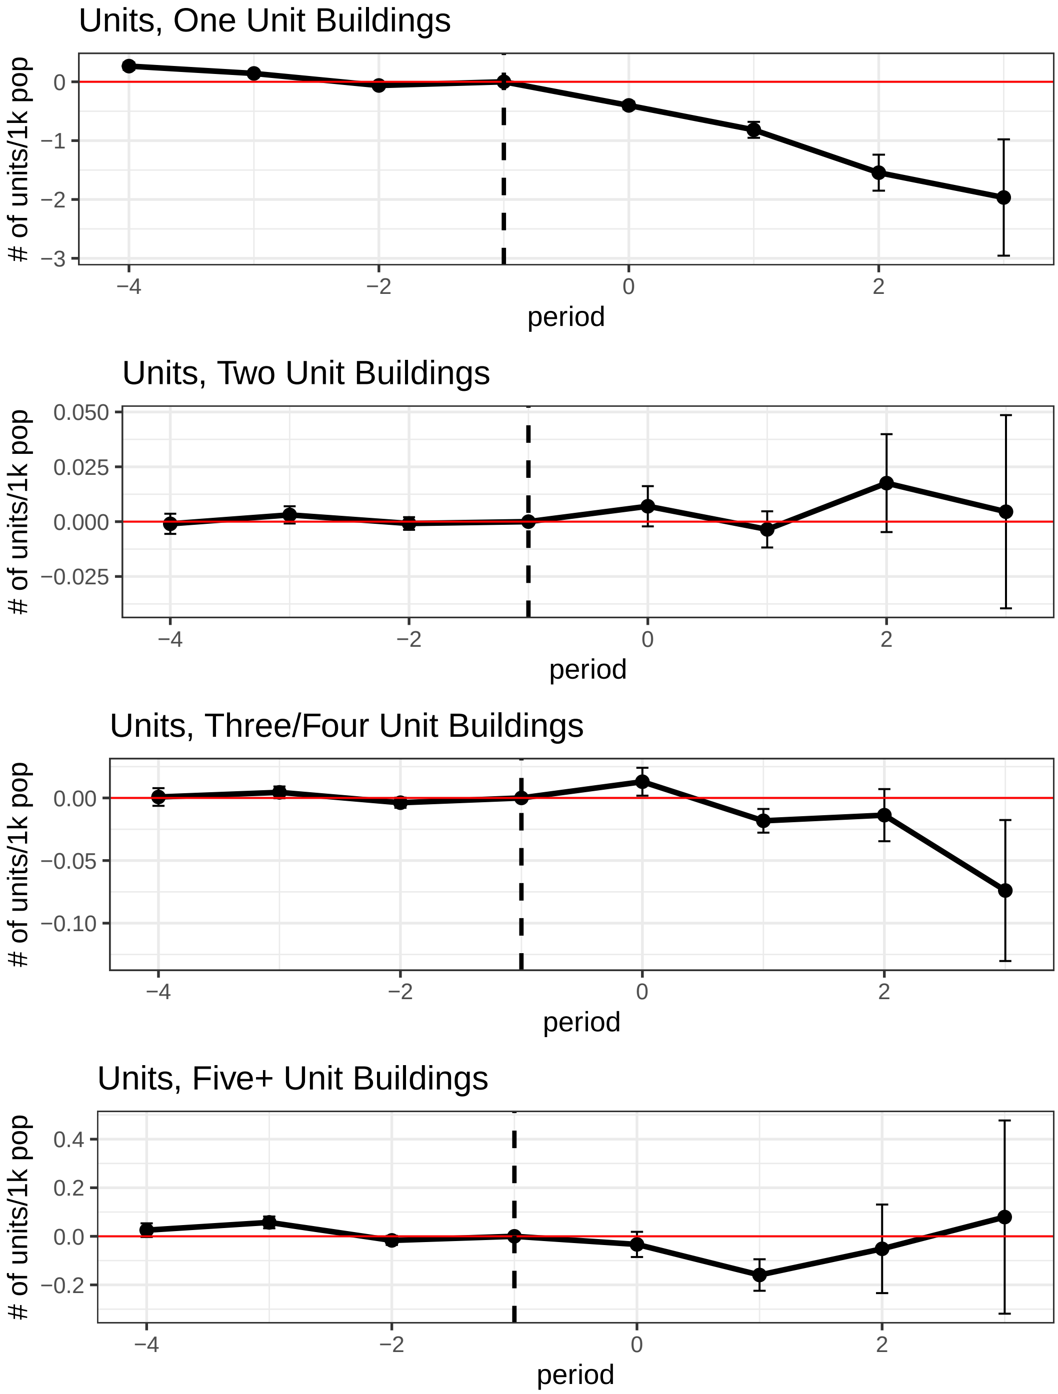

Note: This figure plots the impact of SC on residential permits (intended construction) by building size, with the approach of Gardner (2022) and specification (2). The four panels examine the impact on permitted units of one-unit, two-unit, three/four-unit, and five/more-unit buildings. Estimations in this figure are based on the national sample. We use total permitted units per 1,000 residents as the outcome variables. 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: New Construction, Without Square Footage Imputation

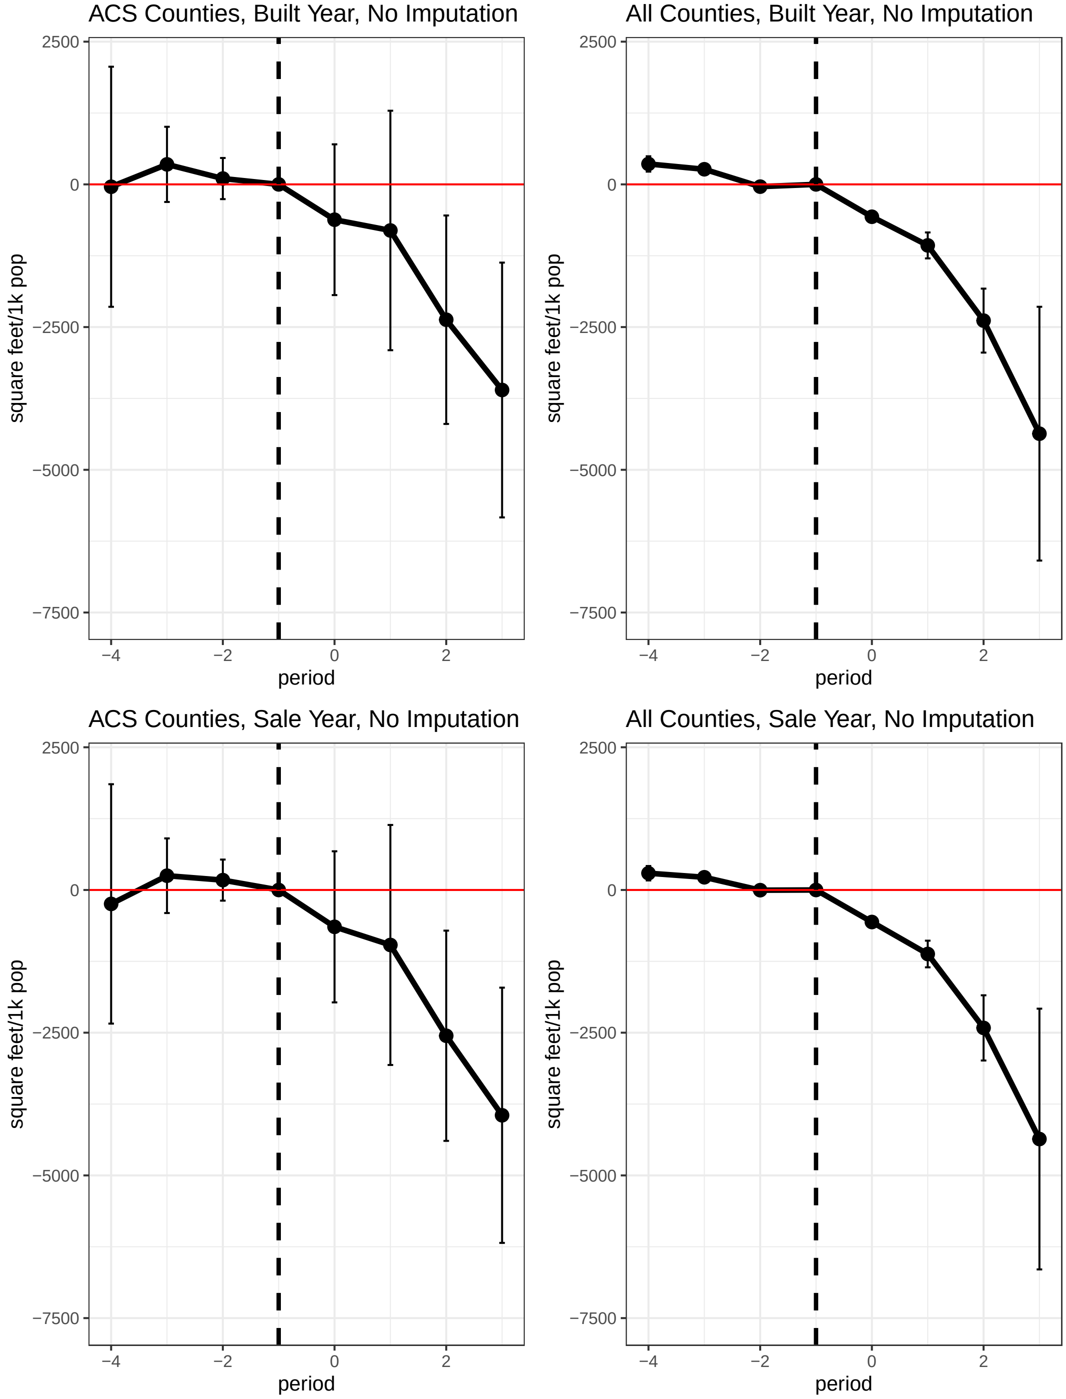

Note: This figure plots the impact of SC on residential construction activity measured by observed new construction (completed new construction), with the approach of Gardner (2022) and specification (2). This figure repeats the analysis of Figure 10 but does not impute square footage for any observation missing that information. 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 one period is a year.

Figure A6: New Construction Attribute Shifts, ACS

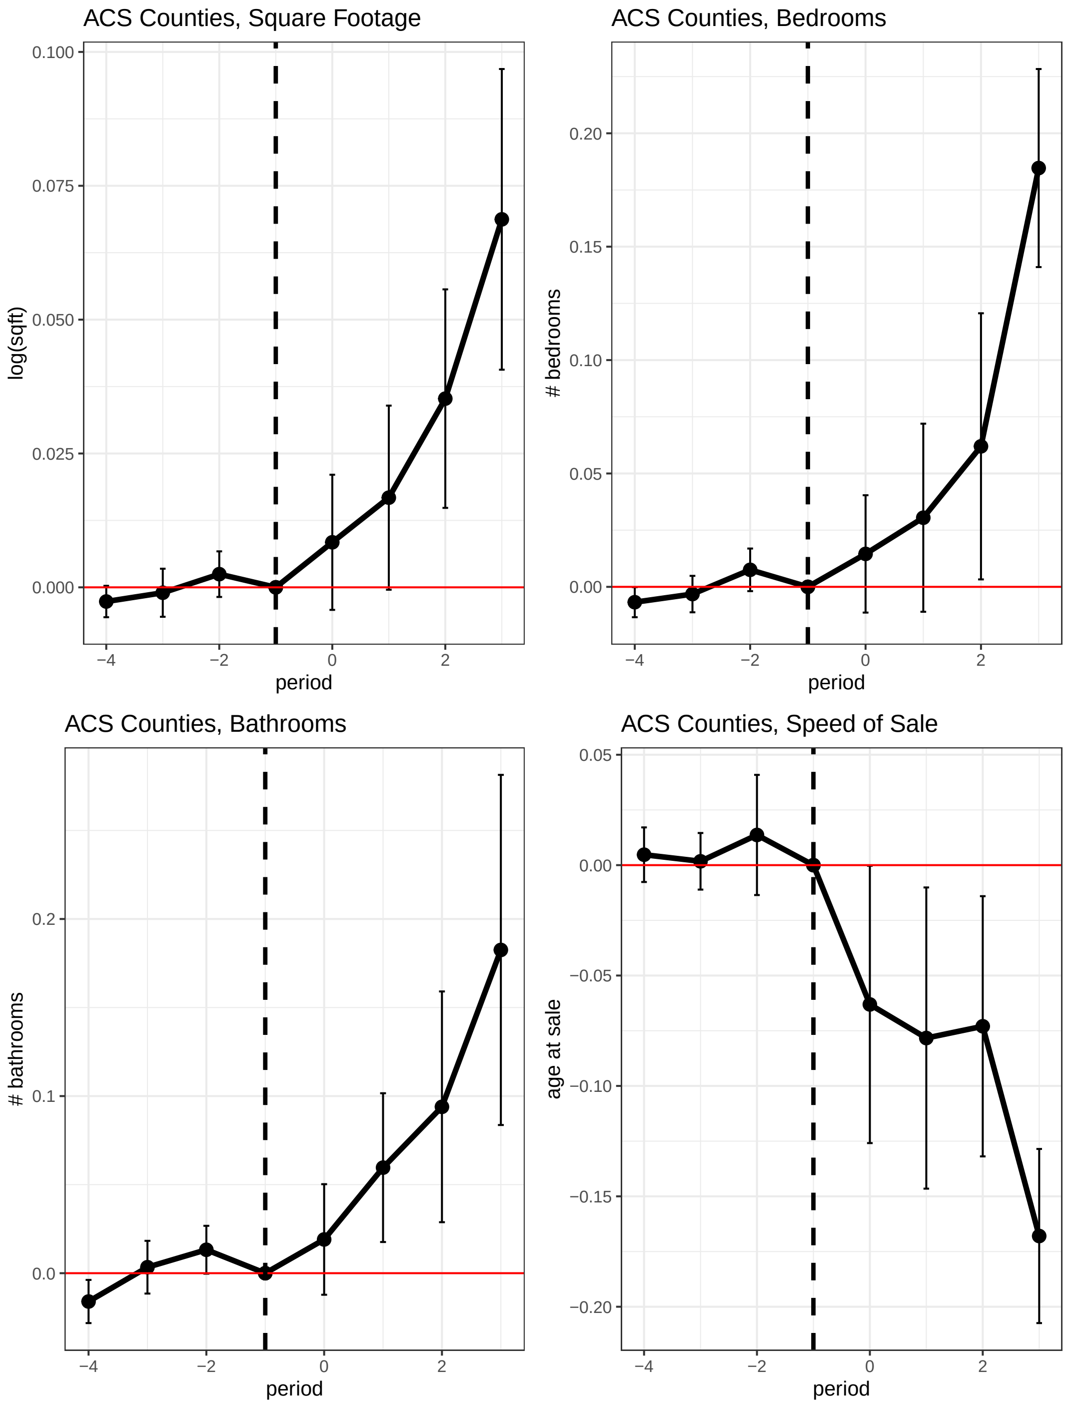

Note: This figure plots the impact of SC on various hedonic attributes of newly constructed properties, with the approach of Gardner (2022) and specification (2). Estimations in this figure are based on the set of counties which are separately identifiable in ACS data. In each panel, 95% confidence intervals 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 A7: New Construction Attribute Shifts, All Counties

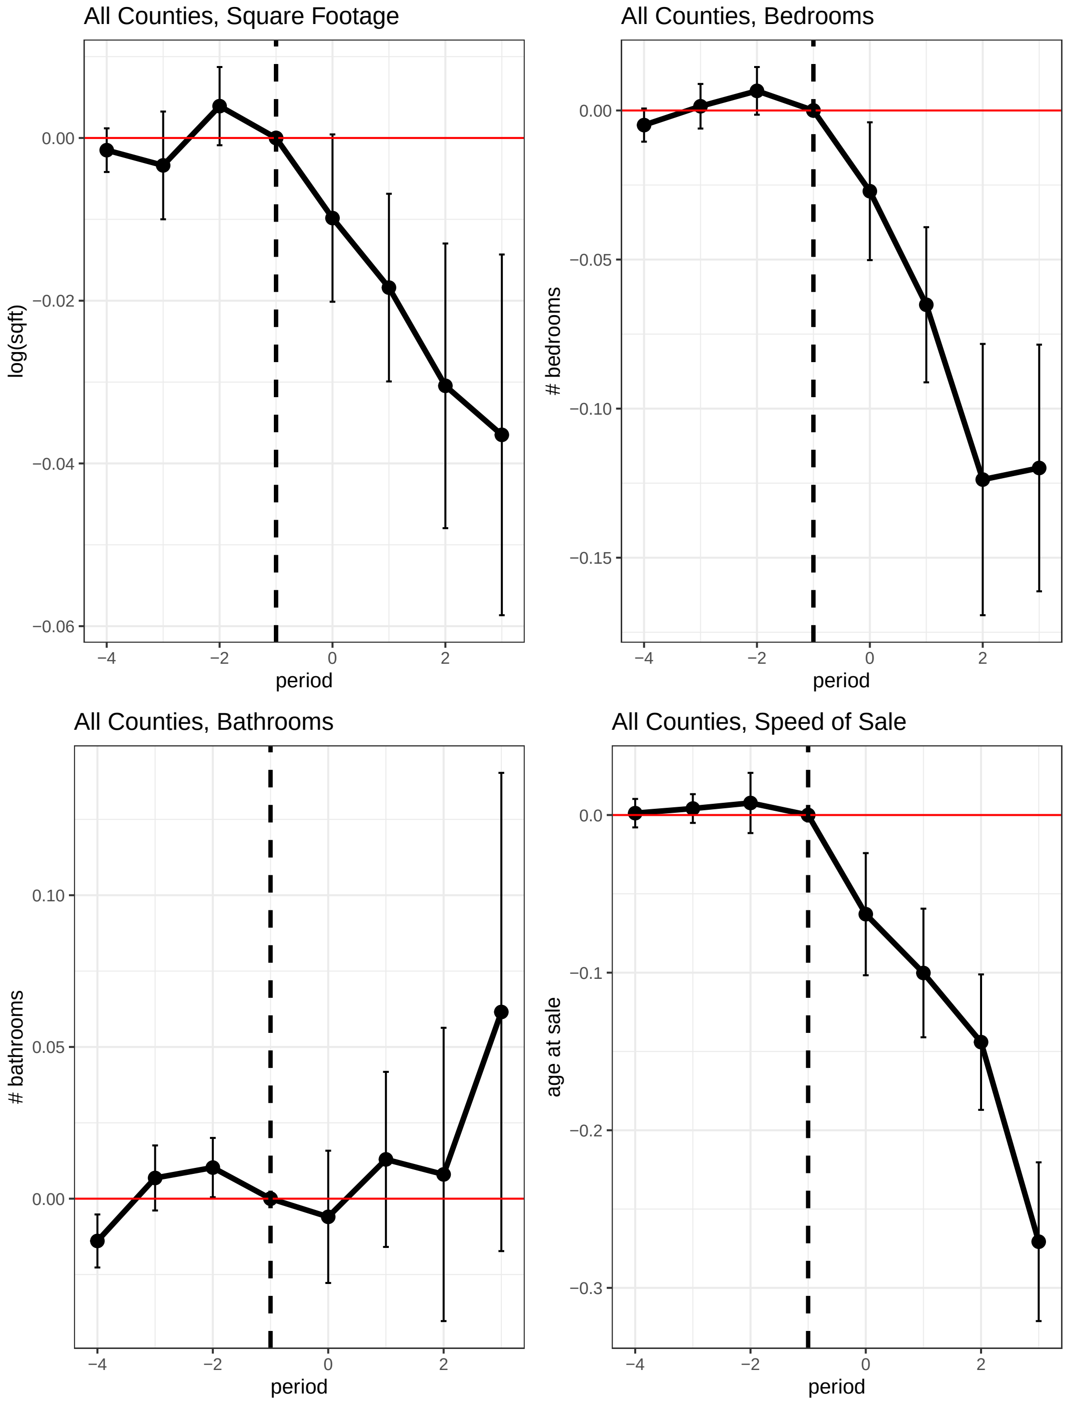

Note: This figure plots the impact of SC on various hedonic attributes of newly constructed properties, with the approach of Gardner (2022) and specification (2). Estimations in this figure are based on the full national sample. In each panel, 95% confidence intervals are plotted around the point estimations,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

Figure A8: Price Impact and Size Adjustment: Condos and Duplexes

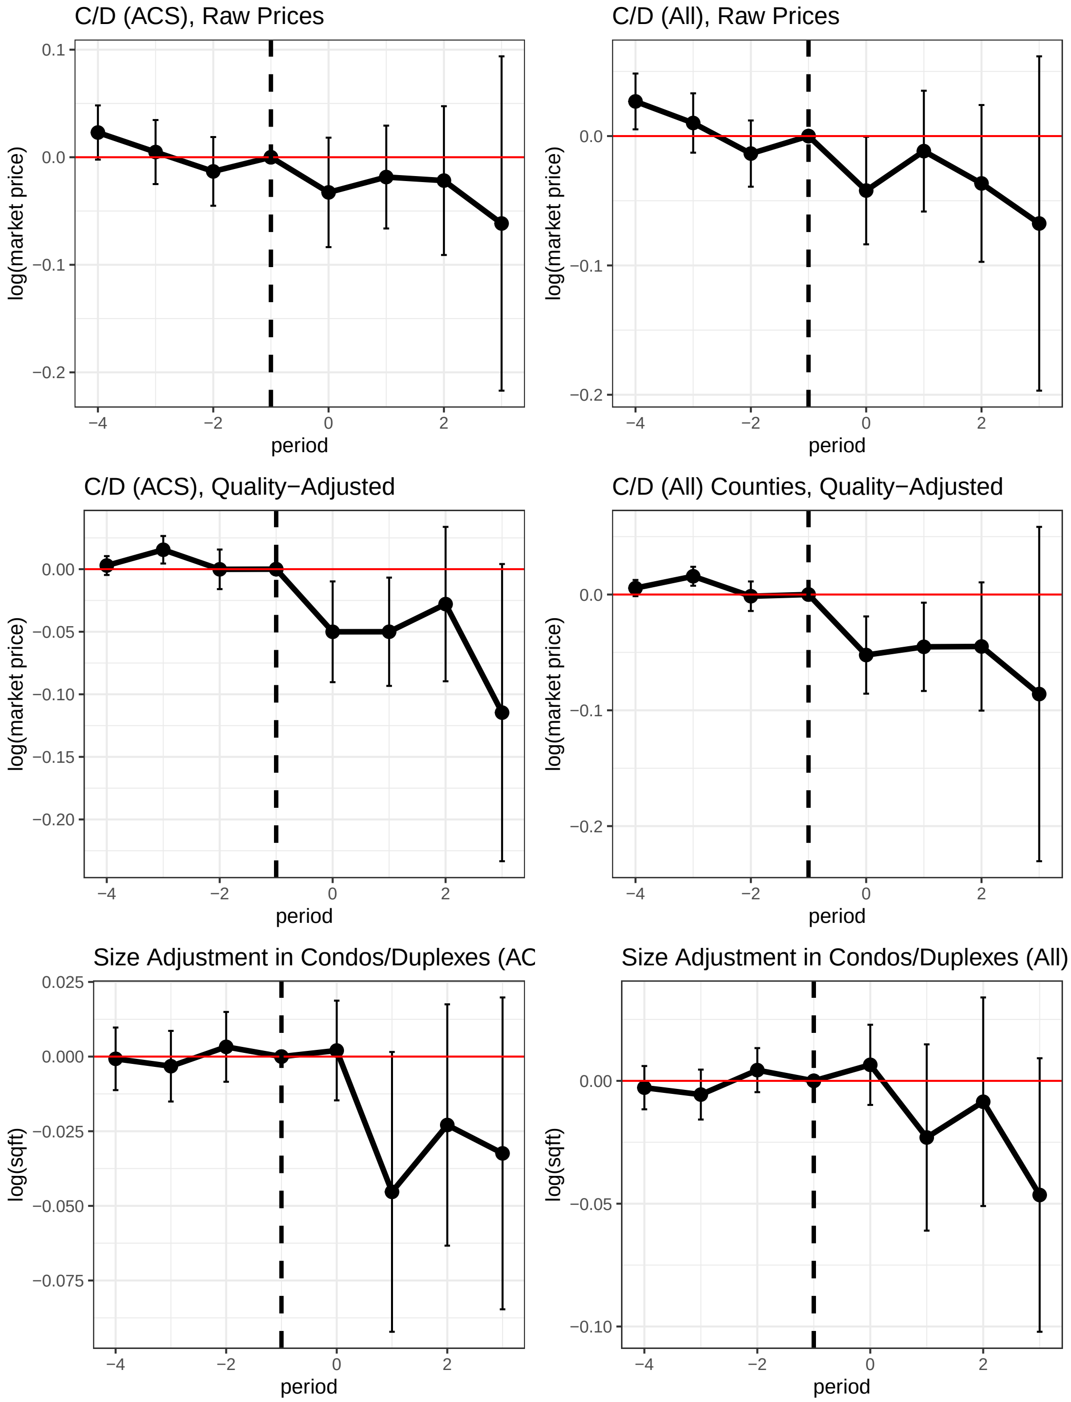

Note: This figure plots the impact of SC on market prices and size for condos and duplexes, with the approach of Gardner (2022) and specification (2). The first row uses raw prices. The second row adds hedonic controls and shows the impact on quality-adjusted prices. The third row shows the effect on unit size. The left columns use counties separately identifible in the ACS data, and the right column uses 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 A9: Quantity Impact for Condos and Duplexes

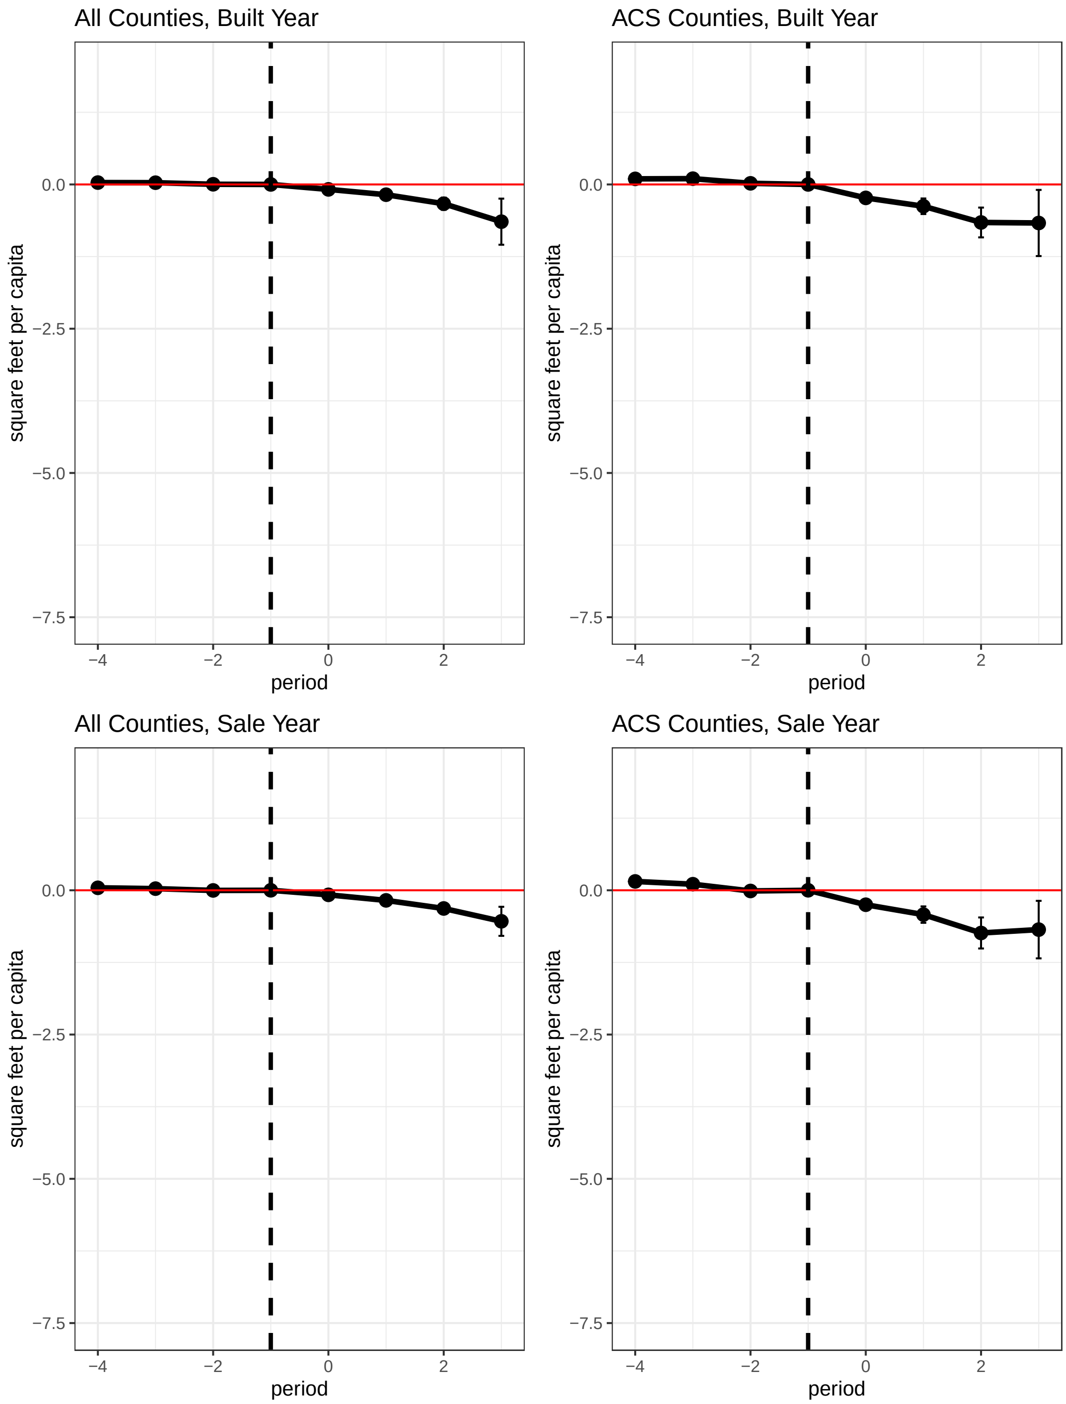

Note: This figure plots the impact of SC on condo and duplex construction, with the approach of Gardner (2022) and specification (2). The outcome variable is total square footage normalized by 2005 county-wide population. 95% confidence intervals are plotted around the point estimations, and the standard errors are clustered at the county level. The unit of one period is a year.