The Impact of Opportunity Zones on Housing Supply

By Benjamin Glasner, Adam Ozimek, and John Lettieri*

The United States suffers from a severe and persistent shortfall in housing. The Opportunity Zones (OZs) tax incentive is one of the most recent federal efforts to catalyze private investment in housing construction, with a specific focus on designated low-income communities. Using modern difference-in-differences methods, we estimate the causal effect of OZ designation on the local stock of active and vacant residential addresses, drawing on data from the U.S. Department of Housing and Urban Development's (HUD) Aggregated United States Postal Service (USPS) Administrative Data on Address Vacancies. We find that OZ designation caused a large and sustained increase in housing supply in designated low-income communities. We estimate that the incentive resulted in more than 416,000 new active and vacant residential addresses through Q1 2025, increasing the rate of growth in new housing within designated OZ tracts by 69.8 percent. The average effect per designated tract was 47.5 active or vacant addresses. These net new addresses largely reflect genuinely new development rather than a reallocation of investment activity from nearby non-designated areas. Our findings demonstrate that place-based capital-gains tax incentives can effectively stimulate private investment in local housing supply and help address persistent underinvestment in low-income communities.
Keywords: Opportunity Zones, Housing, Place-based Policy

Opportunity Zones (OZs) were established under the Tax Cuts and Jobs Act of 2017 with the goal of attracting long-term private investment into designated low-income communities across the United States.[1] As with earlier place-based and community development policies such as the New Markets Tax Credit (NMTC), the stated purpose of the OZ incentive is to improve economic outcomes in distressed areas. However, the design of the policy marked a substantial departure from legacy programs. Whereas previous incentives typically relied on capped

* Benjamin Glasner: Economic Innovation Group, benjamin@eig.org. Adam Ozimek: Economic Innovation Group, adam@eig.org. John Lettieri: Economic Innovation Group, john@eig.org. Thanks to Kenan Fikri, Cardiff Garcia, Nathan Goldschlag, Sarah Eckhardt, August Benzow, Jiaxin He for their support and advice in the completion of this project. This analysis would not have been possible without the data provided by Alexander Din with the U.S. Department of Housing and Urban Development. Thanks to Dan Garrett for his help in contextualizing Opportunity Zones relative to alternative policies. We have received no external funding in relation to this work, though the Economic Innovation Group was the source of the initial white paper proposing Opportunity Zones.

[1]Opportunity Zones were designed with a focus on low-income communities, but did allow for the selection of some contiguous tracts as well. Roughly 2 percent of OZ designations were low-income community adjacent tracts, while not themselves being low-income communities.

tax credits awarded through competitive, centrally administered processes, the OZ incentive was structured as an uncapped, market-driven mechanism intended to decentralize capital allocation and rely on private market actors to identify investment opportunities.

The OZ design raised a central policy question: could a market-oriented model "improve on the previous track record of place-based tax policies," which frequently failed to generate investment at scale or align private incentives with long-term local economic outcomes (Corinth and Feldman, 2024)?

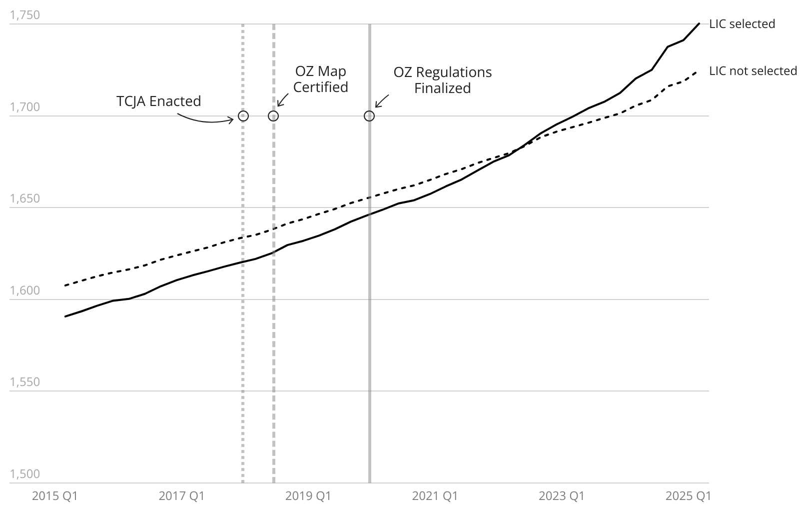

Figure 1. Average Active and Vacant Residential Addresses per Tract, Q1 2015 to Q1 2025

Note: In this figure, we are plotting the average count of residential addresses, active and vacant, in designated LIC tracts and undesignated and nonbordering LIC tracts.

In this paper, we evaluate one fundamental dimension of that question: the impact of OZ designation on a community's housing supply. Residential real estate constitutes a major category of OZ investment activity, and housing development exhibits economic characteristics—capital intensity, long asset duration, and sensitivity to after-tax returns—that align closely with the structure of the OZ incentive. If the theoretical mechanism of Opportunity Zones is functioning as intended, OZ designation should lead to increases in housing development within designated communities relative to similar non-designated areas.

Using administrative data from the U.S. Department of Housing and Urban Development's USPS Administrative Data on Address Vacancies, we measure changes in local housing supply as captured by net growth in active and vacant residential addresses. We find that OZ designation has had a statistically and economically significant effect on the stock of residential addresses, increasing net new active and vacant residential addresses in a designated low-income community

(LIC) tract by 47.5 on average as of Q1 2025. We estimate that 41% of all new residential active and vacant addresses built within treated tracts during this period can be directly attributed to the incentive, equivalent to a 69.5% net increase in new housing supply within our sample.[2] Generalizing our average effect estimate across the full count of 8,764 census tracts designated as OZs across states, territories, and the District of Columbia implies that the policy resulted in roughly 416,000 new residential addresses as of Q1 2025.

I. What are Opportunity Zones?

Opportunity Zones (OZs) are a novel example of a federal place-based tax incentive. OZs were established under the Tax Cuts and Jobs Act of 2017 with the goal of redirecting latent private capital toward long-term equity investment in economically distressed communities by reducing the tax friction that discourages investors from realizing and reinvesting capital gains.[3]

The OZ incentive is not a tax credit. Instead, it allows individual and corporate taxpayers to defer and partially reduce capital gains tax liabilities when such gains are reinvested in Qualified Opportunity Funds (QOFs), which deploy capital into qualifying assets located in designated census tracts. The third and most significant tax benefit—a full exemption on any appreciation in the value of a qualifying investment held for at least ten years—links investor benefits to the long-run economic performance of the underlying assets and is intended to enhance the after-tax returns of projects that might otherwise be financially marginal or unattractive.

While the OZ incentive was enacted in late 2017, the OZ designation process occurred in mid-2018 using eligibility criteria derived from the "low-income community" (LIC) definition established under the New Markets Tax Credit (NMTC) program. Eligible census tracts were generally those with poverty rates above 20% or median family incomes below 80% of the relevant area median.[4] State governors were authorized to nominate up to 25%[5] of their eligible tracts for designation,

[2]Our analytic sample comprises 7,580 designated low-income community (LIC) tracts out of a total of 8,764 designated tracts. We estimate the average effect of OZs as 47.5 new active/vacant addresses per tract. Multiplying by 7,580 LIC OZ tracts yields 360,048 new addresses caused by OZs within the sample. The total number of new addresses within this sample of 7,580 tracts was 875,528 from Q1 2019 to Q1 2025, implying that 41.12% of all new residential active/vacant addresses in OZs were the result of the policy. The within-sample effect estimate implies that 515,480 new active and vacant addresses would have been added over this period absent the policy, so the additional 360,048 represents a 69.8% increase.

[3]The broad concept behind OZs was first outlined in a white paper by Bernstein and Hassett (2015) published by the Economic Innovation Group (EIG). In 2016, a bipartisan and bicameral group of legislators introduced the first OZ legislation, entitled *The Investing in Opportunity Act*, which was later incorporated with minor modifications into the Tax Cuts and Jobs Act (TCJA).

[4]The statute also allowed a small share of non-LIC tracts to be designated under a "contiguous tracts" exception. Such tracts must be contiguous with a designated LIC and have a median family income that does not exceed 125% of the adjacent LIC's median family income. In total, 2.6% of all OZ designations were made under this exemption.

[5]The statute included a special rule for Puerto Rico, automatically designating all qualifying tracts as Opportunity Zones.

subject to certification by the U.S. Department of the Treasury.[6] In total, 8,764 tracts were designated as Opportunity Zones across states, territories, and the District of Columbia, representing approximately 12% of all U.S. census tracts and home to about 10% of the national population, with designations to remain in effect until December 31, 2028. These tracts displayed higher levels of economic distress relative to national averages: on average, they had a poverty rate of 29%[7] and median family incomes approximately 40% lower than the national median.[8] Over 97% met the statutory LIC definition, and more than 70% also satisfied the Treasury Department's more stringent "severely distressed" criteria (Fikri and Lettieri, 2018). Moreover, despite the broad latitude governors were afforded, designated tracts exhibited consistently worse economic characteristics than those that were eligible but not selected (Fikri and Lettieri, 2018).

The zone designation process was completed in June 2018 when the U.S. Secretary of the Treasury certified nominated tracts. Subsequent regulatory guidance was issued in three waves between October 2018 and December 2019.

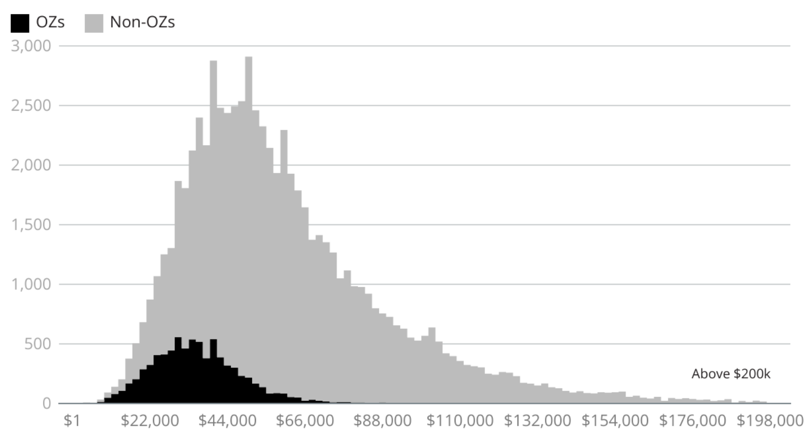

Figure 2. Number of OZs and Non-OZs grouped by Median household income Band

Note: 2017 ACS 5-year estimates

To qualify for OZ tax benefits, investments must be made through a QOF and meet statutory requirements intended to ensure that the incentive supports new productive activity rather than mere transfers of ownership or financial engineer-

[6]This selection process by governors raised questions about the comparability of selected and unselected tracts due to the endogenous nature of treatment assignment. For more information, see Frank, Hoopes and Lester (2022). We address this issue through the use of conditional parallel trends and test for the validity of the parallel trends assumption in our analysis.

[7]U.S. Census Bureau's American Community Survey, 2012–2016 5-Year Estimates.

[8]The median family income of OZs was $42,400 while the national median income was $67,900.

ing. Specifically, a qualifying investment must either initiate the "original use" of property within an OZ—such as through new construction—or satisfy the "substantial improvement" test, which requires the taxpayer to at least double the adjusted basis of an existing asset following acquisition. These provisions are economically significant, as they constrain eligible activity to projects that increase the local capital stock.

In the context of housing, these requirements directly incentivize new development or major rehabilitation, rather than simply rewarding property transactions. Moreover, because the greatest tax benefit accrues only after a long hold period, the incentive structure encourages patient equity capital and favors projects that generate value through long-term appreciation, a characteristic that closely aligns with the economic nature of housing development.

OZs share a broad policy objective with earlier place-based tools, including the NMTC and Enterprise Zones, which aim to stimulate investment and economic activity in distressed areas (Bernstein and Hassett, 2015; Neumark and Simpson, 2015). However, the OZ incentive departs from these models in several critical respects. The first is the delegation of selection authority to governors described above, which is consistent with the theory that states had an informational advantage in identifying low-income communities most likely to benefit from private capital formation. The second key contrast is in the structure of the incentive itself. Unlike the NMTC, which is subject to an annual allocation cap and distributed through a competitive application process, the OZ incentive is uncapped and by-right: any qualifying investment automatically receives the tax benefit. These low administrative barriers could enable capital aggregation and rapid deployment, allowing investors to diversify across multiple assets and geographies. In addition, whereas NMTC investments predominantly originate from large corporate taxpayers motivated by fixed tax credit awards (Fikri, Benzow and Lettieri, 2023; *see also placeholder* **?**), participation in OZs has come from a broader mix of individual and institutional investors (Coyne and Johnson, 2023) whose returns depend on the subsequent performance of the underlying assets. The uncapped and investor-led structure has the potential to provide a greater scale and speed of capital mobilization (*add citation here*).

The scale and geographic reach of OZ investment reflect these foundational design differences. As of the 2022 tax year, Qualified Opportunity Funds reported holdings of $89 billion in qualified OZ property across more than two-thirds of all zones (Corinth et al., 2025). Investments in QOFs have been raised from a large number of investors spread across every U.S. state (Kennedy and Wheeler, 2021). IRS administrative data indicate that by 2020, approximately 21,000 individuals and 4,000 corporations had participated in the incentive (Coyne and Johnson, 2023). Compared to the NMTC program—which channels more than 95% of its investment from corporate taxpayers and is constrained by an annual allocation cap (Fikri, Benzow and Lettieri, 2023)—OZs have mobilized capital at a substantially larger scale and across a significantly greater number of communities in spite of the policy's much tighter geographic constraints.[9] From 2018 through 2020, roughly 3,800 tracts received at least one OZ investment, equivalent to the total number of tracts that saw NMTC investments over the first 18 years of that program's existence (Fikri, Benzow and Lettieri, 2023). Corinth et al. (2025) observe a similar pattern in more recent data: between 2019 and 2022, 5,669 tracts received OZ investments, compared to 1,259 tracts that received NMTC investments over the same period. This elevated level of participation is consistent with the policy's structural features and indicates that the incentive was capable of operating at a scale sufficient to influence local economic outcomes, including housing supply.

The OZ policy's design features establish a clear theoretical mechanism by which OZs may influence local economic outcomes, and in particular the housing market. By improving the after-tax return profile of long-term, equity-financed investments, the OZ incentive can provide a direct channel through which capital formation may increase in designated areas. This mechanism is particularly relevant in the context of housing markets, because housing development is capital intensive and characterized by long investment horizons, and therefore should be sensitive to changes in tax-adjusted returns.

II. Literature Review

Research to evaluate the economic impacts of the OZ incentive in the years immediately following the TCJA should focus on early outcomes that are "upstream" from revitalization (Fikri and Glasner, 2023). Much of the literature has taken the opposite approach, looking for "downstream" outcomes that could only be expected far down the road of revitalization.

For example, using restricted American Community Survey microdata, Freedman, Khanna and Neumark (2023) find no impacts on employment, income, or poverty in the immediate aftermath of OZ designation.[10] Unfortunately, the data they use ran from 2013 to 2019—a period ending just as OZ regulations were being finalized. Moreover, OZs are a community reinvestment incentive, not a direct subsidy to low-income individuals that could plausibly result in immediate and significant decreases in local poverty or income. A study of individual-level outcomes as of 2019, therefore, does little to illuminate the effectiveness of OZs; it is looking for effects before any could reasonably be expected to exist.

Atkins et al. (2023) find no overall impact on job postings using Burning Glass data, but do find small positive effects in urban areas, where the majority of OZ investment has gone. Arefeva et al. (2024) find a positive impact on employment

[9] NMTC investments can be made in any LIC, whereas OZ investments are limited to designated LICs only, which at the time of designation in 2018 represented only about 25% of all LICs nationwide.

[10] As new buildings are constructed, local poverty at the tract level could shift partly due to a dilution of poverty (more residents above the poverty line moving in), a substitution of poverty (those in poverty moving out while individuals above the poverty line move in), and a reduction in poverty (improvements in neighborhood characteristics increasing the incomes of residents who were in poverty before the OZ designation occurred). These contrasting mechanisms have yet to be tested in the literature.

using Infogroup establishment data through 2021. A recent working paper by Freedman, Kouchekinia and Neumark (2025) using LEHD Origin–Destination Employment Statistics (LODES) for 2013–2022 also reports significant increases in employment within designated tracts, but not for residents of a designated tract.

Studies looking at upstream effects tend to find positive results when focused on variables with a strong theoretical justification based on the nature of the OZ incentive: outcomes related to the early stages of community redevelopment. A working paper by Wheeler (2022) found that OZ designation resulted in a significant increase in the likelihood of development activity, based on an examination of building permits across a sample of 47 large cities and 12,000 neighborhoods from January 2014 through June 2022. This effect was particularly strong in communities with more available land and in-fill opportunities, a more elastic housing supply, and lower home values. Wheeler (2022) estimates that tracts designated as OZs saw an increase in the probability of new development in a given month of 2.9 percentage points, a 20.5% increase from baseline.

Other studies suggest positive impacts on real estate development—and, in particular, multifamily residential buildings (Coyne and Johnson, 2023; Corinth et al., 2025; Sciarretti, 2023). New residential construction could have a range of impacts on important long-run local outcomes and would be one of the first signs of broad revitalization.

Some upstream studies, however, focus on outcomes that lack a clear theoretical interpretation. For example, housing prices could be a plausible upstream outcome to study because they can change in anticipation of economic activity. But prices can also change with the supply and quality of goods, making the expected direction of change unclear. If the primary effect of OZs is to increase housing supply, for example, this could lower prices in both OZs and in surrounding tracts. If OZs trigger an increase in amenities, then amenity-driven demand could outweigh supply impacts and increase prices. The lack of a clear prediction on prices makes it impossible to interpret the price effects found in a study if the same study fails to account for quantities or the nature of goods—for example, Chen, Glaeser and Wessel (2023).[11]

Wheeler (2022) measures the effects of OZ designation on both prices and development activity. He finds not just a rise in development activity, as already discussed, but also an appreciation in home values. This offers a sign that amenities may have improved, making homes and neighborhoods more attractive to buyers. Interestingly, Wheeler (2022) also finds no increase in rents as a result of OZ designation, suggesting that supply effects outweigh amenity effects for rental units.

Transaction volume is another upstream outcome that lacks a theoretical jus-

[11]Interestingly, when Wheeler (2022) explored the null effect reported by Chen, Glaeser and Wessel (2023), he found that their null result was driven by the use of price growth rates rather than price levels or log price levels. Wheeler (2022) reported a 3.4% increase in median home values in a subset of urban OZs from 2017 to 2020.

tification for its usefulness, in isolation, as a measure of whether OZs are having their intended effects. Feldman and Corinth (2023)[12] and Sage, Langen and Van De Minne (2023) investigate the impact of OZs on a dataset that comprises nearly the full universe of commercial transactions above $2.5 million. But OZs only increase the returns to investments that are new construction or include "substantial improvements" equal to the initial value of investment.[13] Aggregate transaction volume across OZs might appear unchanged, but the nature of the underlying transactions could be very different as a result of the policy. A shift toward redevelopment may well be lost in the noise of ineligible and unrelated transactions. Even if transaction volume is meaningfully changed by the incentive, these aggregate measures may miss it in the noise of unqualified activity. When focused on the more narrow question of multifamily transaction volume, **?** find a positive effect despite measuring the full universe of transactions rather than those intended for reinvestment.

One initial question regarding the effectiveness of the OZ incentive model was whether its lack of any requirement for government intermediaries to pre-approve or direct investments would result in investors avoiding the vast majority of designated tracts in favor of a small share that are the least distressed and most attractive. The available evidence instead shows that a large and growing share of OZ tracts are seeing investment, and that "OZs [are] providing a large amount of investment to distressed areas" (Corinth et al., 2025).

As of 2022, fully two thirds of designated tracts had already received investment from a QOF, and OZ investment activity was inversely related to economic well-being at the tract level (Corinth et al., 2025).[14] Indeed, even without the sort of government intermediation required to approve investments and award tax benefits found in the NMTC program, OZ investments target areas with similar levels of economic distress. Analysis of Coyne and Johnson (2023)'s findings

[12]Feldman and Corinth (2023) also introduce a regression discontinuity (RD) design that uses both the tract poverty rate and median income as running variables, marking a departure from the difference-in-differences designs favored by most work on OZs. In a traditional RD, a comparison is made on either side of a treatment threshold. However, in Feldman and Corinth (2023), a multivariate measure of OZ eligibility cutoffs is used. Given the Real Capital Analytics (RCA) data employed, the RD is designed to detect whether OZ *eligibility* led to changes in commercial investment on either side of a multivariate eligibility threshold. In a simplified univariate context, an RD would test for differences between census tracts with a 20.1% poverty rate and those with a 19.9% poverty rate (i.e., on either side of the 20% eligibility cutoff), for example. This approach creates three issues. First, an RD is less well-suited to detect treatment effects across the entire sample, including in higher-poverty census tracts where the impact of the incentive may be less marginal or less meaningful, thus offering limited capacity for heterogeneity analyses. Second, an RD relies on precision in the running variable, yet aggregated poverty rates lack consistency in measuring the depth of poverty, compromising our ability to identify two tracts that are truly similar in terms of poverty exposure. Third, discontinuity models struggle to control for spillover effects across geographic units, which other studies have shown to be significant for OZs (Arefeva et al., 2024; Wheeler, 2022).

[13]See Internal Revenue Service (2024).

[14]Corinth et al. (2025) find that, at the tract level, OZ investment and NMTC investment both targeted a similar level of distress. When they expand the analysis to the county level, they observe that a greater share of investment went to low-income tracts embedded in less distressed counties compared to NMTC investments. It is unclear how much of this result is driven by price differentials across counties given the breadth and scale of OZ investment.

on tracts that received OZ investment through 2020 reveals that they were, on average, in the 87th percentile for poverty, the 81st for median family income, and the 80th for unemployment among all U.S. tracts. These results suggest that investor behavior—not simply OZ designation—largely aligns with the intent of the policy to target areas that lag far behind the typical U.S. community in key economic metrics.[15]

Corinth et al. (2025) find that OZ investment tended to target census tracts with a higher volume of commercial transactions from 2013 to 2017, suggesting that the OZ incentive structure favors places that demonstrated investment potential over ones "not ripe for productive investment on their own." It is not clear, however, that this is out of step with either the policy's intent or the optimal use of a community development capital-gains incentive. If, for example, OZ investment were instead concentrated in places without the capacity to use it, a separate criticism could be made that taxpayer dollars were subsidizing unproductive investments, which has long been a critique of place-targeted subsidies (Glaeser and Gottlieb, 2008).

Corinth et al. (2025) interpret the pre-existing transaction volume in OZ tracts as evidence of the "OZ tax incentive rewarding investment that would have occurred in the absence of the program." On its own, however, such data cannot establish whether OZ designation caused subsequent investment to increase. This is precisely the question this paper addresses in the context of residential addresses.

In short, the results of the literature are consistent with an initially positive effect of OZs on upstream factors of redevelopment. It remains unclear if these upstream factors will translate into economically significant impacts on downstream outcomes.

Our paper contributes to the literature on two important dimensions of federal policy. The first is the design of place-targeted tax incentives—in particular, whether a novel capital-gains tax incentive model can succeed in motivating investment in low-income communities at a meaningful scale. The second is the literature on federal policy to increase housing supply. In recent decades, the government has shifted its housing supply-side policy from directly providing it through public housing projects to subsidizing construction (Collinson, Ellen and Ludwig, 2015). The largest supply subsidy is the Low-Income Housing Tax Credit (LIHTC), which promotes the development and rehabilitation of affordable rental housing. LIHTC awards overwhelmingly go to corporate investors, with an estimated cost of roughly $14 billion in 2025 rising to nearly $16 billion in 2028 (Joint Committee on Taxation, 2024). There is also growing awareness that zoning reform is an important policy lever for housing supply (Gyourko, Hartley and Krimmel, 2021; Been, Ellen and O'Regan, 2025). As zoning is largely set at the state and local level, however, there has been relatively little federal policy along this important margin.

[15]Ranked from lowest to highest levels of need, per Fikri, Benzow and Lettieri (2023).

III. Data

A. Data Sources

The primary dataset used in this analysis is the quarterly count of residential, business, and other addresses in a given census tract. When new buildings are constructed, the count of addresses in the neighborhood will grow. An exception arises when older and potentially dilapidated buildings are torn down and replaced with new buildings with the same number of units, in which case the amount of new construction would be underestimated by net address growth. While it will miss revitalizations, we use net address growth as a lower-bound estimate of new construction.

We measure net address growth using the U.S. Department of Housing and Urban Development's (HUD) Aggregated United States Postal Service (USPS) Administrative Data on Address Vacancies.[16] These data are gathered by the USPS to facilitate mail delivery and provide quarterly counts of addresses serviced by USPS. There are several types of addresses provided, which are defined by HUD as follows:

Total Number of Addresses reflects all addresses (residential and business) that are recorded by the USPS.

Total Vacant Addresses are addresses that delivery staff on urban routes have identified as being vacant (not collecting their mail) for 90 days or longer.

Total No-Stat Addresses are addresses that can be classified as "No-Stat" for many reasons, including:

- rural route addresses that are vacant for 90 days or longer;
- addresses for businesses or homes under construction and not yet occupied;
- addresses in urban areas identified by a carrier as not likely to be active for some time.

The data are collected by postal workers with the primary goal of supporting the efficient delivery of mail, not for research purposes.

One potential issue with the data, which is not problematic for our analysis, is that it can be subject to volatility as a result of changes in USPS policy. Most notably, in 2011 and 2014, address counts increased significantly due to the introduction of "Move to Competitive (MTC) Street Addressing for PO Boxes," which allowed customers to register PO boxes as street addresses. This change, however, preceded the designation of OZs by a large enough margin to be unconcerning. To avoid any pollution to our analysis, we focus on data quarters following the first quarter of 2015.

[16]The data we use run from the first quarter of 2014 to the first quarter of 2025. They were provided to us by Alex Din, a Social Science Analyst at the U.S. Department of Housing and Urban Development. The data provided were HUD's internal 2020 tract-definition standardized data.

Another data issue is that tracts with high growth or high decline both tend to have high rates of no-stat addresses. No-stat addresses are both empirically more volatile and also difficult to interpret as genuine new housing units. However, the data allow for active and vacant addresses to be combined into a total that excludes no-stat addresses. This is the approach we take.

Finally, USPS also captures an "other" category of address type, which in many ways is a black box. Since this analysis is focused on housing as the outcome of interest, we exclude both "business" addresses and "other" addresses.

Despite these issues, the administrative nature and regularity of the data offer a clear advantage. As compared to the American Community Survey (ACS), for example, this is not a sample but represents the universe of addresses serviced by USPS. In addition, we can compare HUD's address counts to external data sources. Changes in HUD's active residential addresses for the U.S. track closely to decadal changes in occupied housing units from the Decennial Census and to estimates of new residential construction from the Census Survey of Construction (SOC) and the Building Permits Survey (BPS).[17]

We also use data on which census tracts were eligible for and designated as OZs. One complication is that, while OZs were based on the 2010 census-tract boundary definitions, we use the 2020 standardized HUD data and map those data back to the 2010 tract definitions.[18]

To support our analysis and introduce conditional parallel trends, we measure a variety of socioeconomic outcomes at the tract level using 5-year American Community Survey (ACS) data from 2012 to 2023. For quarters in 2024 and 2025, we carry forward the most recent values observed in the ACS data. These measures include tract-level poverty rates, median household income, unemployment rates, the share of prime-age residents, and the share of a tract's housing stock classified as "solo detached" to proxy for zoning and neighborhood characteristics. We also include an index of local zoning regulations to ensure that we are not contrasting different regulatory environments.[19] These data are intended to help ensure that tracts are compared to valid control units when constructing our difference-in-differences analysis.

B. *Geographic heterogeneity*

One of the contributions of this analysis beyond the national scope of the address count assessment is the testing for geographic heterogeneity in the effect of

[17]This validation serves as a robustness check for this paper. Replication code is available in the project's GitHub repository (linked here). We use the change in HUD active residential addresses from the second quarter of 2010 to the second quarter of 2020 to align with the April 1 timing of the Decennial Census. The correlation is $0.X$, and rises to $0.9x$ if we exclude negative values.

[18]The 2020 standardized HUD data were provided by Alexander Din with the U.S. Department of Housing and Urban Development. These data are not publicly posted, though they should align with a cleaned and standardized version of the online data posted by HUD. Individuals looking to replicate our results are welcome to contact Benjamin Glasner or Alexander Din.

[19]See Bartik, Gupta and Milo (2024), Replication code is available in their GitHub repository.

OZs. Using a novel tract-level geographic coding scheme defined by a population-weighted assignment, we classify census tracts into one of the following categories: large urban, mid-sized urban, small urban, suburban, small town, and rural. We test for the effect of OZs both in aggregate and across distinct geographic subsets to explore how the policy's effectiveness may vary across types of places.

Classifying a tract along the urban–rural spectrum is no simple task. Block-group population data from the 2017–2021 American Community Survey (ACS) are paired with 2021 locale classifications from the National Center for Education Statistics (NCES)[20] to calculate the locale in which most people in a tract live. We then determine where a tract sits on the urban–rural spectrum according to the community characteristics where residents tend to live. We favor this classification scheme given our research focus on residential addresses.

The urban definitions in the NCES framework are adopted here to create three groups of urban tracts: large, mid-sized, and small. At the other end of the spectrum, the NCES definitions for the continuum of rural and small towns are collapsed into two definitions, depending on whether the majority of tract residents are in a small town or a rural area. Between these two poles, suburban tracts are based on the share of the population living in a medium or large suburban area. The explicit definitions are as follows:

Large urban At least 50% of the tract population is in a large urban area, *and* the tract is in a large urban county. Classified as suburban if at least 50% of the tract population is in a large urban area but the tract is not in a large urban county.

Mid-sized urban At least 50% of the tract population is in a mid-sized urban area, *and* the tract is in a mid-sized urban county.

Small urban At least 50% of the tract population is in a small urban area, *and* the tract is not classified as large or mid-sized urban.

Suburban If at least 50% of the tract population is in an urban suburban area of any size, the tract is classified as suburban (excluding those already classified as urban) regardless of county type. If at least 50% of the tract population is in a small town area, the tract is classified as suburban if in an urban or suburban county; otherwise it is classified as small town.

Small town At least 50% of the tract population lives in a small town of any size, *and* the tract is not classified as suburban or urban.

Rural More than 50% of the tract population lives in a rural or small town area of any size, *and* the tract is not classified as urban, suburban, or small town.

Around 500 tracts were not assigned an NCES locale and were given their county classification.

[20]See National Center for Education Statistics Locale Classifications (2023).

IV. Methodology

Due to the geographic variation in OZ designation, the simultaneous designation of tracts and finalization of regulations, and the characteristic-dependent selection of eligibility, we favor a difference-in-differences approach that is flexible enough to incorporate conditional parallel trends and produce dynamic effect estimates over the observed post-treatment period.

In particular, we use the doubly robust difference-in-differences estimator from Callaway and Sant'Anna (2021) (CSDID). We use a balanced sample for the difference-in-differences estimates of all 2010 census tracts with cross-walked treatment definitions that were either eligible low-income communities (LICs) or ineligible communities. We exclude tracts that were eligible via their contiguous status.

LIC tracts that were designated OZs are the treated group, while non-designated LIC tracts and ineligible non-LIC tracts form the control group. We exclude the roughly 2% of OZ tracts that were non-LIC (per the "contiguous tracts" special rule) from the primary analysis to better target the research question of interest: can a capital-gains tax incentive drive investment activity toward low-income and distressed places?

To help guard against a violation of the parallel trends assumption, we construct conditional parallel trends using each tract's poverty rate, median household income, the share of single-family housing, and a measure of local zoning rules within the tract in 2017.[21] Because OZs were enacted in late December 2017, we classify the first quarter of 2018 as the first treated period of the analysis.[22]

Notation. Let i index tracts and t year-quarters. Outcomes Y_{it} are: (i) residential *active and vacant* addresses excluding no-stat (business and "other" address types are excluded); (ii) log of that level; and (iii) a winsorized year-over-year growth rate. For (iii), we define the raw growth as

$$g_{it} \;=\; \frac{Y_{it} - Y_{i,t-4}}{Y_{i,t-4}},$$

and apply symmetric winsorization at pre-specified tails (held fixed across specifications).

We work with a balanced quarterly panel from 2015Q1 through 2025Q1 constructed from HUD–USPS administrative counts mapped to 2010 tract bound-

[21]We test for a violation of the parallel trends assumption using each of our conditioning variables and highlight the results in the appendix section *Impact of Conditional Trends on CSDID Estimates.*

[22]When best to define the beginning of treatment is somewhat unclear given the policy rollout of Opportunity Zones. While our analysis adopts the passage of the TCJA as the first point of treatment, it is also true that little effect could have occurred before the OZ map itself was certified. In an effort of caution—though it down-weights our average effect estimates over the post-treatment period—we adopt the TCJA treatment definition. This aligns us with the literature in general. Further, by defining treatment in the earlier sense, we limit the contamination of our conditional parallel trends, which may have been at risk if we allowed for post-TCJA and pre–map-certification quarters to enter our control period.

aries. Let $D_{it} = 1\{t \geq G_i\}$ denote treatment, where G_i is the tract's first treated quarter, taken to be 2018Q1 following the passage of TCJA.

The treated group comprises LIC tracts designated as OZs; the comparison pool comprises non-designated LIC tracts and ineligible non-LIC tracts.[23] Tracts designated under the "contiguous" rule (non-LIC) are excluded from the analysis. Let X_i denote tract covariates used for conditioning (poverty rate, median household income, solo-detached housing share, zoning index); ACS covariates are taken from 2012–2023 5-year files and carried forward for 2024–2025 when needed, while conditioning uses the baseline (pre-2018Q1) snapshot.

Define potential outcomes $Y_{it}(1)$ and $Y_{it}(0)$. The cohort–time average treatment effect on the treated (ATT) is

$$\mathrm{ATT}(g,t) \equiv E[Y_{it}(1) - Y_{it}(0) \,|\, G_i = g,\ t \geq g]\,.$$

We report aggregate post-treatment effects by averaging $\mathrm{ATT}(g,t)$ over (g,t) with observation-weighted schemes, and dynamic effects by event time $\ell = t - g$,

$$\mathrm{ATT}^{\mathrm{ES}}(\ell) \;=\; \sum_{g} \omega_{g,\ell}\, \mathrm{ATT}(g, g+\ell),$$

normalizing a pre-period (e.g., $\ell = -1$) to zero when plotting event studies.[24]

TWFE. For continuity with prior OZ work, we estimate the canonical two-way fixed-effects regression

$$Y_{it} \;=\; \beta^{\mathrm{TWFE}} D_{it} \;+\; X_{it}'\delta \;+\; \alpha_i \;+\; \gamma_t \;+\; \varepsilon_{it}, \tag{1}$$

with tract and period fixed effects and tract-clustered standard errors. Under both a staggered adoption scenario and time varying continuous control variables within X_{it}, β^{TWFE} can be impacted by negative weights, distorting the interpretation of the true effect estimate.(Goodman-Bacon, 2021; Sun and Abraham, 2021; De Chaisemartin and d'Haultfoeuille, 2024); we therefore treat the TWFE as a benchmark, and rely on alternative estimators of the DID for causal interpretation.

CSDID. We implement doubly robust $\mathrm{ATT}(g,t)$ estimates conditioning on baseline tract characteristics and both aggregate to an overall post-treatment effect and an event study by relative time (using a universal base period and normalizing the base pre-period to zero). Standard errors follow the `did` package's formulas.

Matrix completion. To relax additive FE assumptions, and allow for local effect estimates within each individual tract, we also estimate matrix completion counterfactuals for $Y_{it}(0)$ using pre-treatment and never-treated donor cells,

[23] When estimating the TWFE approach, we do not include ineligible tracts within the control group as they represent a poor set of tracts for comparison.

[24] This is done using the `did` option for a universal base period rather than a "varying" base period.

yielding tract-time effects $\widehat{\tau}_{it}$ and post-treatment averages. These $\widehat{\tau}_{it}$ are used in the later test on additionality.

Under an interactive fixed-effects specification for untreated outcomes,

$$Y_{it}(0) = \alpha_i + \gamma_t + \boldsymbol{\lambda}_i' \boldsymbol{f}_t + u_{it},$$

we impute $\widehat{Y}_{it}(0)$ via matrix completion on donor cells and compute $\widehat{\tau}_{it} = Y_{it} - \widehat{Y}_{it}(0)$ for $t \geq G_i$; we summarize with average post-treatment effects and map $\widehat{\tau}_{it}$ for heterogeneity.[25]

A. Estimator Integration

Our analysis adopts a layered strategy. We use a two-way fixed effect (TWFE) approach that provides continuity with the previous literature, even if flawed, to help ground our effect estimates. The CSDID supplies the main causal inference under well-documented assumptions, handling both the conditional parallel trends necessary for causal identification and documenting the time varying effect of OZ designation at a low computational cost. We also introduce a matrix completion (MC) estimator that probes the stability of results when we relax additive fixed-effect restrictions Xu (2017). The MC estimator also opens the door to more direct measures of treatment effect heterogeneity thanks to the tract-time, $\widehat{\tau}_{it}$, treatment estimates which are impossible in the TWFE and CSDID approaches.

Convergent estimates of the aggregate effect of designation across all three methods strengthen confidence that the observed effects are indeed attributable to the Opportunity Zone designation rather than to unmodelled tract-level shocks.

V. Results

A. Descriptive statistics

Prior to OZ designation, annual growth in residential addresses in OZ tracts lagged consistently behind growth in the rest of the country. Following designation in mid-2018, annual growth rates started converging with those in non-OZ communities. After OZ regulations are finalized at the end of 2019, we can see a sharp and sustained increase in growth as OZ tracts quickly close the gap. By 2023, OZ tracts have caught up with non-OZ tracts, and by Q1 2025, they have pulled ahead by a meaningful margin.

One plausible reason for this break in trend is that low-income census tracts overall may have seen an increase in demand, and the increase in housing supply in OZ tracts may simply reflect this broader shift. However, prior to the implementation of the OZ incentive, a steadily declining share of addresses was apparent

[25]Implementation of the matrix completion strategy uses teh r package `fect`, with the options selected for two-way effects, cross-validated rank, and bootstrap uncertainty.

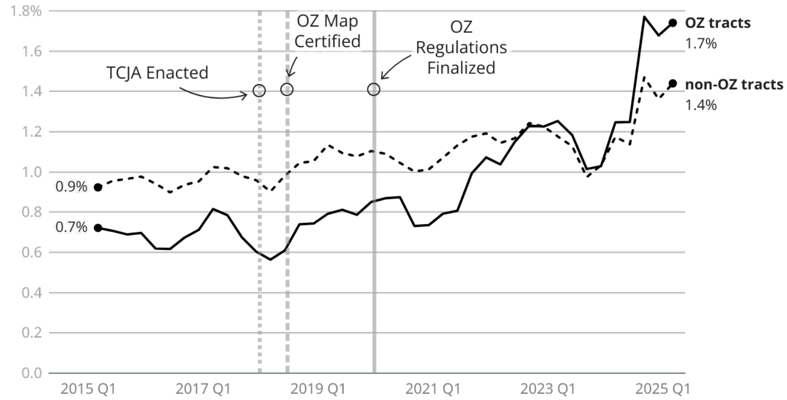

FIGURE 3. ANNUAL GROWTH RATE IN HOUSING UNITS, Q1 2015 TO Q1 2025

Note: In this figure, we are plotting the four quarter growth rate in the count of residential addresses, active and vacant, in designated LIC tracts and undesignated tracts.

across all low-income tracts, demonstrating years of low rates of residential investment. This is true of both OZ designated and non-designated low-income tracts. After implementation, however, OZ tracts have seen their share of all U.S. residential addresses grow for the first time in over a decade. Meanwhile, low-income communities that were not designated as OZs are continuing to shrink relative to the rest of the country.

While the aggregate data are consistent with the OZ incentive driving a substantial increase in residential housing investment, a more careful analysis is needed to establish causality. The key comparison is between LIC tracts that were or were not designated OZs. Our treatment group is a balanced panel of 7,580 designated LIC tracts observed quarterly from Q1 2014 to Q1 2025. The treated tracts in our sample are contrasted with the 23,057 undesignated LIC tracts, offering a comparable control group when using the two-way fixed effect estimator. When we use the CSDID or matrix completion method, both of which allow for better control unit selection on observable characteristics, we also include ineligible tracts in the control group, an additional 30,694 tracts.

Before the introduction of the OZ incentive, selected and unselected tracts within our sample had a similar stock of residential addresses, with the median values of active and vacant residential addresses of 1,534 and 1,510, respectively. Selected tracts were poorer on average, with a median family income of $35,095, as opposed to $40,542, and a poverty rate of 27.5 percent, 5.6 percentage points higher than unselected tracts. The unemployment rate was also higher among selected tracts, sitting 1.8 percentage points above the unselected unemployment rate of 8 percent.

Together, the descriptive statistics preceding the introduction of OZs paint a

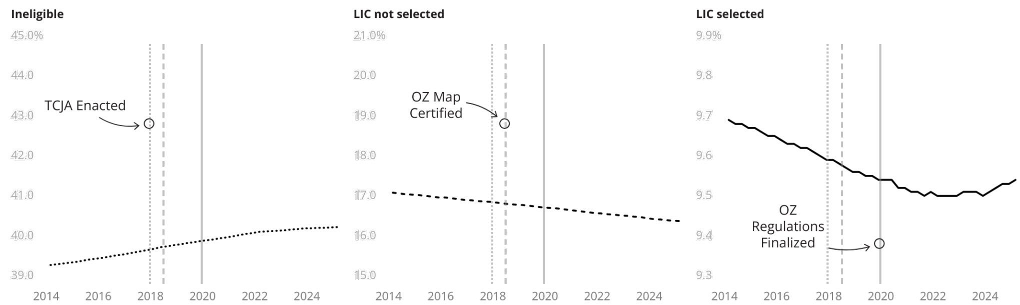

Figure 4. Change in the share of all active and vacant residential addresses by Opportunity Zone designation category

Note: In this figure, we are plotting the share of all active and vacant addresses within each of the three designation categories used in our treatment and control samples.

picture of places in severe economic need. Tracts that were designated as OZs were among the poorest in the country and looked likely to continue to experience economic segregation as investment activity shifted further away. Following designation, we can see descriptive evidence of a shift. The increase in both the share and the growth rate of active and vacant residential addresses, specifically among treated tracts, indicate a likely effect of the policy.

B. *Effect Estimates*

We evaluate OZ effects on local housing activity using three related outcomes built from USPS–HUD address counts: (i) the level (the count of active and vacant residential addresses), (ii) the log level (the natural log of that count), and (iii) the annual growth rate of the count (winsorized at the 1st percent and 99th percent).

Each outcome illuminates a different facet of the response. The level measures absolute changes in address units. The log-level tests for the comparability across tracts with different baseline sizes and reduces scale-driven variance. The annual growth rate captures the pace of change, highlighting timing and acceleration patterns around designation and construction pipelines.[26]

Taken together, levels speak to unit counts, logs to proportional impacts, and growth to dynamics. Because address counts net out one-for-one unit replacements, positive net growth should be read as a lower bound on new construction; consistent findings across all three outcomes strengthen the interpretation that observed changes reflect genuine movements in local housing supply.

[26]Winsorization in the growth rate prevents outliers or spikes from dominating estimates which rely on averages across treated tracts.

Table 1—Descriptive Statistics as of Q1 2017

Outcome	**LIC selected**	**LIC not selected**	**Contiguous selected**	**Contiguous not selected**	**Ineligible**
Active/Vacant	**1,534**	**1,510**	**1,784**	**1,712**	**1,797**
Active	1,439	1,442	1,738	1,667	1,767
Median					
MFI	35,095	40,542	52,832	55,000	77,297
Pov. Rate	27.5%	21.9%	13.1%	11.6%	6.6%
Prime-Age Share	38.8%	39.5%	38.3%	38.6%	39.2%
Unemp. Rate	9.8%	8.0%	6.4%	5.5%	4.5%
Share					
Large urban	28.4%	29.1%	16.2%	14.2%	15.3%
Mid-sized urban	6.2%	6.7%	0.6%	3.1%	2.3%
Small urban	7.5%	6.3%	3.0%	3.7%	3.6%
Suburban	27.1%	32.0%	25.8%	35.2%	58.2%
Small town	9.9%	6.1%	6.0%	4.3%	3.1%
Rural	20.9%	19.9%	48.5%	39.5%	17.5%
Number of Tracts	7,580	23,057	167	10,007	30,694

Note: Using 2020 standardized USPS data that are cross-walked to 2010 tract definitions. For the purposes of our analysis, we exclude contiguous tracts.

Table XX shows the average treatment effect across the post-treatment period for all three outcomes of interest and across our three estimators. It is clear that OZ designation produces consistent positive and significant effects on the supply of housing across our three outcomes and estimators of interest. Annual growth in the count of residential addresses rises by 0.1–0.2 percentage points. In levels, OZs are associated with roughly 15 to 20 additional new residential addresses per tract on average in the post-treatment period. In logs, effects are 0.9–1.2 percent. These effects are robust to the exclusion and inclusion of neighboring tracts, suggesting spillovers to neighbors, or the pull of projects from neighboring tracts, both LIC and non-LIC, are not driving the results.

These estimates, however, are averaged over the entire post-treatment window, which is not ideal when time is part of the treatment "dosage." As noted by Fikri and Glasner (2023), OZ impacts are likely to ramp up only after key regulatory milestones, capital raising and deployment, and the inherent time-to-build (often 18–24 months for multifamily).

It is reasonable to anticipate that early post periods contain little effect, mechanically diluting later, larger impacts in the average ATT. In such settings, event-time estimates are better suited to assess the policy's trajectory.

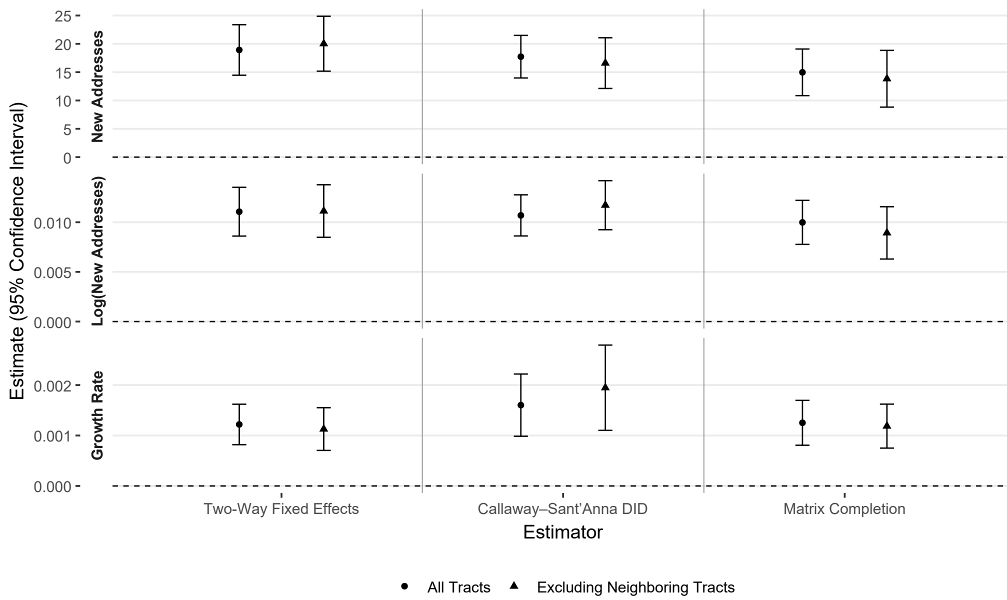

FIGURE 5. ESTIMATED AVERAGE TREATMENT EFFECTS IN THE POST-TREATMENT PERIOD

Note: Points show estimates; bars show 95% confidence intervals.

We also note that the CSDID effect estimate on the count of active and vacant residential addresses may indicate that some neighboring tracts are seeing a small reduction in address growth, resulting in local effect estimates that are positively inflated by a negative spillover, comparing the 17.7 increase in address count to the 16.6 after excluding neighboring tracts. Because of this, we opt to exclude neighboring tracts when presenting the results that follow.

We can see that across all three of our estimators, OZs resulted in a similar positive effect on residential addresses. Due to the doubly-robust nature of the CSDID estimator, capacity for propensity score matching, computational load, and integrated dynamic effect estimates, we favor these estimates for aggregate effect calculations.

Leveraging the CSDID methodology, we explore the heterogeneity of the effect across geographic typology. We also validate the parallel trends assumption using a universal base period event study design. Results for the average effect across the post-treatment period and the effect in the final period of the analysis are presented in Tables XX and XX.

We find that OZ designation results in a statistically significant increase in the number of active and vacant residential addresses in aggregate. The average

Outcome Variable	Active and Vacant	Address Growth Rate	log(Active and Vacant Residential)
All	**16.728 *** **	**0.2007 *** **	**0.012 *** **
	(2.246)	(0.0431)	(0.001)
Large urban	29.54 ***	0.1478	0.021 ***
	(4.319)	(0.0749)	(0.002)
Mid-sized urban	26.651 *	0.7252	0.014
	(11.12)	(0.3843)	(0.007)
Small urban	29.224 ***	0.3123	0.027 ***
	(9.984)	(0.3339)	(0.01)
Suburban	10.306	0.1092	0.008
	(5.072)	(0.0663)	(0.004)
Small town	-2.007	-0.0147	0.003
	(6.656)	(0.1364)	(0.003)
Rural	9.776	0.2647 *	0.004
	(4.631)	(0.0864)	(0.003)

Table 2—Average effect estimates over the full treated period

Note: Regression results from the Callaway and Sant'Anna difference-in-differences (CSDID) approach. Using 2020 standardized Census tracts cross-walked to 2010 definitions. Results include conditional parallel trends accounting for tract poverty rate, median household income, share of housing classified as solo-detached, and an index of local zoning. Control groups exclude bordering census tracts to avoid contamination of the effect estimate. Reported effects for the count of active and vacant addresses are average treatment effects over the full post-treatment period. *Significance*: * 95%; ** 98%; *** 99%.

effect size in the final period of our sample is 47.5 new residential addresses per OZ tract.[27] This average masks significant variation in the effect size across geographic typology.

Across geographic typologies, the pattern is unambiguously urban-led. Large, mid-sized, and small urban tracts account for the bulk of the aggregate effect in both the post-period average and the final-period snapshot. In levels, the post-period averages are 29.5 additional residential addresses per treated tract in large urban areas, 26.7 in mid-sized urban areas, and 29.2 in small urban areas, each lifting the all-tract average to 16.7. By the end of the sample, these differences widen: the final-period effects reach 78.7 (large), 73.0 (mid-sized), and 88.7 (small). It should be unsurprising that more urban tracts have an advantage on the level effect given their greater concentration of people.

Growth-rate and log-level specifications corroborate the level results while also opening the door to the suburban effect and the mid-sized urban acceleration. Averaged over the post-period, the address growth rate effect is modest in large urban tracts (0.15) but substantially higher in mid-sized urban tracts (0.73), small urban tracts (0.31), and even rural tracts (0.26). In the final period, mid-sized urban tracts reach an annual growth rate effect estimate of 1.34, more than

[27]A single unit increase in the level of active and vacant addresses in our effect estimate translates to a single new address present in a tract that is either actively occupied or vacant, either short-term or long-term. Units that are under construction or that are determined to be "no-stat" by USPS are not included in the analysis's outcome measures.

Outcome Variable	Active and Vacant	Address Growth Rate	log(Active and Vacant Residential)
All	**47.5 *** **	**0.4802 *** **	**0.031 *** **
	(4.854)	(0.065)	(0.003)
Large urban	78.701 ***	0.6218 ***	0.052 ***
	(9.746)	(0.1326)	(0.006)
Mid-sized urban	73.024 ***	1.3442 ***	0.035 ***
	(19.656)	(0.4626)	(0.01)
Small urban	88.664	0.4047	0.084
	(41.822)	(0.5559)	(0.044)
Suburban	36.206 **	0.2725	0.022 **
	(11.655)	(0.1267)	(0.007)
Small town	6.35	0.3035	0.013
	(12.884)	(0.1664)	(0.007)
Rural	23.45	0.3421	0.009
	(10.078)	(0.1588)	(0.005)

Table 3—Average effect estimates in the final period of analysis

Note: Regression results from the Callaway and Sant'Anna difference-in-differences (CSDID) approach. Using 2020 standardized Census tracts cross-walked to 2010 definitions. Results include conditional parallel trends for tract poverty rate, median household income, solo-detached share, and a local zoning index. Control groups exclude bordering census tracts. Reported effects show the event-study estimate in the final observed period, to avoid down-weighting due to continued treatment dosage. *Significance*: * 95%; ** 98%; *** 99%.

doubling the effect on large urban tracts (0.62). The log-level specification tells a similar story. Taken together, these patterns suggest that mid-sized urban places combine sizable late-post level gains with faster proportional growth, whereas large urban places accumulate larger absolute additions with steadier proportional changes.

Small urban and suburban tracts also contribute positively but with distinct temporal profiles. Small urban tracts post a sizable positive post-period average in levels (29.2) and a clearly positive average log effect, indicating sustained accumulation across the window; however, their final-period point estimates are noisier in both levels and growth rates. Looking at the event study across all three outcomes in Figure XX helps illuminate this dynamic showing steady positive effects growing across the post treatment period, even as the geographic typology sample is noisy. Suburban tracts look comparatively smaller, with a lower average effect over the post-period in levels and growth, but the final-period level and log effects become statistically meaningful, consistent with a longer lag between designation and observable address additions, particularly in a less concentrated population center.

Small town and rural tracts trail the urban categories in both magnitude and precision. Point estimates in the final period are positive in levels and logs for both, but they remain imprecisely estimated. The one clear exception is the rural growth-rate average over the full post-period, which is positive and statistically significant. This combination, statistically detectable average growth through the post-treatment window but an imprecise late-post level effect, implies that rural tracts experienced modest, broadly distributed, and sporadic increases. One possible explanation for this is a less consistent investment profile among designated LIC rural tracts. As OZ designation does not guarantee investment activity, it is possible the event study is capturing the inconsistent nature of rural investment relative to the thicker markets in more urban tracts.

Our findings comport with previous research examining building permits across a sample of larger cities (Wheeler (2022)). That study found that OZ designation caused a "large and immediate" effect on the likelihood of development in a given tract and significant spillovers into non-OZ tracts nearby, but did not estimate the number of new addresses created. Survey data from Novogradac[28] of a subset of OZ investment reveals nearly 200,000 housing units built or scheduled to be built (as opposed to the completed, net new addresses in our findings) with OZ investment, but the nature of this survey data cannot establish causality. Similarly, private data from RealPage[29] found that OZ communities have more than doubled their national share of market rate multifamily housing since the policy was enacted, but did not establish causation or analyze fiscal cost. An analysis of Real Capital Analytics (RCA) apartment construction data[30] also shows OZ tracts going from lagging ineligible tracts in the pre-treatment period to outpacing them, and pulling farther ahead from tracts that were eligible but not selected.

One key point of contention in the literature is about the timing of when OZ impacts on neighborhood outcomes would begin to materialize empirically. Using CSDID's dynamic estimates, we can assess when effects begin, the rate of change in the post-treatment period, and the validity of the control group using the pre-treatment period.

The event study structure of the effect estimates makes two things clear.

First, our treatment and control groups appear to be an appropriate fit. The pre-treatment period, defined as the quarters preceding when TCJA was passed, exhibits no clear violation of the parallel trends assumption.[31] This falls in line with the aggregate descriptive figure on annual growth rates across OZ and non-OZ tracts.

Second, the growth in residential addresses across all tract types appears to grow after regulations were finalized. An estimate of the aggregate treatment effect across all periods underestimates the effect of OZs, downweighted by early periods when construction is likely still underway. As such, we are confident in the reporting of the final period of the analysis as our primary effect estimates.

Taking these results, we can estimate an aggregate effect on the number of

[28]See Watkins (2025).
[29]See Parsons (2024).
[30]See Costello (2025).
[31]We found that the parallel trends assumption held both with and without the inclusion of conditional parallel trends for both the logged count of active and vacant addresses and the growth rate of addresses, but the unconditional estimate failed to meet the parallel trends assumption for the raw count of active and vacant addresses. As such, we present effect estimates using the conditional parallel trends.

active and vacant residential addresses caused by OZs.

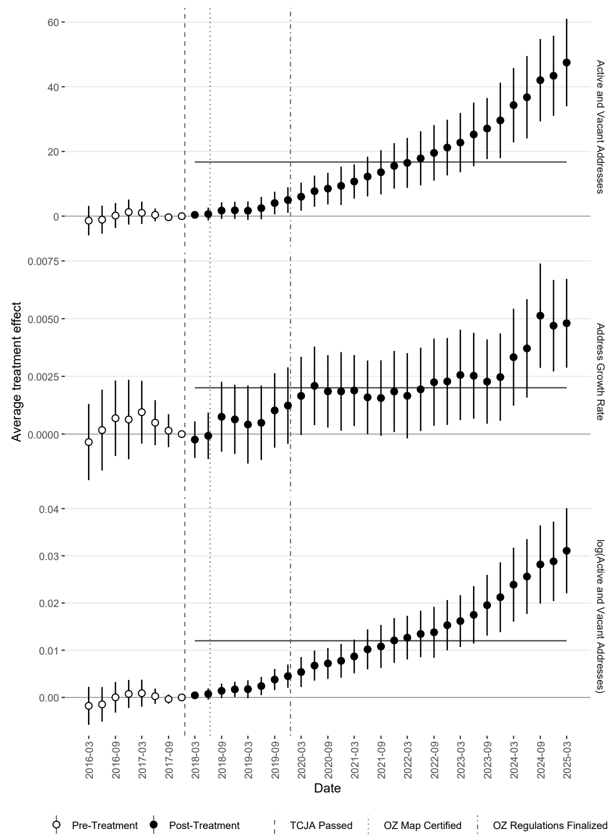

Figure 6. Event Study results by outcome across all treated tracts in sample

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.

Using the CSDID average effect of 47.5 additional active or vacant addresses per OZ tract and an analytic sample of 7,580 LIC OZ tracts, we estimate 360,048 addresses attributable to OZs since the passage of TCJA. That implies 41.1 percent of all new active or vacant addresses added in OZ tracts, 13.4 percent among LIC tracts, both designated and non-designated, and 3.8 percent of national net additions within our sample over this period were caused by the OZ incentive.

From the finalization of OZ regulations to Q1 2025, OZ tracts within our sample added 875,528 net new active or vacant addresses, or 32.7 percent of the 2,679,295 added in LIC-eligible tracts regardless of designation. Across the full analysis sample, which represents XX percent of the census tracts in the country, OZ tracts account for 9.3 percent of the 9,434,830 new addresses. Absent the OZ incentive, designated tracts would have accounted for only 5.5 percent. For comparison, in the pre-policy period (Q3 2014–Q3 2019), OZ tracts represented 6.5 percent of 5,718,562 net new active or vacant addresses across our sample. Taken together, these figures demonstrate that the OZ incentive reversed the trajectory of designated areas in terms of national housing growth, resulting in an increase in national share rather than a continued decline.

A longstanding critique of place-based policies is that they may simply displace or reallocate activity that would have happened outside of targeted geographies, leading to little or no meaningful changes on net. We find, however, that the housing gains in OZ tracts are largely additional on net to their wider areas rather than an indication of displaced investment and construction activity. In an inclusive total that adds statistically significant neighbor-tract effects, the treated total falls, at its minimum, to 96.6 percent of the unconditioned OZ effect at a distance radius of roughly 5 km, implying limited local reallocation. Moreover, the small shortfall at 5 km is almost entirely offset by net growth in new addresses at larger radii. This estimate is discussed in more detail in the appendix.

We can go further. Because our estimate represents a restrictive sample of OZs, focusing solely on a balanced panel of tracts, and prioritizing the effect estimate from LIC OZs, our aggregated effect estimate of 360,048 new addresses is an undercount. If we assume a consistent treatment effect across all 8,764 OZ tracts across states, territories, and the District of Columbia, we can estimate that the OZ incentive resulted in 416,290 new residential addresses.[32]

C. *Other considerations about aggregate effect size*

There are a number of reasons our results may understate the scale of the OZ incentive's effect on housing development. We exclude units that were under construction as of Q1 2025, for example, as well as completed developments that do not yet show up as USPS addresses. Because the effect is rising as of the end of our study period, there is strong reason to expect that even larger cumulative effects will be observed in subsequent quarters as in-progress developments are certified for occupancy. Moreover, as we are only modeling the net change in residential addresses, our findings exclude units that have been substantially rehabilitated as a result of the OZ incentive, or dilapidated housing that was demolished and replaced. Lastly, there is evidence to suggest that OZ designation of "contiguous tracts" — ones that were not LICs but nevertheless eligible for designation under a special rule — caused a larger increase in residential addresses than what we observe on average in designated LIC tracts.[33]

[32]This analysis leverages HUD's data on addresses within states and the District of Columbia and does not include addresses in territories. If we exclude OZs that are not in states or the District of Columbia, we are left with 7,826 designated tracts. This exclusion yields an alternative estimate of 371,735 new addresses caused by OZs. However, this estimate ignores any potential effect on addresses within territories and is likely to underestimate the total effect.

[33]It is likely that contiguous, non-LIC, designated OZ tracts saw even greater increases in net new active and vacant addresses. Been et al. (2025) found that "In New York City, of the 306 total tracts

VI. Conclusion

There is broad consensus on the need to expand U.S. housing supply. A large body of evidence demonstrates that increasing local housing supply lowers rents or slows rent growth (Been, Ellen and O'Regan, 2025). This applies across the full spectrum of new housing—not only subsidized units for lower-income residents (Lettieri, 2021). New market-rate construction, for instance, tends to loosen overall housing markets, benefiting middle- and lower-income households as well (Bratu, Harjunen and Saarimaa, 2023). By easing pressure on the housing stock, additional supply can reduce eviction rates and mitigate displacement risk associated with neighborhood revitalization (Dawkins, 2024).

Our study directly addresses two central questions about the housing effects of Opportunity Zones (OZs) that are broadly relevant to any tax-incentive-based development policy: (1) did the policy induce behavioral or economic changes that would not have occurred "but for" the incentive, and (2) did those changes occur at an economically meaningful scale? The evidence indicates yes on both counts.

We find that OZs are a significant housing-supply policy, catalyzing new residential investment that would not otherwise have materialized in targeted areas. Using HUD data derived from USPS address counts, we estimate the causal effect of OZ designation on the total number of active and vacant residential addresses and examine variation across geographies. The results reveal a large — and still growing — increase in housing supply within OZ communities. Even as of Q1 2025, the effect size remains on an upward trajectory, suggesting that the cumulative impact will continue to expand as more projects reach completion.

We observe a distinct pre- versus post-designation shift in housing growth, addressing the concern that OZs merely subsidize activity that would have occurred regardless of the incentive. Prior to designation, low-income tracts that became OZs were shrinking as a share of the national housing stock — a clear marker of relative economic decline. Following designation, their trajectory reversed: housing growth in OZ tracts has since outpaced that of non-OZ communities nationwide.

These findings highlight an essential consideration for interpreting the broader OZ literature: causal effects on real-asset formation, such as housing, unfold gradually in line with standard development timelines. Studies limited to early post-designation periods likely underestimate the magnitude of the policy's effects, or perhaps even miss those effects entirely.

More broadly, the results speak to the efficacy of the OZ model's distinctive design. Unlike earlier place-based programs, OZs rely on decentralized, market-driven capital allocation rather than top-down project selection. The substantial

that the State designated as OZ, 14 were contiguous but not low-income tracts. In those 14 tracts, rates of completed units reached 16 percent after the program was operational, compared to 6 percent in other OZs."

increase in housing supply documented here suggests that this approach can successfully mobilize private capital for local revitalization at scale.

Our study fills a critical gap in the empirical literature on both Opportunity Zones and place-based tax incentives more generally by quantifying their impact on housing development. Beyond advancing understanding of OZs themselves, these results contribute to ongoing debates about how tax policy, housing markets, and private investment can be harnessed to improve local outcomes.

REFERENCES

Arefeva, Alina, Morris A. Davis, Andra C. Ghent, and Minseon Park. 2024. "The Effect of Capital Gains Taxes on Business Creation and Employment: The Case of Opportunity Zones." *Management Science*, mnsc.2022.03223.

Atkins, Rachel M. B., Pablo Hernández-Lagos, Cristian Jara-Figueroa, and Robert Seamans. 2023. "JUE Insight: What is the impact of opportunity zones on job postings?" *Journal of Urban Economics*, 136: 103545.

Bartik, Alexander, Arpit Gupta, and Daniel Milo. 2024. "The costs of housing regulation: Evidence from generative regulatory measurement." *Available at SSRN 4627587.*

Been, Vicki, Hayley Raetz, Brad Greenburg, and Matthew Murphy. 2025. "How Federal Programs Shape Housing Policy in New York City." The Stoop*, NYU Furman Center Blog.*

Been, Vicki, Ingrid Gould Ellen, and Katherine O'Regan. 2025. "Supply Skepticism Revisited." *Housing Policy Debate*, 35(1): 96–113.

Bernstein, Jared, and Kevin A. Hassett. 2015. "Unlocking private capital to facilitate economic growth in distressed areas." *Economic Innovation Group.*

Bratu, Cristina, Oskari Harjunen, and Tuukka Saarimaa. 2023. "JUE Insight: City-wide effects of new housing supply: Evidence from moving chains." *Journal of Urban Economics*, 133: 103528.

Callaway, Brantly, and Pedro HC Sant'Anna. 2021. "Difference-in-differences with multiple time periods." *Journal of econometrics*, 225(2): 200–230. Publisher: Elsevier.

Chen, Jiafeng, Edward Glaeser, and David Wessel. 2023. "JUE Insight: The (non-) effect of opportunity zones on housing prices." *Journal of Urban Economics*, 133: 103451. Publisher: Elsevier.

Collinson, Robert, Ingrid Gould Ellen, and Jens Ludwig. 2015. "Low-income housing policy." In *Economics of Means-Tested Transfer Programs in the United States, Volume 2.* 59–126. University of Chicago Press.

Corinth, Kevin, and Naomi Feldman. 2024. "Are Opportunity Zones an Effective Place-Based Policy?" *Journal of Economic Perspectives*, 38(3): 113–136.

Corinth, Kevin, David Coyne, Naomi E. Feldman, and Craig Johnson. 2025. "The Targeting of Place-Based Policies: The New Markets Tax Credit Versus Opportunity Zones." National Bureau of Economic Research.

Costello, Jim. 2025. "Opportunity Zone Legislation Helped Spur New Investment." *LinkedIn post*, Head of Real Estate Economics, MSCI; Chief Economist, MSCI Real Assets.

Coyne, David, and Craig Johnson. 2023. *Use of the Opportunity Zone Tax Incentive: What the Data Tell Us.* US Department of the Treasury, Office of Tax Analysis.

Dawkins, Casey J. 2024. "Land Use Regulations, Housing Supply, and County Eviction Filings." *Journal of Planning Education and Research*, 44(3): 1719–1729. Publisher: SAGE Publications Inc.

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2024. "Difference-in-differences estimators of intertemporal treatment effects." *Review of Economics and Statistics*, 1–45.

Feldman, Naomi, and Kevin Corinth. 2023. "The Impact of Opportunity Zones on Commercial Investment and Economic Activity."

Fikri, K., and B. Glasner. 2023. "Are opportunity zones working? What the literature tells us." *Unpublished working paper. Economic Innovation Group.*

Fikri, Kenan, and John Lettieri. 2018. "The State of Socioeconomic Need and Community Change in Opportunity Zones." *Economic Innovation Group.*

Fikri, Kenan, August Benzow, and John Lettieri. 2023. "Examining the Latest Multi-Year Evidence on the Scale and Effects of Opportunity Zones Investment."

Frank, Mary Margaret, Jeffrey L. Hoopes, and Rebecca Lester. 2022. "What determines where opportunity knocks? Political affiliation in the selection of opportunity zones." *Journal of Public Economics*, 206: 104588. Publisher: Elsevier.

Freedman, Matthew, Noah Kouchekinia, and David Neumark. 2025. "A Longer-Run Evaluation of the Employment Effects of Opportunity Zones."

Freedman, Matthew, Shantanu Khanna, and David Neumark. 2023. "Jue insight: The impacts of opportunity zones on zone residents." *Journal of Urban Economics*, 133: 103407. Publisher: Elsevier.

Glaeser, Edward L., and Joshua D. Gottlieb. 2008. "The economics of place-making policies." National Bureau of Economic Research.

Goodman-Bacon, Andrew. 2021. "Difference-in-differences with variation in treatment timing." *Journal of econometrics*, 225(2): 254–277.

Gyourko, Joseph, Jonathan S. Hartley, and Jacob Krimmel. 2021. "The local residential land use regulatory environment across US housing markets: Evidence from a new Wharton index." *Journal of Urban Economics*, 124: 103337. Publisher: Elsevier.

Internal Revenue Service. 2024. "Opportunity zones frequently asked questions | Internal Revenue Service."

Joint Committee on Taxation. 2024. "Estimates of federal tax expenditures for fiscal years 2024–2028, JCX-48-24, December 11." Joint Committee on Taxation.

Kennedy, Patrick, and Harrison Wheeler. 2021. "Neighborhood-level investment from the US opportunity zone program: Early evidence." Working Paper 4024514.

Lettieri, John. 2021. "Moving Beyond Flawed Critiques Of the O-Zone Incentives." *Tax Notes Federal*, 171(13).

National Center for Education Statistics Locale Classifications. 2023. "National Center for Education Statistics Locale Classifications."

Neumark, David, and Helen Simpson. 2015. "Place-based policies." In *Handbook of regional and urban economics*. Vol. 5, 1197–1287. Elsevier.

Parsons, Jay. 2024. "Opportunity Zones are working as planned. The number of apartment units built in an OZ jumped 3x between 2016 and 2022, and then DOUBLED in 2023. In 2023, 20% of all new apartments in the U.S. were located in an OZ. . . ." *Tweet*, Posted 2:57 PM. Twitter/X post.

Sage, Alan, Mike Langen, and Alex Van De Minne. 2023. "Where is the opportunity in opportunity zones?" *Real Estate Economics*, 51(2): 338–371.

Sciarretti, John. 2023. "Novogradac-Tracked QOFs Show Sharp Decline in Investment in First Quarter."

Sun, Liyang, and Sarah Abraham. 2021. "Estimating dynamic treatment effects in event studies with heterogeneous treatment effects." *Journal of econometrics*, 225(2): 175–199.

Watkins, Jason. 2025. "Nearly 200,000 Homes Financed by QOFs Tracked by Novogradac."

Wheeler, Harrison. 2022. "Locally optimal place-based policies: Evidence from opportunity zones." *Unpublished working paper. https://hbwheeler. github. io/files/JMP_HW. pdf.*

Xu, Yiqing. 2017. "Generalized synthetic control method: Causal inference with interactive fixed effects models." *Political Analysis*, 25(1): 57–76.

Appendix

A1. Impact of conditional trends on CSDID estimates

This section reports how alternative sets of tract-level covariates used to construct conditional parallel trends affect the estimated impact of Opportunity Zone designation on residential address outcomes. We re-estimate the Callaway and Sant'Anna (2021) difference-in-differences (CSDID) model across four control specifications: (i) no controls; (ii) poverty rate and median household income; (iii) poverty, income, and the share of single-family ("solo detached") housing; and (iv) poverty, income, single-family share, and a tract-level zoning index. For each specification we report (a) the overall post-treatment average treatment effect on the treated (ATT) and (b) the effect in the final observed period. Estimates are shown for the full sample and by geographic typology. Outcomes mirror the main text: the level of active + vacant residential addresses, the log level, and the annual growth rate (winsorized at the 1st/99th percentiles).

Without controls, overall effects for the full sample are negative in levels (-24.011) and logs (-0.003) despite a small positive effect on annual growth (0.0021). Allowing the estimator to condition on poverty and income flips the level (18.352) and log (0.013) effects positive and significant. Adding the single-family share and then the zoning index leaves the overall results stable; growth remains positive and similar across all conditioned sets ($\approx$ 0.0019–0.0020). This pattern indicates that unconditioned comparisons are biased and that modest conditioning removes that bias without attenuating effects. Figure A1 shows the dynamic paths behind these averages. The unconditioned specification exhibits clear pre-treatment deviations from zero, with systematic trends in the pre-period that violate the parallel trends assumption. By contrast, specifications that condition on tract characteristics display flat and small pre-period estimates within simultaneous bands, supporting their validity. Post-treatment, conditioned series rise steadily and align closely across control sets, indicating robustness to reasonable choices of conditioning.

Final-period estimates show the same stability with larger magnitudes. For the full sample, the preferred specification with all four controls yields a level effect of 47.500, a log effect of 0.031, and annual growth of 0.0048. Using poverty+income only gives 51.037 in levels and 0.032 in logs; adding single-family share and zoning slightly reduces magnitudes while preserving significance. Geographic heterogeneity in the preferred final-period estimates matches the main text. An assessment of parallel trends is central in difference-in-differences. Here, the no-controls event studies fail that test; the conditioned designs pass it and produce consistent estimates. We therefore adopt the all-controls specification as preferred and report the full sensitivity set for transparency.

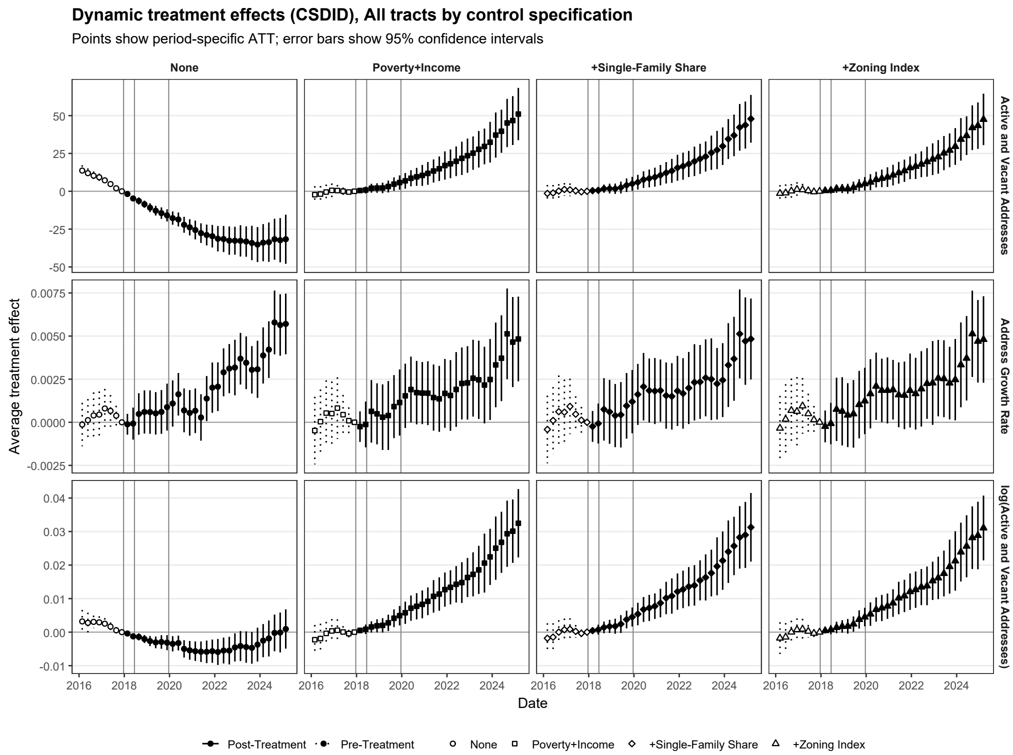

Figure A1. Event Study results by outcome and control groups

Note: Points show ATT; error bars show 95% confidence intervals.

A2. Geographic impact and additionality

The core result presented in the main findings is local: OZ designation raises housing supply inside designated tracts. The next question is about additionality. Did this investment activity within designated tracts come from genuinely new investment, or did some of these investments shift across tract borders, implying a smaller net address effect?

We address this in two parts. First, we check that results are not an artifact of estimator choice or sample definition. We do this by re-estimating the TWFE, CSDID, and MC approaches on both the broad sample and the "no-neighbor" sample used in the main text. Effects are positive and of similar order across designs and samples, which reduces concern that border tracts alone drive the findings.

Second, we assess the additionality concern directly, that OZ investment merely reshuffles activity from nearby tracts. We do this by comparing treated effects to "inclusive" totals that add up statistically significant spillovers within growing radii around OZ boundaries via an imputation approach. MC estimates provide tract-specific counterfactual paths. We take final-period tract effects from within the MC estimate, group tracts by 0.2-km distance bins to the nearest OZ boundary, and stack effects by tract status (designated OZ; LIC non-designated—border and non-border; ineligible—border and non-border). The figure shows the cumulative curve of each of these binned effect estimates. A displacement story would show sizable negative effects just outside OZs that offset treated gains.

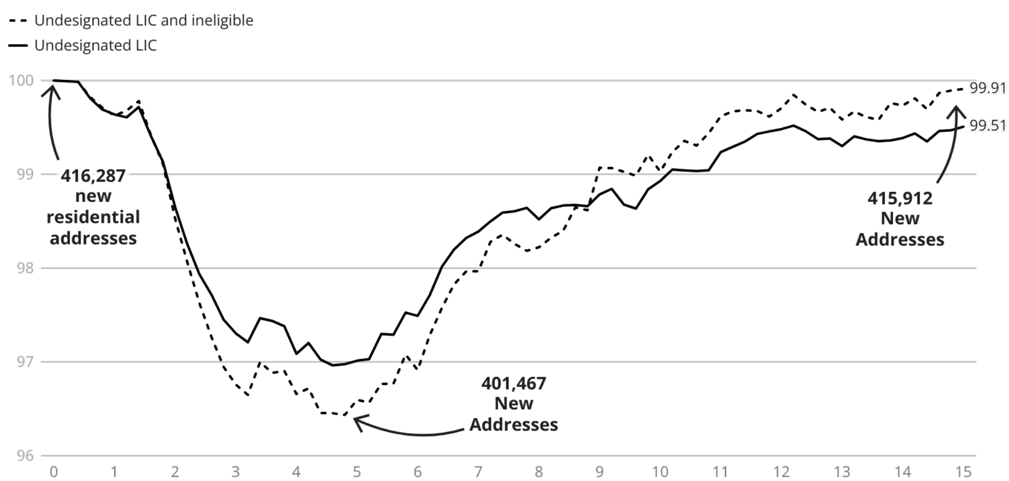

Figure A2. Cumulative net effect by distance from Opportunity Zone boundary for statistically significant final period tract effects

Note: Inclusive cumulative totals within radius r, in 0.2-km bins to 15 km. Indexed to treated total = 100 among only treated tracts. The treated baseline (100 at 0 km) is the sum of tract-level matrix-completion (MC) effects for designated LIC tracts in the final period. For each radius r, the numerator adds up significant MC effects ($p < 0.05$) for neighbors within r (binned to the nearest 0.2 km) and includes the treated baseline. Thus, points below 100 indicate net crowd-out within r; points near 100 indicate limited local spillovers; points above 100 indicate positive spillovers to neighbors.

Using the significant-only universe for MC effects, we can create an estimate of the risk to additionality and contrast it to our preferred baseline treated total implied by CSDID, 416,287 addresses across 8,764 OZ tracts. Using the MC effect estimates, we cumulate statistically significant, at an α of 0.05, effects from treated tracts and neighbors and index the running total to the treated baseline.

The pattern shown in Figure A2 is not consistent with meaningful displacement. Within 2.0 km, the inclusive total reaches 98.5 percent of the treated baseline (scaled total 410,158). At the inclusive 5.0 km point, where most local reallocation concerns should be strongest from this visual trend, the inclusive total is 96.6 percent (402,113). By 15.0 km, the inclusive total is 99.9 percent (415,914), effectively one-for-one with treated-only gains. These magnitudes imply that the vast majority of housing supply growth inside OZs is net rather than borrowed from nearby low-income tracts.

This evidence aligns with the main design choice to exclude neighboring tracts when identifying the treated effect, yet it also offers a conservative "all-in" check: when we explicitly add back significant neighbor effects within wide radii, the aggregate conclusion barely moves. The small gap at 5 km suggests limited crowd-out close to OZ borders, but it closes by 15 km. In short, the data do not support the conclusion that OZs raise construction inside zones primarily by draining activity from adjacent places. Instead, the results indicate substantial net additions to area housing supply. This result aligns with the findings from **?**.

A3. Supporting Figures

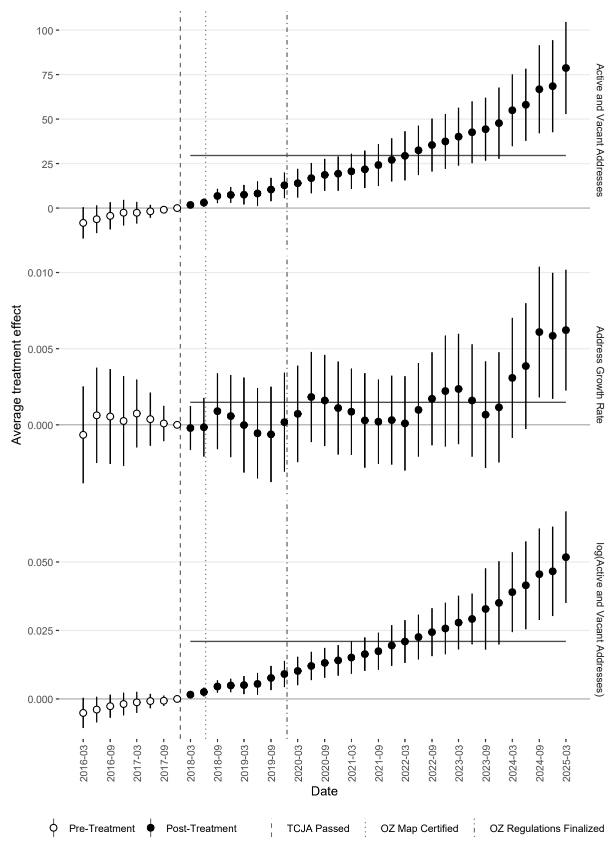

FIGURE A3. EVENT STUDY RESULTS BY OUTCOME ACROSS LARGE URBAN TREATED TRACTS IN SAMPLE

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.

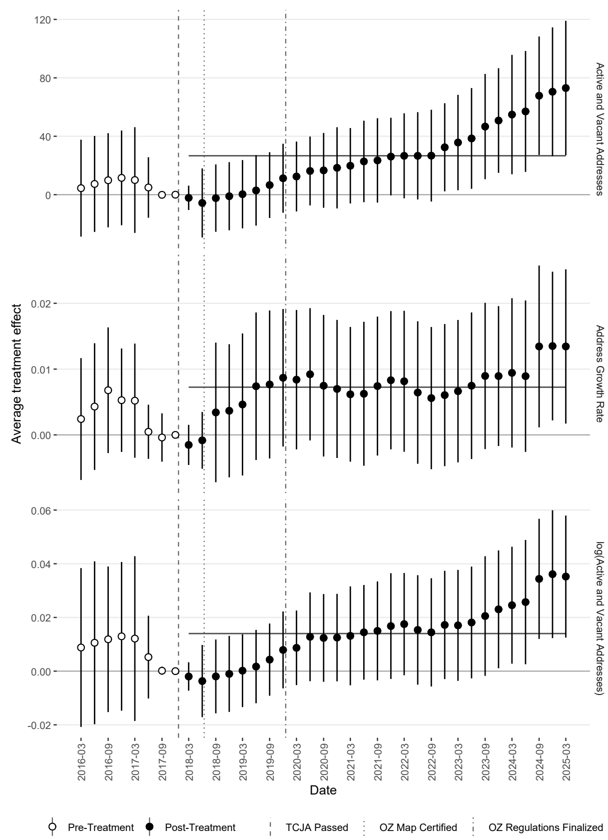

Figure A4. Event Study results by outcome across medium urban treated tracts in sample

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.

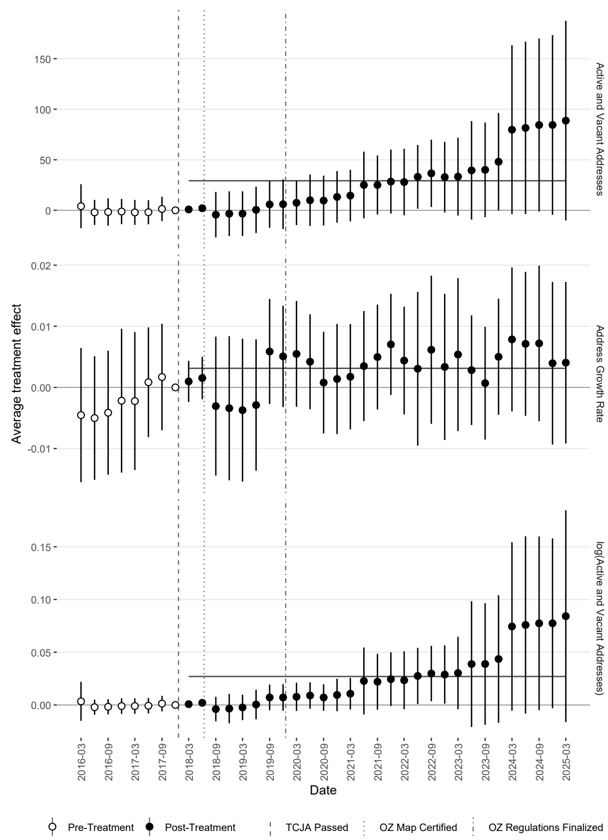

Figure A5. Event Study results by outcome across small urban treated tracts in sample

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.

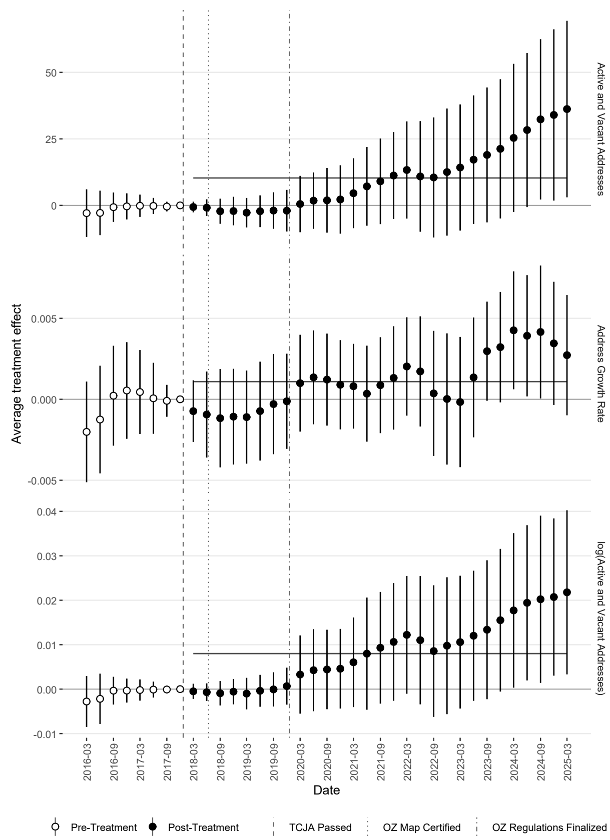

Figure A6. Event Study results by outcome across suburban treated tracts in sample

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.

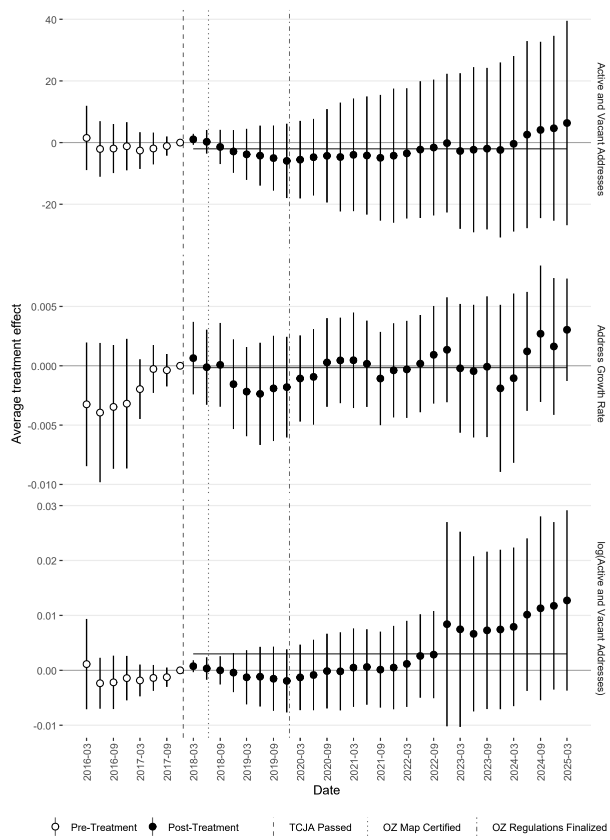

Figure A7. Event Study results by outcome across small town treated tracts in sample

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.

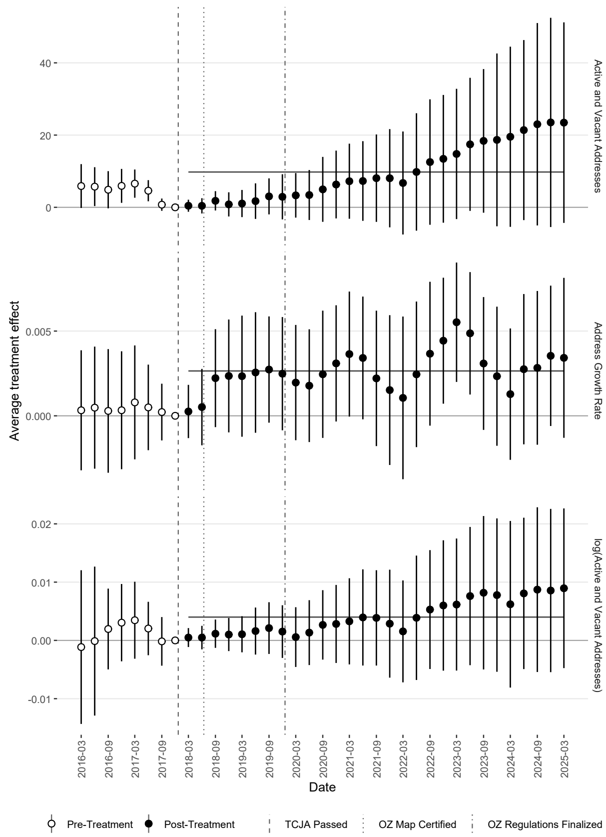

Figure A8. Event Study results by outcome across rural treated tracts in sample

Note: Points show ATT; error bars show 95% confidence intervals; solid line is post-period average ATT.