# Local Spillover Effects of Density Bonus Policy on Housing Production: Evidence from Massachusetts *

Aja Kennedy

January 21, 2026

Click here for latest version.

**Abstract**

Many metropolitan areas in the United States suffer from a housing crisis characterized by housing prices that are unaffordable to low- or moderate-income residents. One way policymakers address this problem is to increase supply by incentivizing housing production, especially production of income-restricted housing. A "density bonus" policy is one such intervention. It incentivizes production of dense, mixed-income housing developments by relaxing restrictive local zoning regulations and providing a streamlined permitting process for developers who include income-restricted units in their proposed residential developments. Massachusetts Chapter 40B, the country's oldest density-bonus policy, relaxes zoning regulations through a state-level override of local restrictions on density. Using Chapter 40B as a case study, I analyze local crowd out effects of mixed-income housing developments in high-income neighborhoods by estimating spillover effects of these developments on local housing construction using a staggered differences-in-differences technique to compare areas with approved 40B developments to areas where approval has not yet occurred. I find that, within 0.8 miles of a Chapter 40B development, approval reduces local single-family construction by about 5 percent over the course of 10 years following approval. Comparing the average number of units in a development with the average negative spillover

*This work benefitted from feedback from the Brownbag series at The Federal Reserve Bank of Boston, the 2025 Association of Public Policy Analysis and Management Conference, 2025 American Economic Association Summer Pipeline Conference, and seminars with the Neubauer Family Economics and Public Policy program at Tufts University.This work has been supported (in part) by Grant G-2401-46906 from the Russell Sage Foundation and a dissertation fellowship from the Federal Reserve Bank of Boston. Any opinions expressed are those of the principal investigator alone and should not be construed as representing the opinions of the Russell Sage Foundation, the Federal Reserve Bank of Boston, or the Federal Reserve system. Any errors or omissions are the sole responsibility of the author.

effect, I estimate that incentivizing one additional 40B unit adds about 0.95 of a unit to the local housing stock, meaning many 40B units are not locally inframarginal.

## 1 Introduction

Many major U.S. cities are experiencing a persistent housing affordability crisis which generates a host of economic and social problems including displacement of lower income households, long commute times for workers, and difficulties for employers in attracting and retaining low-wage employees (Blumenberg and Wander, 2023). At the same time, high productivity areas limit housing production by implementing restrictive local zoning regulations. These policies artificially restrict housing supply and raise housing prices above the cost of construction, introducing an inefficiency that exacerbates housing cost burdens. Duranton and Puga (2023); Glaeser and Gyourko (2018). One understudied policy that seeks to address both affordability and inefficient restrictions on supply is a "density bonus." Such policies relax restrictions on residential density for developers who agree to cap prices on a share of units in their proposed residential development and restrict occupancy of those units to low- or moderate-income households. Thus, density bonus policies generally incentivize mixed-income developments in places that otherwise would only have less dense and higher cost housing.

Economists often analyze market responses to government interventions, including those in the housing policy space. The economic and urban planning literature has investigated a variety of spillover effects of affordable housing production, primarily focusing on housing price spillovers, but also estimating impacts on homeowner turnover, displacement risk, and evictions. This paper contributes to this literature by paying special attention to local spillovers on housing production. Though general equilbrium effects of housing production can influence an entire commuting zone (i.e., a metropolitan statistical area), many housing policy decisions are made at the local level, making estimation of local effects highly relevant for public policy. In evaluating policies with supply-side incentives, a key concern is whether or to what extent additional supply crowds out supply that would have been provided by the market. In the context of affordable housing development, policymakers are likely interested to learn what share of units produced under a policy are inframarginal units, that is, units that would have been produced by the market absent government intervention. This paper estimates a crowd-out effect of mixed-income developments, focusing on crowd-out that occurs in close proximity to developments produced under a density bonus policy. Using as a case study developments constructed under Massachusetts Chapter 40B, this paper asks how approval of these mixed-income developments affect the quantity or type of housing that is constructed locally.

Mixed-income housing development may influence nearby residential construction through two main channels: a direct, supply-side effect, and an indirect demand-side effect. Both mechanisms operate through local prices, which determine the revenue a developer expects to make from building housing on a given plot of land local to a mixed-income development. The direct, supply side effect lowers local housing prices when an approved mixed-income development adds units to the local housing supply, particularly when units in the mixed-income development are closer substitutes for housing that would have been built in absence of the policy. The indirect, demand side effects also play a role through changing expected revenues developers anticipate from local housing markets. Approval of a dense, mixed-income development could affect local demand for housing if it changes local amenities such as increasing traffic congestion (i.e., negative spillover on amenities) or local business activity (i.e, positive spillover on amenities). Additionally, if residents have a distaste for living near dense housing, income-restricted housing, or residents of these units, then the mixed-income development itself could affect local housing prices. The sum of these supply and demand side effects will determine observed spillovers of mixed-income housing development on local housing construction, which I characterize as local crowd-out effects.

To empirically evaluate these issues, I study Massachusetts Chapter 40B, the oldest density bonus policy in the country. Chapter 40B came into effect in 1969 and since then has produced more income-restricted housing than any other program in Massachusetts Citizens Housing and Planning Association (2014). Importantly, Chapter 40B provides a pathway for a state-level override of local zoning board decisions, promoting a regional approach to housing production. Detailed data on Chapter 40B units makes this program an ideal case study of the impacts of density bonus policy. Chapter 40B offers relaxed density restrictions and a streamlined permitting process for residential projects that set aside at least 20-25% of units in the development to be sold or rented to low- or moderate-income occupants at below-market price. To understand how this policy affects local housing supply, I rely on the staggered timing of approval of developments under Chapter 40B and compare places with 40B approvals to places where a development has not yet been approved (i.e., differences in differences with staggered treatment timing.)

I find no definitive evidence of local crowd out of housing generally, but some crowd-out of local single-family housing. In the ten years following approval of a Chapter 40B development, 5 percent fewer single-family homes are constructed within 0.8 miles of a development approval location. This mild crowd-out only moderately undermines the overall supply-boosting effects of Chapter 40B. Using a back-of-the-envelope calculation to express subtract total number of crowded out units from total 40B units approved, I find that one 40B unit approval adds 0.95 unit to the housing stock after accounting for local spillovers within 0.8 miles and 10 years of 40B development approval.

This paper contributes to two strands of the literature: one examining spillover effects of housing production, and another quantifying crowd-out effects of affordable housing incentives. It is most closely related to Blanco and Sportiche (2025), who find that only large 40B developments decrease property values in the immediate vicinity and have small effects on local resident turnover. My study extends this work by analyzing the effect of mixed-income housing development on the *quanitity* of housing produced locally instead of effects on local prices.

By focusing on crowd-out effects (i.e., quantities of housing), this paper sheds light on how policies like Chapter 40B change the local housing stock, which is a first order concern for a supply-side incentive policy. From a theoretical perspective, measuring spillover effects of 40B on the number of housing units also contributes to the body of work that evaluates mechanisms that influence local prices. This work also aligns with the framework in Pennington (2021), who evaluates the impact of exogenous new construction on displacement, outlines supply-side and demand-side forces that contribute to spillovers, and finds evidence that the supply effect reaches geographically further than the hyper-local demand side effects. Notably, this paper did not find any spillover effects generated by affordable housing.

Other work on spillover effects of affordable housing generally focus on housing produced under the Low-Income Housing Tax Credit (LIHTC), a federal program that subsidizes mixed-income development through distributing tax credits. LIHTC, the largest affordable housing development incentive program in the U.S., differs from MA Chapter 40B in several key ways. First, LIHTC involves tax credits instead of relaxed zoning restrictions, and second, in practice LIHTC units are typically built in urban as opposed to suburban areas, where housing produced under density bonus policies are normally located. LIHTC units are also exclusively rental units (Chapter 40B can be rental or owner-occupied) and have deeper affordability requirements than Chapter 40B. [1] Diamond and McQuade (2019a) and Baum-Snow and Marion (2009) find heterogeneity in spillovers, finding that property values and housing prices *increase* near LIHTC developments in low-income areas, and housing prices *decrease* local to LIHTC units in high income areas. Baum-Snow and Marion (2009) also find increased homeowner turnover in areas surrounding a LIHTC development.

Finally, this paper contributes to limited literature on crowd-out effects of income-restricted housing. Empirical evidence in this literature has also generally relied on evaluations of housing produced under the

[1] A households must make 80% of or less of area median income to be eligible for an income-restricted Chapter 40B, where the relevant restriction is 60% AMI for LIHTC units.

Low-Income Housing Tax Credit. Sinai and Waldfogel (2005) analyzes market-rate and affordable housing supply across U.S. metropolitan areas and find that after accounting for crowd-out, subsidized housing units add between 1/3 and 1/2 of an additional unit of housing, and crowds out private market rate housing to a lesser extent in more populous markets. In other words, in more populous markets, subsidized housing is better able to add to the housing supply rather than completely replace market-rate housing units that would have been constructed otherwise. Eriksen and Rosenthal (2010) finds substantial crowd-out of market rate rental housing, but less impact on owner-occupied housing (though estimates of the latter are imprecise). Soltas (2024a) finds that the Low-Income Housing Tax Credit causes development to occur *sooner* but does not add many units to the housing supply, since many of the units constructed under LIHTC are ultimately inframarginal.

This study of Chapter 40B advances the crowd-out literature by examining effects at a hyperlocal level, capturing effects adjacent to developments. Scant previous attention has been paid to the spatial distribution of crowd-out at this level of geography. Baum-Snow and Marion (2009) comes closest, estimating the effect of LIHTC units on local residential construction (within 1km rings). They find that one LIHTC unit adds 0.8 rental units to the local housing stock, but only 0.37 rental units in gentrifying areas.

The remainder of the paper is organized as follows. Section 2 provides background on the Chapter 40B policy. Section 3 outlines a conceptual framework for understanding how supply- and demand-side factors contribute to local spillover effects of housing development. Section 4 walks through estimation strategies for quantifying spillovers, setting up the staggered differences-in-differences specification. Sections 5 and 6 describe the data used for empirical estimation and present the results, including heterogeneity analysis. Section 7 concludes with discussion of implications of findings and areas for further research.

## 2 Massachusetts Chapter 40B

Massachusetts Chapter 40B is a statute enacted in 1969 which promotes the production of mixed-income housing developments in higher-income suburban areas by relaxing restrictions on density for qualifying residential developments and approving these developments under a streamlined permitting process. A proposed residential development can qualify to be approved under Chapter 40B if at least 20-25% of its units are income-restricted. More specifically, 25% of units in a development must be sold or rented to households making 80% or less of the area median income (where a location's area median income is determined by guidelines published by US Department of Housing and Urban Development). Alternatively, a development can qualify for Chapter 40B by selling or renting 20% of units to households making 50% or less of the area median income. In each case, Chapter 40B requires that the price of income-restricted units be sold or rented at or below a price cap set by the policy, which are prices that are "affordable" to a household making 80% (or 50%) of the area median income, with some variation based on bedroom size of the unit/household size. In this context "affordable" means sale or rent prices are set such that a household spends 30% or less of their income on housing costs (i.e., rent or mortgage and property taxes). Aside from income-restricted units, the remaining 75% (or 80%) of units in a 40B development may be rented or sold at market prices, without price restrictions. Thus, most Chapter 40B developments are not 100% income restricted, but instead are mixed-income developments.

**Zoning and Chapter 40B Production**

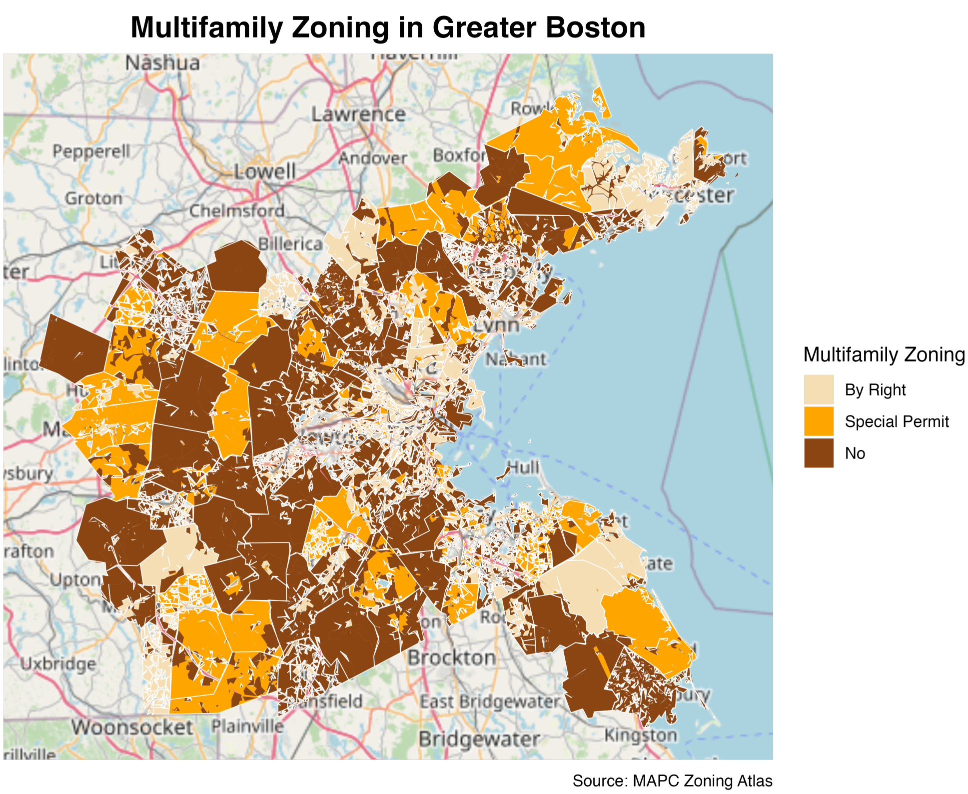


Figure 1: Metropolitan Area Planning Council Zoning Atlas

For a developer, the main advantage of using Chapter 40B is that it allows for production of housing that is more dense than would otherwise be allowed in much of Massachusetts, especially in Eastern Massachusetts/Greater Boston. Figure 2 shows that in much of Greater Boston, even a 2-unit condo would not be allowed as-of-right. Chapter 40B allows for this type of structure to be constructed in many more places. Since the most salient feature of the Chapter 40B policy is the relief it provides on density restrictions, we expect it will most commonly be used in areas that do not otherwise allow dense housing. Descriptive evidence shows this is indeed the case. The Metropolitan Area Planning Council provides a zoning atlas that summarizes zoning code regulations across much of Greater Boston. Using this data, I take a look at how 40B production and housing production overall map onto existing local zoning regulations (see: appendix table 3). This data does not cover the entire state of Massachusetts and so does not provide a comprehensive picture of where 40B units are produced, but for developments within Metropolitan Area Planning Council's catchment area, it does allow some analysis of zoning characteristics of places where 40B developments and other housing are approved or constructed. [2].

The summary statistics show evidence of what we would expect ex-ante. Among multifamily addresses in the Warren Group housing inventory, 40B addresses are more likely than non-40B addresses to be located in areas where multifamily housing is not normally allowed. If we look at addresses in the Warren Group data, only 16 percent of units that are identified as 40B multifamily units are built in places were any multifamily permitting is allowed as-of-right. 55 percent are built in places where no multifamily housing is (usually) allowed [3]. By contrast, 64 percent of non-40B multifamily units are built in places where multifamily housing is allowed as-of-right. Only 22 percent of them are constructed in places where multifamily housing is not (normally) allowed.

### The Permit Approvals Process

A second important feature of Chapter 40B is that it allows developers to take advantage of a relatively conslidated approvals process. Rather than apply for many different permits required in the housing production process, developments approved under Chapter 40B get a "Comprehensive Permit" which combines several of these permits into a single application process. A more detailed explanation of the general process of getting a 40B permit is outlined in appendix 14. All proposed developments must first be deemed eligible to be approved under 40B. To get an official eligibility decision, a developer applies to get a project eligibility letter one of the state's authorizing agencies (such as MassHousing, Massachusetts Housing Partnership, or Massachusetts Executive Office of Housing and Livable Communities). After a project is deemed eligible, the developer proposes the project to the respective municipality's local zoning board of appeals for approval.

[2]National Zoning Atlas has since put together information on zoning in the entire state of Massachusetts, though it is anticipated that this data will not become available for researcher use until early 2026 or later.

[3]Even in places with restrictive zoning regulations, variances can be allowed, to relax regulations for specific cases, so many aspect of zoning regulations are not hard and fast rules across the board

At this stage, the municipality can approve or deny the project, or request modifications.

If a local zoning board denies or modifies a proposal, in some cases the developer has the opportunity to appeal to the state to override local decisionmaking. If the proposed development lies within a municipality that does not have "enough" subsidized housing according to the state's official count, then the developer can appeal to the state to have its project permitted without the need for local approval. This override is possible in municipalities that do not already have at least 10% of their housing counted in the state's Subsidized Housing Inventory.

In practice, the Chapter 40B law can be understood to allow for construction of multiifamily, below market rate housing in suburbs of the Greater Boston area where zoning regulations are such that this type of housing would not otherwise exist. Aside from formal regulation however, local opposition to dense or affordable housing can slow or discourage housing production. It is not uncommon for the proposal of a Chapter 40B development to incur opposition from local residents Fisher (2008). Though 40B developments are far from immune from experiencing local opposition, the fact that the policy allows for state-level override of local decisions provides a mechanism for overcoming these two local barriers to housing development (i.e., zoning regulation and local opposition). The Chapter 40B policy does not involve any direct cost to the government beyond administrative costs for oversight of its implementation. Chapter 40B may be paired with development subsidies, or constructed using no subsidy at all. [4]

**Location of MA Chapter 40B Developments**

Chapter 40B is sometimes referred to as a "fair share" housing policy since it especially aims to incentivize affordable housing development in places which do not already have their "fair share" of afforable housing.

In the case of MA Chapter 40B, having a "fair share" of affordable housing in a municipality is defined as meeting the 10% threshold. 40B applies everywhere in the Commonwealth of Massachusetts except for in the City of Boston, which has particular zoning rules. Boston already has well above it's fair share of affordable housing. As of 2021, over 19 percent of Boston's housing stock is income-restricted City of Boston (2021).

This feature of the Chapter 40B is an important distinction, since Chapter 40B tends to incentivize housing in places where income-restricted housing and dense housing would not typically be prevalent. Sportiche et al. (2024) documents how Chapter 40B developments are, on average, constructed in areas that are higher income than the typical Massachusetts neighborhood, and that are higher income than typical areas where

[4] There are similar "density bonus"policies in other states, such as the Mount Laurel doctrine in New Jersey.

affordable housing is constructed under the Low Income Housing Tax Credit (LIHTC), the most prominent affordable housing incentive program in the United States. Because of this and other differences between Chapter 40B and LIHTC, we might not expect findings from the more abundant literature on LIHTC housing to be immediately generalizable to a policy like Chapter 40B, motivating the importance of study of affordable housing incentives beyond LIHTC.

# 3 Conceptual Model

There are two broad mechanisms–one on the supply side and the other on the demand side–that explain why approval of a Chapter 40B development might impact nearby housing construction. The first mechanism, on the supply side, is a more traditional version of crowd-out, where at least some of the 40B units themselves are inframarginal, substituting for a housing unit that may have otherwise been bulit absent the 40B policy. If some 40B units are inframarginal, then on average each 40B unit adds fewer than one housing unit to the local housing stock. The second mechanism, on the demand side, operates through the impact of housing development on local amenities such as traffic congestion. The 40B development itself could also be an amenity (or disamenity) if local would-be residents have a taste (or distaste) for living near a dense, mixed-income housing development. If the presence of a nearby 40B development affects local amenities and thus local demand for housing, developers may anticipate this change in demand and factor this into decisions about whether or when to build housing in these areas.

**Chapter 40B Developments Crowding out Local Development**

The first mechanism that explains why 40B development might affect local housing production is through a traditional crowd-out channel.

Figure 2, an adapted version of an figure appearing in Eriksen and Rosenthal (2010), illustrates this mechanism. Figure 2 displays supply and demand curves for housing units. Prices here are labeled as rents $R$, though rent price can also correspond to monthly mortgage costs in the case of owner-occupied units. Quantities $Q$ refer to quantities of housing units. In this model, the policy incentive (i.e. Chapter 40B in the case of this paper, Low-Income Housing Tax Credits in the case of Eriksen and Rosenthal (2010)) pushes the housing supply curve to the right, from $S$ to $S'$ as the policy incentive induces production of a number

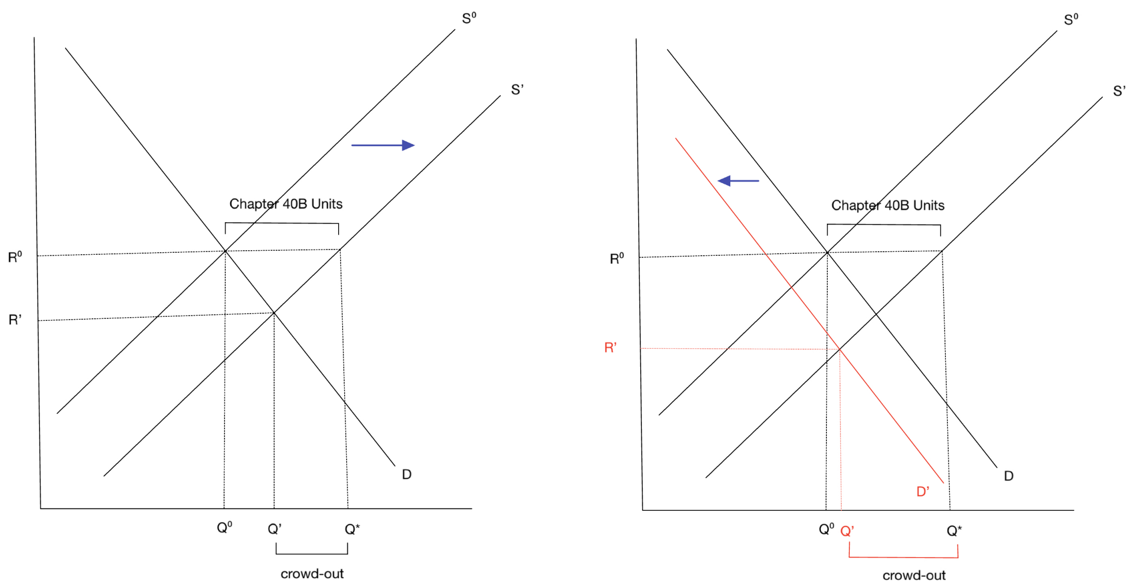


Figure 2: Chapter 40B developments can affect local housing supply in two different ways. First, they push the supply curve to the right, potentially crowding out local housing construction (left). Second, in areas close to an approved development, demand for housing may shift in response to local amenity changes (right).

of housing units equal to the difference between $Q*$ and $Q^0$. Because Chapter 40B relieves a regulatory barrier to housing production, it allows a developer an option to construct more housing units than would be possible without the policy. Because demand is not perfectly elastic, however, the shift of the supply curve does not push the quanitity of housing out to $Q*$. Instead, at equilibrium, housing is produced at level $Q'$, and the difference between quantities $Q*$ and $Q'$ represent crowded out housing units. These are units that may have otherwise been produced but are not, since they are crowded out by Chapter 40B units. These units are those that would have been a close substitute for the 40B units that were produced, or they may be units that were not produced if Chapter 40B development crowded out local resources (e.g., local administrative permitting capacity).

If a developer chooses to build 40B units in an area where housing units would have been constructed otherwise (i.e., even absent the 40B policy), some of the housing produced might simply be inframarginal, so that housing production is simply pushed forward in time, and other local housing production is diminished after approval of a 40B development. (This logic follows findings on LIHTC crowd-out from Soltas (2024a).) Similarly, if would-be local developers anticipate 40B housing to be a reasonable substitute for what they might have built in the area, these would-be developers may choose to pursue projects at a different time or place, viewing the local housing market as already being crowded with new housing options. Finally, 40B housing production may simply occupy local resources such that it makes it difficult to produce more housing locally.

40B production might take up the good land for housing production or might crowd out the bandwidth of local planning departments such that planning departments are unable to engage with additional proposed housing projects.

**Spillovers on Local Amenities**

Previous literature has established that affordable housing development affects local housing prices through the effect of housing production on local amenities (Diamond and McQuade (2019a)). Diamond and McQuade (2019a) describes this disamenity as a distaste for living near afforable housing. In the case of Chapter 40B we might also think of this preference as a distaste for living near dense housing. Other potential impacts on local amenities include impacts on traffic congestion or availability of municipal resources. If housing development affects local amenities and thus local prices, we can extend this logic and hypothesize that changes in local amenities and prices could also impact a developer's decision about whether or not to construct other housing in the area. For example, if affordable housing represents a local disamenity, developers may anticipate lower demand for housing in the area, which would decrease a developer's expected profits from constructing housing, making them less likely to choose to build more housing locally.

This mechanism is also illustrated in Figure 2. If producing Chapter 40B units produces a local disamenity, then the local demand curve shifts to $D'$ so that, all else equal, households are less willing to pay higher rents (or monthly mortgage costs) to live in the local area. Equilibrium rents fall, and only developers willing to supply at lower rents $R'$ will supply housing. Those who may have otherwise supplied local housing only at higher rent $R^0$ will no longer chose to build in the local area. (They may choose not to build, or instead build elsewhere.) 40B's impact on local amenities could affect the local political climate toward housing development. If local residents experience a disamenity due to approval of a 40B development, local opposition may build to local housing production and prevent the construction of additional housing in the area.

To test for empirical evidence of one mechanism v. another, this paper examines heterogeneity of treatment effect between 40B units that are likely to elicit one type of effect v. another. Compared to multifamily 40B units, single family 40B units are more similar to the type of residential unit that is constructed in Massachusetts cities and towns absent the 40B policy. Thus, I argue they are more likely to exhibit traditional crowd-out. On the other hand, since multifamily 40B units represent a more dense type of housing and are more different from exisiting local housing, they are more likely to affect local amenities in the area through local residents preference (or distaste) for leaving near dense or income-restricted housing.

## 4 Empirical Strategy

To examine local spillover effects of Chapter 40B developments on local housing supply, an ideal experiment would randomize locations of Chapter 40B developments such that we would be able to compare places where 40Bs developments have been approved to places where 40B developments have not been approved. In practice, however, locations of 40B developments are non-random and likely correlated with housing production outcomes, and thus we may be concerned that a simple comparison of places with versus without 40B developments would violate a parallel trends assumption. We cannot reasonably assume that an area with a 40B development would have had the same trajectory of local construction outcomes as a non-40B location in the absence of the 40B policy. When compared to places without 40B approvals, we might expect places that have 40B approvals to be places that are where land is available for residential development, or are otherwise more profitable for housing development.

Because areas with 40B approvals are not likely comparable to areas without, it is necessary to consider an approach which relies on a different comparison. While it is not reasonable to assume the *location* of a 40B development is random, it is a weaker assumption to assume the *timing* of approval of a 40B development is random, and use an identification strategy that relies on variation in *timing* of approval of a 40B development, rather than variation in location. My main specification will compare areas with approved 40B developments to places that will get 40B developments approved, but have not as of yet. If the timing of 40B approvals are orthogonal with respect to trends housing production outcomes, this comparison will produce an unbiased estimate of the treatment effect of approval of a 40B development.

To accomplish this, my methodology relies on differences-in-differences with staggered treatment timing. This approach compares areas with approved 40B developments (i.e, treated areas) to areas which have not yet had 40B developments approved (i.e., not-yet-treated areas). This is the differences-in-differences methodology described in Callaway and Sant'Anna (2021), which was also used as an alternative specification to quantify spillover effects of 40B developments on local prices in Blanco and Sportiche (2023). In this set up there is no pure control group, as all areas are eventually treated. Primary outcomes of interest are the quantity of housing units constructed in each year. This staggered differences in differences approach avoids the "forbidden comparison" pitfalls of the static two-way fixed effects model, which, in a staggered treatment setting, can attenuate treatment effect estimates by putting negative weights on some comparisons included in estimation, specifically on "forbidden" comparisons between later treated units and earlier treated units.

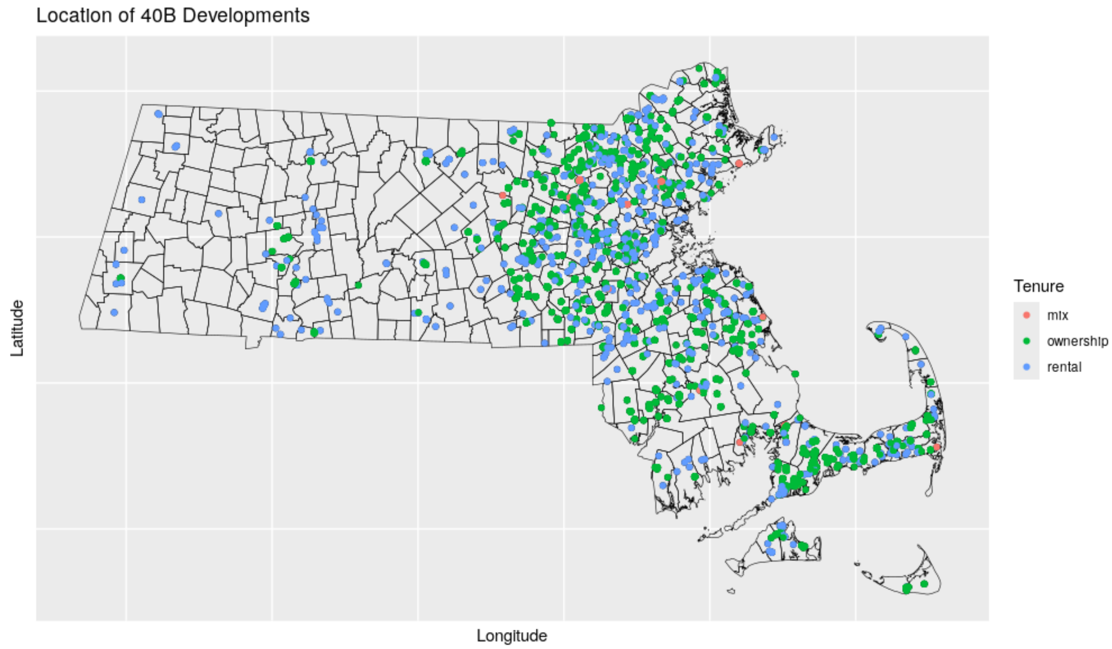


Figure 3: Location of Chapter 40B Developments

**Setup**

I set up my regression at the development area by year level, where each development area is defined as the area lying within 0.8 miles of the site where a 40B development is eventually approved (e.g., Figure 4). To define each development area, I begin by using the GPS coordinates provided in the data on Chapter 40B developments, which includes a GPS coordinate for each unique development. Next, addresses from the 40B data were grouped by development, and the maximum distance between any two addresses in that development was calculated. If the maximum distance exceeded 500m, the development was manually searched and inspected on Google Maps. If latitude and longitude were found to be incorrect they were updated. If addresses looked erroneous, they were edited or dropped from the development. If addresses within a single development were clustered in several distinct areas, a secondary "cluster" was created to distinguish between the localities. Of the 1,129 developments available in the data, 1,117 consist of a single cluster of housing units, nine developments included two separate clusters each, and three developments had three clusters each. Development areas were be drawn around individual development "clusters," although from this point forward in the paper, I will refer to them as "development areas".

For each development area, the timing of treatment is defined as the year in which the respective 40B development's building permit is approved. A more detailed explanation of the 40B approval process is outlined

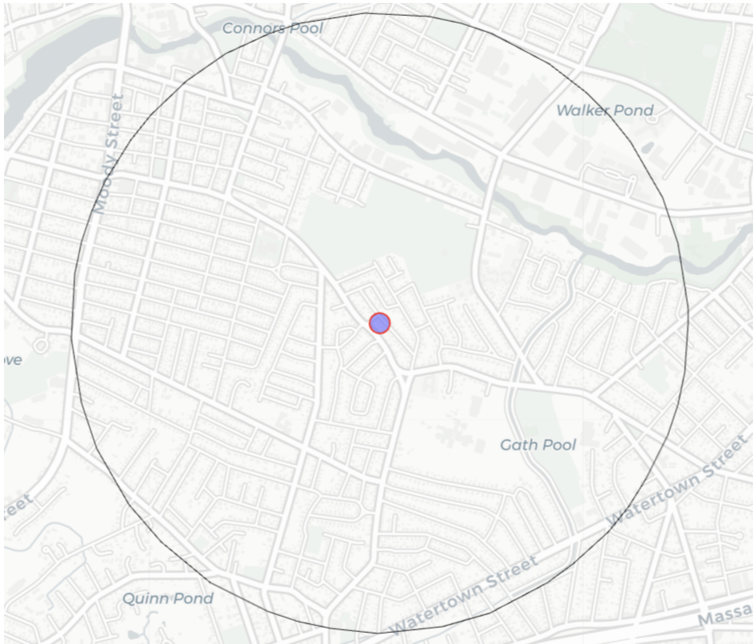


Figure 4: Example of a development area; and area of 0.8 mile radius drawn around a GPS coordinate defining a cluster of units within a 40B development.

in Appendix 8 Figure 14. Outcomes for each development area and each year are tabulated by taking counts or averages. For example, the primary outcome of interest, number of newly constructed residential units, is calculating by adding up all housing units in the data that are constructed within the defined development area $i$ and in the respective year $t$. For all outcomes (including transaction outcomes), observations corresponding to the respective development area's 40B units themselves are excluded, in order to avoid erroneously capturing mechanical (i.e.., direct) effects of the policy and instead capture spillovers.

**Static Two-Way Fixed Effects**

One typical differences-in-differences setup for estimating a treatment effect is a two-way fixed effect model such as equation 1. Here, $i$ indexes a treated unit (i.e., development area) and $t$ indexes a calendar year. $y_{it}$ is a time variant outcome (e.g. number of new units constructed in area $i$ and year $t$). $approved_{it}$ is a treatment for a unit that is treated (i.e., gets a 40B development approved) in time $t$ or before. $\gamma_i$ and $\delta_t$ are unit and calendar year fixed effects.

$$y_{it} = \tau \cdot approved_{it} + \gamma_i + \delta_t + \epsilon_{it} \tag{1}$$

When all units are treated at the same time, the $\tau$ coefficient quantifies an unbiased treatment effect as long as the parallel trends assumption holds.[5] As outlined in Goodman-Bacon (2021), however a problem arises when units are treated at *different* times, such as in our setting. The $\tau$ parameter can be understood to be a weighted sum of many 2x2 differences in differences comparisons between treatment and control units pre and post treatment, and in a staggered treatment setting, some of those 2x2 comparisons will consist of comparisons between treated units and units that were treated at an earlier time period, and will put negative weights on that difference. Secondly, the $\tau$ parameter doesn't reveal information about the timing of when treatment effects occur, which maybe important in some contexts. Instead $\tau$ estimates a single treatment effect for all post-treatment periods and assumes all relevant effects are captures within time periods $t$ in the data.

**Event Study**

A second way we might be interested in estimating a treatment effect is by using a simple event study, such as that defined in equation 2. The notation in this setup is similar to that of equation 1, except that $j$ indexes the year relative to treatment (i.e., treatment exposure time) and $approved_{i,t-j}$ indexes a unit in a time period $j$ years after treatment. Coefficients $tau_j$ quantify the difference in mean outcome between units $j$ years after approval and mean outcome in a reference year $j = -1$.

$$y_{it} = \sum_{\substack{j \neq -1 \\ -10 \leq j \leq 10}} \tau_j \cdot approved_{i,t-j} + \gamma_i + \delta_t + \epsilon_{it} \tag{2}$$

A distinct difference between the event study approach (i.e., equation 2) and static two-way fixed effects (i.e., equation 1), is that the event study allows us to estimate dynamic treatment effects, to quantify how outcomes evolve over time in response to treatment. In order to estimate a causal effect, this setup requires an assumption that, absent treatment, post-treatment outcomes would have remained the same as the reference year (i.e., the same as year $j = -1$). It also requires an assumption of no anticipation, such that units are not affected by the treatment prior to year $j = 0$. Importantly, however, the event study setup does not include a comparison with a control group, and is not able to take advantage of a parallel trends assumption, which would be a weaker assumption than assuming that outcomes would be identical across time (net of a time trend) absent the treatment. To be able to improve identification, we take advantage of a differences-in-differences estimation with staggered treatment timing.

[5] $\mathbb{E}[Y_{i,t_{post}}(0) - Y_{i,t_{pre}}(0) \mid approved_{i,t} = 1] = \mathbb{E}[Y_{i,t_{post}}(0) - Y_{i,t_{pre}}(0) \mid approved_{i,t} = 0]$

**Differences-in-Differences with Staggered Treatment Timing**

$$ATT(g,t) = \mathbb{E}_{\omega}[Y_{i,t} - Y_{i,t_{pre}}|G_i = g] - \mathbb{E}_{\omega}[Y_{i,t} - Y_{i,t_{pre}}|G_i > max\{g,t\}] \tag{3}$$

The parameters estimated in staggered differences in differences estimation are reflected in equation 3. For each group of development areas treated in year $g$, This approach is analogous to estimating a 2x2 differences in differences specification for each treatment cohort $g$ in each calendar year $t$, comparing means between treatment and control group outcomes in year $t$ net of baseline outcomes in a reference year (i.e., pre-period). The reference year is one year before approval in the setup presented in this paper. In each $ATT(g,t)$ 2x2 estimation, a treatment cohort's control group is defined as all areas that are treated after the treatment group (i.e., development areas in cohorts $g' > g$), and that have also not been treated as of year $t$ (i.e., development areas in cohorts $g' > t$).

There are many ways to aggregate estimates $ATT(g,t)$ into a parameter(s) that is of policy interest. Here we are primarily interested in estimating dynamic treatment effects, so for each year relative to treatment $j = t-g = [-10, 10]$, we take a weighted average of all relevant $ATT(g,t)$ where the weight of each $ATT(g,t)$ is equal to the number of observations in the respective treatment cohort $g$.

$$ATT(j) = \sum_{g} \omega_g \cdot [ATT(g,t)|t - g = j] \tag{4}$$

Since $j$ indicates the treatment group's length of exposure to the treatment, $ATT(j)$ estimates the effect of a 40B development on the number of residential units constructed within the surrounding buffer area at a time $j$ years after the 40B development is approved. $ATT(j) < 0$ would indicate that in the $j$th year after 40B development approval, there are fewer housing units constructed in the local area when compared to areas which do not yet have a 40B development, implying crowd out in relative year $j$. $ATT(j) > 0$ on the other hand, would imply crowd *in*.

In this approach, the key assumptions are, again, parallel trends and no anticipation. In the staggered treatment setting, the parallel trends assumption needs to hold for each treatment-control group pair used for estimation of $ATT(g,t)$, such that absent treatment, outcomes for control group $g$ evolve parallel to outcomes in control group $g' > \max\{g, t\}$. [6]. Intuitively, this means the timing of approval of Chapter 40B developments is orthogonal with respect to trends in construction activity in the local area. If, for example, Chapter 40B developments are only constructed in local areas at a time when those areas are already experiencing increases [decreases] in construction activity, this would violate the identifying assumption and lead to estimates of $ATT(j)$ that are biased upward [downward] and erroneously measure an effect such that Chapter 40B developments cause increases [decreases] in local construction. In order to address the concern that the data exhibit this type of violation of the identifying assumption, I estimate 1,000 versions of the staggered difference in differences design with placebo timing. In these placebo versions, I randomly reassign the approval timing of 40B developments and re-estimate treatment effects (results of this estimation are in appendix 8.2). (Results of placebo timing estimation show that results of this model in this context are not driven by business cycles in housing production. The total average treatment effect for 10 years post-approval is in the 3rd percentile of placebo estimated treatment effects.) The no anticipation assumption $Y_{it}(g) = Y_{it}(\infty)\ \forall\ t \leq g$ means that before construction outcomes are unaffected by the treatment prior to the treatment itself. In our case, this means construction outcomes do not respond to approval of a 40B development prior to approval of a comprehensive permit. One example of a violation of this would be if local developers anticipate 40B permit approval and change their construction activity even before a 40B permit is approved. This type of violation seems unlikely in our setting, since timing of treatment is relatively early in the 40B process (i.e., approval of a permit, before construction begins), and outcomes observed are late in the housing production process (i.e., construction date of a building, which is the last stage of housing production). Local housing production would need to operate on a quick time frame for construction outcomes to be affected even before a 40B permit is approved.

In the primary specification for this model there is no "never-treated" group, as all units used in analysis are eventually treated. Instead, units treated last serve as a control group for those that are treated earlier. As a robustness check, I also estimate an alternative specification with a constant control group, where the last three treated cohorts served as a control for all treated cohorts. The main outcome of interest counts the number of residential units constructed in each buffer area in each year, excluding units that are constructed within the buffer area's 40B development itself.

[6] Parallel trends assumption: $\mathbb{E}[Y_{i,t}(\infty) - Y_{i,t'}(\infty)|G_i = g] = \mathbb{E}[Y_{i,t}(\infty) - Y_{i,t'}(\infty)|G_i = g']\ \forall\ t \neq t'$ and $g \neq g'$

# 5 Data

Implementation of the empirical strategy specified in the previous sections requires data on the location and timing of approval of Chapter 40B developments. To quantify outcome measures, analysis also requires data on location and timing of overall residential housing production (i.e., including non-40B residential units), in order to quantify housing production in areas surrounding Chapter 40B developments.

**Data on Chapter 40B Developments**

Data on Chapter 40B developments comes from information collected in Blanco and Sportiche (2023), which have been made publicly available on the Harvard Dataverse portal. This dataset is at the address level, where each address contains several units within a single 40B development. Several addresses may correspond to a single development. For analysis, this project checked these data against publicly available information obtained from websites of individual 40B residential developments to make edits to the data as needed.

Housing characteristics in the dataset are at the level of the Chapter 40B development and include the street address(es) of each development, the date a building permit was obtained for that development, the subsidizing agency, the type of development (i.e., single family, condo, townhouse, etc.), the tenure of the development (i.e., ownership or rental), and an estimated count of the number of income-restricted and market rate units in the development. [7]

In total, the data used for analysis include 992 developments permitted 1965-2020 (Figure 5). About half of these developments are rental and about half are owner-occupied. Rental developments are generally much larger and include more housing units than owner-occupied developments. 43,368 units in the data are rental units, and 15,004 are owner-occupied. For about one percent of developments, we also found evidence that the development is not yet completed.

The summary statistics above describe developments that are included in the main estimation sample, but not all 40B developments available in the raw data are included for analysis. Of the 1,128 developments listed in the raw data, 992 are included for analysis and 136 are excluded. The most common reason that

[7] The original data include counts of the number of income-restricted units and a count of the number of "SHI" units in each development. "SHI" units are those that are counted in the Massachsetts Executive Office of Housing and Liveable Communities Subsidized Housing Inventory. This inventory counts all of the units in a 40B rental development, but only includes income-restricted units in 40B owner occupied developments. This means SHI units gives us the total count of units within a 40B rental development, but typically only a share of the units in an owner-occupied development. To estimate total units in owner-occuped developments, we assume that 25% of units in the development are income-restricted, and thus, for owner-occupied developments, the total number of units in the development is four times the number of SHI units. We also checked this information against counts provided in websites for individual developments.

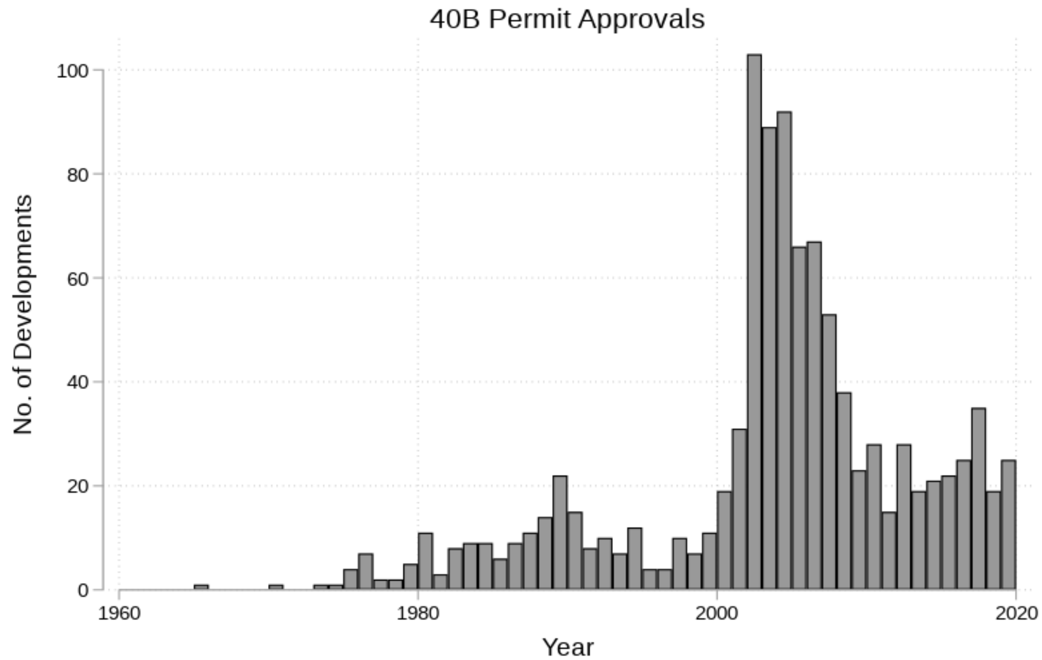


Figure 5: Treatment Timing; First permitting year for Chapter 40B developments.

developments are excluded is because the permitting date is missing. Of the 136 developments excluded, all but 11 are missing a permit date. Table 6 provided summary statistics comparing 40Bs included in analysis to those excluded. Excluded developments are less likely to be owner-occupied developments, are less likely to have been approved through MassHousing, and are more likely to have been approved through the Executive Office of Housing and Liveable Communities.

Finally, for estimation, 40B units are grouped by clusters, where each cluster is a group of units within a single 40B development. A 40B development may consist of one or several clusters. In each devel-

| | Estimation Sample, N = 992 | | Excluded 40Bs, N = 136 | | Comparison | |
|---|---|---|---|---|---|---|
| **Variable** | **Mean** | **SD** | **Mean** | **SD** | **Diff** | **p-value** |
| Total Units | 58.84 | 78.44 | 55.07 | 50.60 | 3.77 | 0.59 |
| Affordable Units | 15.70 | 20.63 | 14.89 | 12.57 | 0.81 | 0.66 |
| Owner-Occupied | 0.48 | 0.50 | 0.36 | 0.48 | 0.12 | 0.01 |
| Single Family | 0.29 | 0.45 | 0.29 | 0.46 | -0.01 | 0.89 |
| MassHousing | 0.38 | 0.49 | 0.18 | 0.38 | 0.20 | 0.00 |
| EOHLC | 0.45 | 0.50 | 0.71 | 0.45 | -0.26 | 0.00 |

Figure 6: 40B Developments in Estimation Sample v. Excluded Developments

opment, if the maximum distance between two addresses within the development was more than 500m, the development was manually inspected via Google maps. If addresses were clustered in distinct areas, addresses were grouped in separate clusters. Buffers (which I refer to as "development areas") are drawn around each cluster of addresses. There are a total of 1,007 unique clusters in the analysis sample.

### Data on the Housing Stock in Massachusetts

For data on housing production outcomes, we use a property inventory from Warren Group, which contains address level information on all residential properties in the state of Massachusetts, including address and GPS coordinates of each address, year of construction for each property, land use codes for each parcel on which the property sits (allowing the researcher to differentiate between single and mutlifamily properties), as well as other characteristics such as assessed value and lot size of properties.

For analysis we summarize housing production information at the development area by year level, counting the number of housing units constructed in each area in each year. Counts of housing constructed in each area exclude units within the area's respective 40B development itself, in order to quantify spillovers without include the direct effect of the policy. Across all 0.8 mile buffer areas in the data, when including only years that are within 10 years (before or after) 40B approval, there are an average of about eleven new residences constructed in each buffer area in each year. About seven of those are single family residences, and about four are multifamily residential units. On average, eight of those units are owner-occupied units.

Table 1: Summary Statistics - Development Area x Year

| | Mean | Median | SD |
|---|---|---|---|
| No. of New Single Family Units | 7.00 | 3.00 | 11.86 |
| No. of New Multifamily Units | 4.24 | 0.00 | 30.60 |
| No. of 40B units | 0.32 | 0.00 | 4.42 |
| No. of Subsidized Housing Units | 0.01 | 0.00 | 0.88 |
| No. of New Owner-Occupied Single Family Units | 5.98 | 3.00 | 10.57 |

N = 70,490

The housing production data used here is imperfect. The data consist of a cross section of residential units that exist as of 2023. This means we may not observe all housing that has been produced over time. If there is a unit that was previously constructed but was demolished prior to 2023, we would not observe that in the data. [8] Secondly, the housing inventory is an imperfect measure of existing housing units. When matching 40B addresses to housing inventory data in order to identify 40B units in housing inventory data,

[8] The Greater Boston area has some of the oldest housing stock in the country, which mitigates concerns we might have that we may undercount production in development areas that may have previously constructed but subsequently demolished residential units.

not all existing 40B addresses were able to match to the housing inventory. In particular, rental units are less likely than owner-occupied units to be able to be identified in Warren Group data. Of 4,328 addresses in the raw 40B data, 3,377 matched to Warren Group data (78% match rate). Only 27% of rental units matched, but 92% of owner-occupied units matched. 85% of mixed tenure developments matched to Warren Group data.[9] Because of this bias in the data, spillovers estimates should largely be understood to be mostly reflective of effects on owner-occupied housing production, and may underestimate any impacts on rental housing production.

**Data on Transactions**

Finally, to add context to main results and to replicate findings from previous literature, I use data on real estate transactions to estimate local spillover effects on real estate prices and sales activity. This data is also provided from the Warren Group and includes real estate transactions from years 2010-2023.

Table 2: Summary Statistics - Buffer Area x Year

| | Mean | Median | SD |
|---|---|---|---|
| No. of Transactions | 246.36 | 138.00 | 343.93 |
| No. of Single Family Sales | 141.41 | 99.00 | 133.51 |
| Single Family Sale Price (Mean, in 1000s) | 534.37 | 431.21 | 350.54 |
| No. of Condo Sales | 69.94 | 16.00 | 226.45 |
| No. of 2-3 Family Res Sales | 19.25 | 2.00 | 46.82 |
| No. of 9+ Apt Unit Sales | 1.09 | 0.00 | 4.98 |
| No. of Non-Residential Sales | 14.67 | 8.00 | 22.65 |

N = 2,090

# 6 Results

## 6.1 Difference-in-Differences with Staggered Treatment Timing

Figure 7 shows dynamic treatment effects derived from estimation of group-time treatment effects estimated using equation 3, which are then aggregated into dynamic treatment effects as described in equation 4. [10] Specifically, the figure shows point estimates and standard errors for $ATT(j)$ for $j \in [-10, 10]$. The x-axis of figure 7 shows treatment exposure time $j$, and for each year $j$ the chart displays a difference-in-difference of the average number of residential units constructed in treatment versus control buffer areas (i.e., treated

[9] Of the 1,144 40B clusters in the raw data, 754 matched to Warren Group data (66% match rate). 37% of rental clusters matched, 99% of ownership units matched, and 59% of mixed tenure clusters matched.

[10] Estimation includes development area and calendar year fixed effects, and standard errors are clustered at the municipal level, since many factors that affect housing markets (zoning regulations, school zone boundaries, etc.) are at the municipal level.

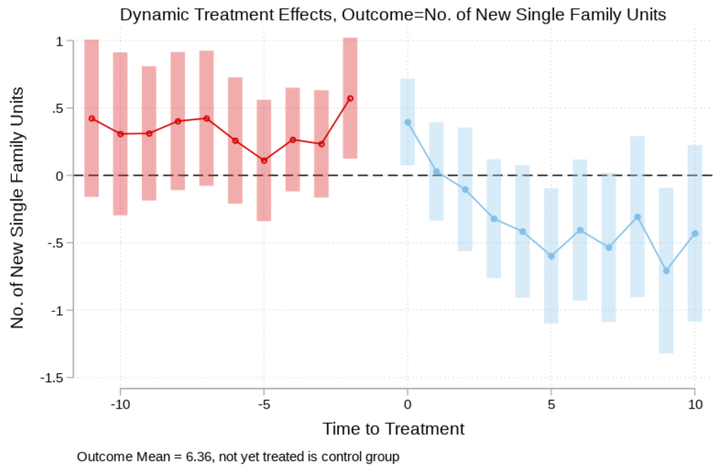


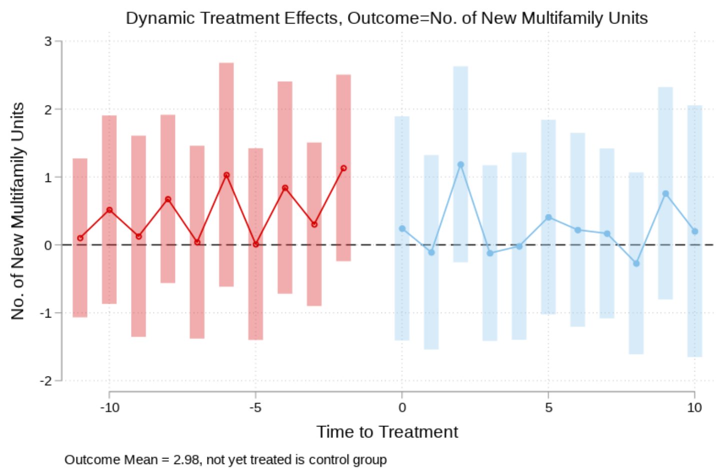


Figure 7: Dynamic treatment effect estimates $ATT(j)$ from equation 4. Effect of Chapter 40B approval on local single family construction (top) and multifamily construction (bottom). Chapter 40B approvals affect single family housing construction, but have no impact on multifamily housing construction.

versus not-yet-treated) for year $j$, net of differences between treated and control units in reference year $j = -1$. (Here, both the pre and post periods show "long differences", which are differences relative to reference period $j = -1$.) I replicate this estimation to look at two different outcomes: number of single-family units constructed in each development area-year, and number of multifamily units constructed in each development area-year.

Results of this regression show clear evidence of local crowd-out of single family units. For multifamily units, results do not detect any effect at all. These results align with what we might expect given zoning allowances typical in Eastern Massachusetts. Because multifamily housing is not allowed as-of-right in much of the area of study, crowd-out effects are most likely to appear in the production of single-family housing. Similarly, if we look at impacts on production of owner-occupied single-family housing (as opposed to all single family housing), results are quite similar, given that absent Chapter 40B, much of the housing that is constructed in the area of study is owner-occupied single-family housing, with limited rental availability (see: Appendix 8.1, Figure 16). Appendix 8.1 shows that 40B approval has no detectable impact on construction of other Chapter 40B developments locally, but do produce crowd *in* effects for other types of subsidized housing. [11]

In the ten years following approval of a 40B development, the crowd-out effect on single famliy housing increases over time. All else equal, an area with an approved 40B development constructs about 0.7 fewer single-family units than an area without a 40B approval during the ninth year after the 40B development is approved, representing a 10 percent decline. [12] Across years 0-10 post treatment, crowd-out leads to about 3.4 fewer single-family units per development area. Given an average of about 70 units produced over this lenght of time period in each development area, this average treatment effect corresponds to about 5 percent fewer single-family units constructed locally within 10 years following approval of a 40B project.

Benchmarking treatment effects against the number of 40B units constructed at within each development area allows me to estimate how many of 40B units are inframarginal with respect to nearbyl housing production in areas close by the 40B development. On average, each site has 8 40B single-family housing units and about 50 multifamily housing units (or, more accurately, about 50 housing units that are not identified as single family in the data). Comparing these estimates to the 10-year spillover effect of -3.4 single family housing units constructed locally suggests that 40B construction generally has had a net positive impact both on local single family housing construction (+4.6) and on housing construction overall (+54.6). When taking into account local crowd-out effects on single-family housing, each 40B unit results in a net increase of 0.95 housing units. [13]

### Heterogeneity

To better understand the mechanisms through which Chapter 40B affects local housing production, I explore heterogeneity in its impact on single-family construction across development types. Specifically, I compare the effects of owner-occupied versus rental 40B developments, single versus multifamily developments, and large versus small developments. If supply-side crowd out is the salient mechanism by which 40B developments affects local housing production, 40B developments that are a close substitutes to existing local housing (e.g., owner-occupied, single family housing) will crowd out local housing production more than developments that

[11] Here, subsidized housing units are those that Warren Group data identifies as being subsidized housing units in their housing stock inventory.

[12] This -0.7 unit treatment effect is relative to a mean of 7 single family units per year, meaning the estimated treatment effect shows a 40B approval reduces local single family housing production by about 10 percent during the ninth year after 40B approval.

[13] $\frac{58-3.4}{58} \approx 0.95$ (rounded values)

differ from housing typically produced locally. Alternatively, if a demand-side amenities effect has salience in this context, large multifamily 40B developments would likely impact local housing production. Results provide evidence of the former mechanism, with no evidence that the demand-side mechanism is as salient in this setting. Finally, I also examine heterogeneity according to the zoning regulations of the area in which housing is constructed, finding that negative spillovers occur in areas where multifamily housing is not allowed as-of-right.

Figure 8 shows owner-occupied 40Bs have more of an impact on local housing than do rental 40Bs. Overall, the single versus multifamily 40B distinction, in contrast, appears to have little impact on observed spillovers (see: Appendix 8.1, Figure 17), suggesting that substitutability between Chapter 40B units and locally available market rate housing operates primarily along the owner versus rental dimension.

[14] Among owner-occupied developments, however, multifamily owner-occupied 40B developments cause much more crowd-out than single family owner-occupied developments (See: Appendix 8.1, Figure 15.). The size of the 40B development does not seem to influence spillover effects. Taken together, these results provide little evidence that the demand side amenity channel impacts local residential construction in the ten years following 40B development approval. Instead, observed spillovers appear to be driven more by 40Bs entering the local market for owner-occupied housing and thus mildly crowding out local construction.

Finally, we might wonder how zoning regulations interact with spillover effects of Chapter 40B approvals. Much of 40B's usefulness rests on the fact that is allows for a state-level override of restrictive local zoning regulations, so how do spillover effects of this policy vary by zoning regulation allowances in areas where 40B is used? To address this question, I estimated dynamic treatment effects for two different outcomes: housing units constructed in areas where multifamily housing is allowed as of right, and housing constructed in areas where multifamily housing is *not* allowed as of right. Due to limitations on availability of zoning data for the full state of Massachusetts, this analysis is conducted only using information on housing units constructed within the Greater Boston area. [15] Results reveal that negative spillover effects occur in areas where multifamily housing is *not* allowed as of right. Logically, this aligns with the main finding, which showed crowd-out of single family housing but not multifamily housing.

[14] Analysis of pre-trends shows a small positive trajectory in single-family construction for areas that eventually approve single-family 40B units, indicating these areas were already buildling locally. In contrast, places that approve multifamily 40B developments do not exhibit the same pre-trend. These are places that are not necessarily already building out single family (or multifamily) housing. Given this pre-trend, evidence suggests that single-family 40Bs crowd out local construction *more* than multifamily 40B developments. Figures bounding estimates accounting for pre-trends are presented in the appendix.

[15] Efforts from the National Zoning Atlas team will likely make data for the full state available to researchers soon!

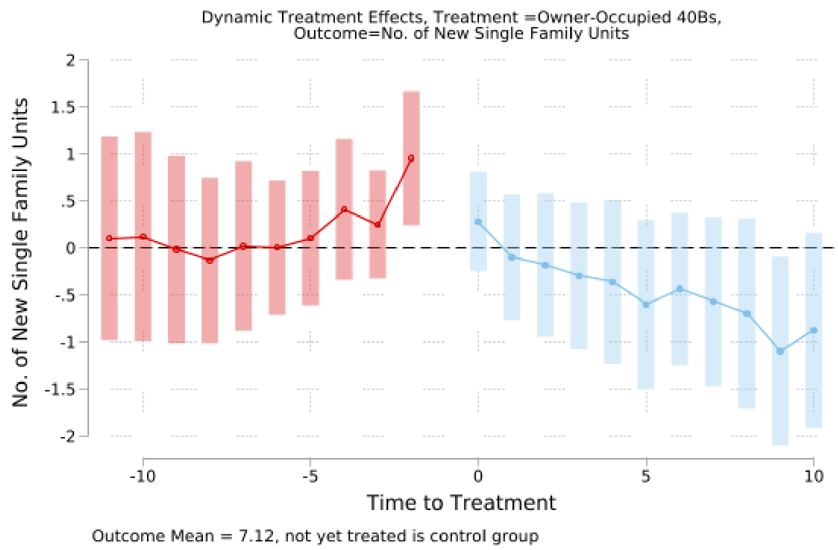


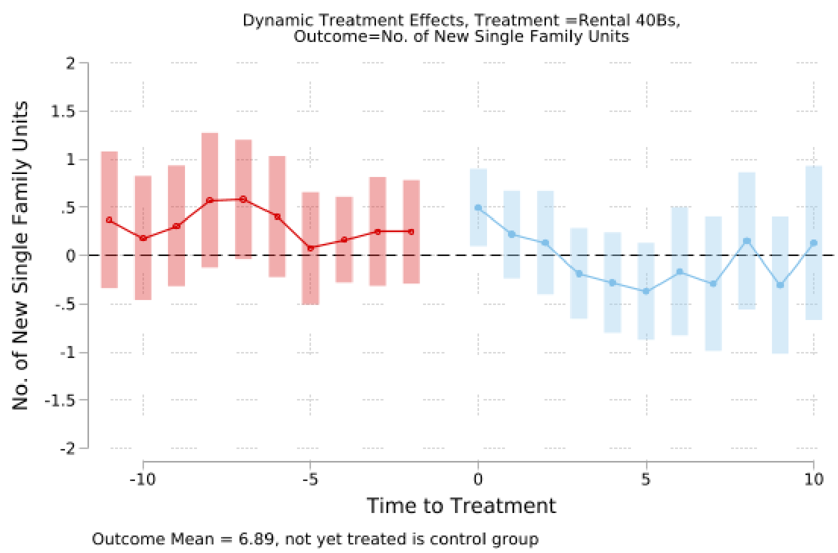


Figure 8: Dynamic treatment effect of approval of an owner-occupied 40B development (top) and approval of a rental 40B development (bottom) on local single family housing construction.

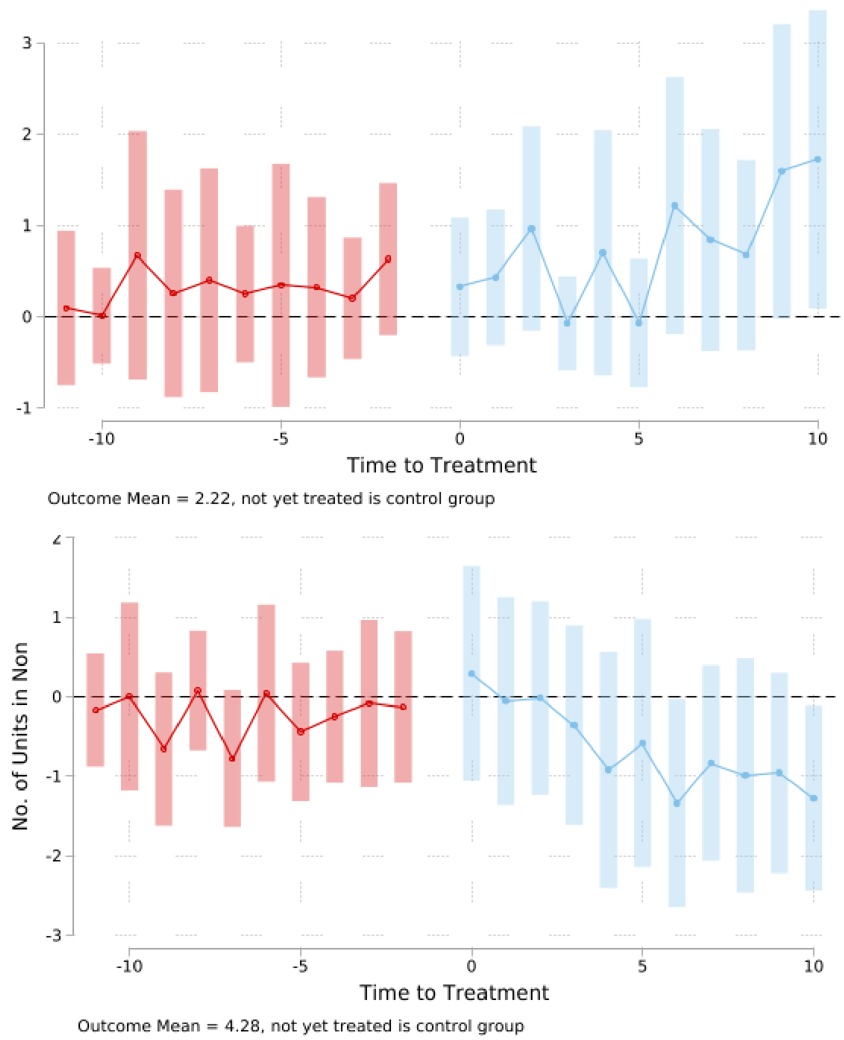


Figure 9: Dynamic treatment effect of 40B development approval on units constructed in areas where multifamily housing is allowed as of right (top) and on units constructed in areas where multifamily housing is not allowed as of right (bottom). In this estimation, due to limitation of availability of zoning data, only housing units lying within Greater Boston are included in analysis.

**Transaction Outcomes**

I next investigate impacts on other local outcomes, focusing on real estate sales activity in areas surrounding 40B developments. Using a differences-in-differences framework, I estimate spillover effects on both the number of sales in each development area-year, and the average sale price of properties transacted in each development area-year. Outcomes only include sale counts and prices for pre-existing housing units, where pre-existing housing units are defined as those which were constructed at least 10 years prior to approval of the nearby 40B development. Note that this analysis is constrained by data availability: while the psuedo-panel I constructed for housing production outcomes extends into well before the 1960s, the transaction records used for analysis begin only in 2010. As a result, the estimates presented here rely on a smaller sample of more recent 40B developments.

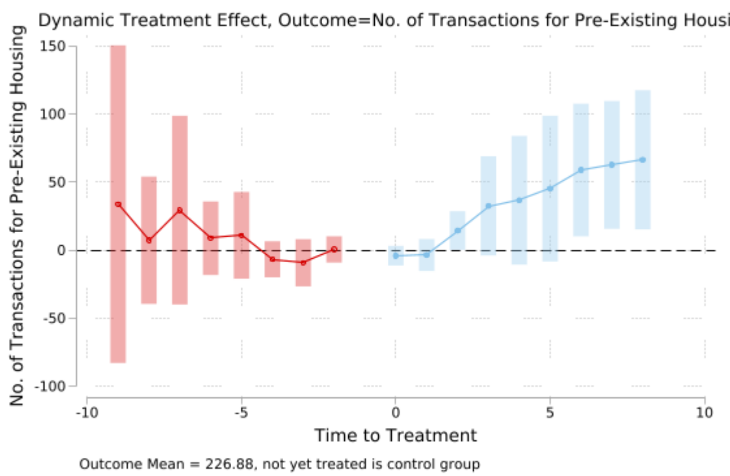


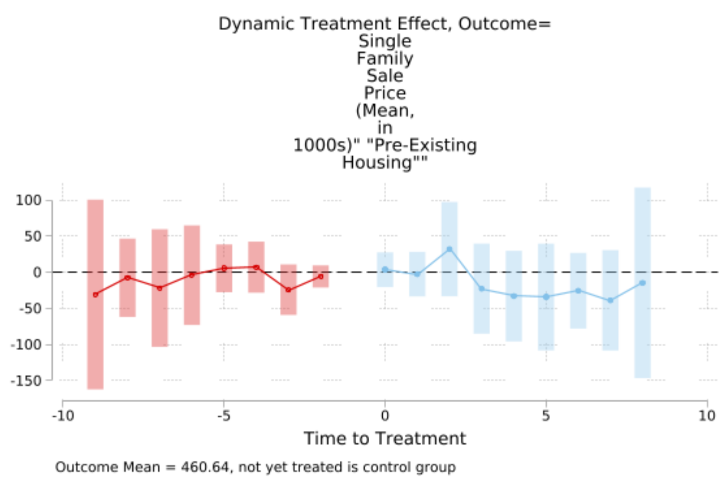


Figure 10: Effect of 40B approval on number of sales of pre-existing local single family home sales (top) and single family sale price (bottom). Pre-existing units are those that were constructed at least 10 years prior to the approval of the nearby 40B development.

Qualitatively consisent with Blanco and Sportiche (2025) and Diamond and McQuade (2019a) findings on spillover effects of affordable housing development, my results show increases in local sales activity and a falling single-family home sale prices following approval of a nearby 40B development. In the eighth year after approval, 40B areas have about 50 more transactions than areas that do not yet have 40B developments approved (Figure 10). This result is largely driven by single-family sales, although we do see some small effect on condos 2-3 unit properties (Figure **??** and appendix 8.1 Figure 19). Average single family sale prices also fall in areas surrounding 40B developments, though the fall in price is not statistically significantly different from prices prior to approval of the 40B development. In the eighth year after 40B approval, the mean single family home sale price drops by about $50,000 more in 40B areas relative to areas that do not yet have a 40B approval. [16]

**Robustness**

The main empirical specification uses a control group that varies for each cohort in the staggered differences-in-differences estimation. To check for robustness, I rerun the analysis using a control group that is constant across all cohorts in every estimation of $ATT(g, t)$. I define the control cohort to be the latest treated cohorts, i.e., development areas in which the 40B development was approved in years 2017-2020. Results are reported in appendix 8.2. This estimation produces qualitatively similar results to my main specification, though it does show treatment groups outperforming the control group on multifamily housing production post-treatment.

I also plot outcome averages for the treatment and control groups in appendix 8.3. Each plot shows average outcomes for a particular cohort of treated buffer areas, compared to its respective not-yet-treated control buffer areas. Here, production fluctuates over time, but in most cases, the control group outperforms the treatment group in single-family home production. In particular, many Chapter 40Bs were produced in the mid-2000s, and thus that treatment group is relatively large. When I plot the year 2006 (the largest treatment group) against its never-treated comparison group, both groups show declining production over time, but single family production declines more post-treatment compared to the control group.

Finally, I follow Roth (2022) and bound the estimates of my main result of a 5% reduction in local single-family home production within a 0.8 mile radius during the 10 years following Chapter 40B approval. This test asks how large a (linear) violation of parallel trends needs to be in order for my given empirical

[16] When calculating average prices in each development area-year, transactions with a price of 0 or 1 are excluded, since these are nominal prices used as place holders for transactions where the listed sale price does not reflect a valuation of the property (e.g. transactions for properties that change ownership between family members).

specification to detect it 50% of the time. I then draw a line with this hypothesized linear trend and re-calculate dynamic treatment effects. This approach rules out spillover effects larger than an 8 percent reduction in local single-family home production [17]. (See: Appendix figure 26.)

# 7 Discussion

Finding equitable ways to increase housing production is fundamental to addressing the housing crisis. This paper evaluated effects of Massachusetts Chapter 40B, a policy designed to increase housing supply through use of a state level override of local zoning regulations. From a policy standpoint, this paper furthers understanding of an understudied policy tool that increases and diversifies the housing stock. This study finds that mixed-income housing developments approved under Massachusetts Chapter 40B have a slight negative effect on nearby local single family housing construction. Two mechanisms may explain this effect. 40B units could be inframarginal, replacing single-family housing that would have been built in the absence of the policy, or they may effect local amenities, and in turn, local housing demand for housing. I find evidence consistent with the former mechanism, and find no evidence consistent with the latter. Further qualitative study could clarify the relative importance of these mechanisms.

Examining spillover effects of Chapter 40B on local housing construction is a step toward understanding the impact of density bonus policies on the overall housing supply. Previous literature has investigated crowd-out effects of the Low Income Housing Tax Credit, which subsidizes the cost of construction of income-restricted housing units without any accompanying change in zoning regulation. This study examines crowd-out effects of an affordable housing production policy that does not involve a subsidy or tax credit, but instead relaxes zoning restrictions in areas that normally would not have dense housing. This density bonus policy allows for more housing units to be constructed on a given plot of land and thus theoretically provides a setup in which fewer units constructed are inframarginal when compared to a policy that incentivizes housing production purely through the use of subsidy.

Negative spillovers estimated in this paper are quite small, suggesting that density bonuses such as Chapter 40B have the potential to add housing units to communities without crowding out housing that would be built absent the policy. Since density bonus policy incurs no direct cost to the government beyond administrative costs, a policy such as this one might be particularly attractive to state governments who wish to find cost-effective solutions for housing production. Note that the analysis only looks at *local* crowd-out, and does

[17] The upper bound of the treatment effect summed across years 0-10 = -5.81. $\frac{-5.81}{70} = -0.083$

not account for, for example, a developer who might substitute away from building near a 40B development and instead constructs housing in a completely different municipality in the Greater Boston. This focus on local effects is appropriate, given that housing policy is often influenced at the local level, particularly in Massachusetts. Furthermore, the analysis relies on counts of housing units and does not take into account how 40B affects detailed characteristics of the housing produced and how changes in type and price of housing affects the characteristics of households in the area. These factors are critical for understanding the overall distributional impact that Chapter 40B's has in local communities in terms of housing availability and affordability, neighborhood composition, and the welfare of local residents. Finally, analysis here is only able to evaluate effects within 10 years following approval of mixed-income developments under density bonus policy, so if sizeable affects occur further into the future, they would not be quantified here.

The estimates presented here relate to literatures on spillover effects of housing production and to literature on crowd out. The most closely related literature evaluates local price spillovers of Chapter 40B developments. Blanco and Sportiche (2025) finds on average no effect of 40Bs on local prices (the papers finds price effects associated with large developments), though these estimates are not quite comparable to the differences presented in this paper. Blanco and Sportiche (2025) estimates differences between the inner most ring to of a development area and the areas 0.5-0.6 miles away from the development. Other closely related papers on spillover effects of housing development include Diamond and McQuade (2019a), which finds a 2.5% fall in housing prices in high-income areas nearby units constructed under the Low-Income-Housing Tax Credit and Pennington (2021), which finds that rents fall \$100 within 100 meters of exogenously constructed new housing in San Francisco, and that exogenous new construction doubles the probability of new construction within 100 meters.

Other estimates of crowd-out in the literature include Soltas (2024b), which finds that for every ten units constructed under the Low-Income Housing Tax Credit, eight are inframarginal. LIHTC mostly moves investment forward in time. Sinai and Waldfogel (2005) found each LIHTC added 1/2 to 1/3 of a new unit to the housing supply after accounting for crowd-out. Eriksen and Rosenthal (2010) found that nearly all LIHTC unit production was offset by a reduction in market rate rentals units. Finally, Anagol et al. (2021) estimated the economic impact of zoning reform in São Paolo and found upzoning increased housing stock by 1.4% and caused a 0.4-0.9% price reduction.

A density bonus is but one option among several policies a government might choose to pursue in order to encourage housing production. The policy is effective because it increases and diversifies the housing stock without crowding out local housing supply and without direct cost to the government. If we compare a density bonus policy to a production subsidy (or tax credit), the lack of direct cost to the government for density bonuses is an especially attractive feature for governments that would have difficulty budgeting for additional subsidies, and or for areas where land prices and construction costs are high and thus incentivizing affordable production through subsidy would be expensive. A density bonus policy would be best suited for an area where pre-existing zoning regulations meaningfully limit production of housing types that would be affordable to moderate income households, such that relaxing that restriction represents a non-trivial change to a developer's options of what and where to build.

Since 2019, several U.S. states have passed upzoning laws that relax pre-existing restrictions on residential density. States such as Oregon, California, Montana and Massachusetts have passed laws that override local zoning and allow for dense housing types as-of-right. In a world where zoning regulations are generally less restrictive, a density bonus policy such as Chapter 40B becomes less relevant. Upzoning laws, however, are notoriously politically difficult to pass, and thus a density bonus policy might serve as an alternative second best solution. Historically, density bonus policies have been unique in their regional level approach to zoning. Massachusetts Chapter 40B has been in effect since 1969 and New Jersey's Mount Laurel doctrine has been in place since 1975, long before political winds shifted to allow for the more recent wave of state level upzoning policies. Furthermore, density bonus policies include an income restriction for a subset of units constructed, and thus provide an extra policy mechanism by which to address affordability more directly through a de facto tax on development.[18] For these reasons, a policy such as a density bonus can prove to be a useful tool in the tool box for state governments who want to encourage housing production.

[18]Better data on pricing of individual units constructed under these policies would be important for understanding more about distributional consequences for housing affordability.

## References


Anagol, S., Ferreira, F., and Rexer, J. (2021). Estimating the Economic Value of Zoning Reform. Technical Report w29440, National Bureau of Economic Research, Cambridge, MA.

Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, 93(5-6):654–666.

Blanco, H. and Sportiche, N. (2023). Local Effects of Bypassing Zoning Regulations in High-Income Areas.

Blanco, H. and Sportiche, N. (2025). There Goes the Neighborhood? The Local Impacts of State Policies that Override Municipal Zoning.

Blumenberg, E. and Wander, M. (2023). Housing affordability and commute distance. *Urban Geography*, 44(7):1454–1473.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-Differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Citizens Housing and Planning Association (2014). Chapter 40B, The State's Affordable Housing Law.

City of Boston (2021). Income Restricted Housing Report.

Diamond, R. and McQuade, T. (2019a). Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development. *Journal of Political Economy*, 127(3):1063–1117.

Diamond, R. and McQuade, T. (2019b). Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development. *Journal of Political Economy*, 127(3):1063–1117.

Duranton, G. and Puga, D. (2023). Urban Growth and Its Aggregate Implications. *Econometrica : journal of the Econometric Society*, 91(6):2219–2259.

Eriksen, M. D. and Rosenthal, S. S. (2010). Crowd out effects of place-based subsidized rental housing: New evidence from the LIHTC program. *Journal of Public Economics*, 94(11-12):953–966.

Fisher, L. (2008). What Gets Proposed, What Gets Approved, What Gets Appealed, and What Gets Built?

Glaeser, E. and Gyourko, J. (2018). The Economic Implications of Housing Supply. *Journal of Economic Perspectives*, 32(1):3–30.

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

Pennington, K. (2021). Does Building New Housing Cause Displacement?: The Supply and Demand Effects of Construction in San Francisco. *SSRN Electronic Journal*.

Roth, J. (2022). Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends. *American Economic Review: Insights*, 4(3):305–322.

Sinai, T. and Waldfogel, J. (2005). Do low-income housing subsidies increase the occupied housing stock? *Journal of Public Economics*, 89(11-12):2137–2164.

Soltas, E. (2024a). Tax Incentives and the Supply of Low-Income Housing.

Soltas, E. (2024b). Tax Incentives and the Supply of Low-Income Housing.

Sportiche, N., Blanco, H., Daepp, M. I. G., Graves, E., and Cutler, D. (2024). Can Fair Share Policies Expand Neighborhood Choice? Evidence From Bypassing Exclusionary Zoning Under Massachusetts Chapter 40B. *Housing Policy Debate*, pages 1–33.

# 8 Appendix

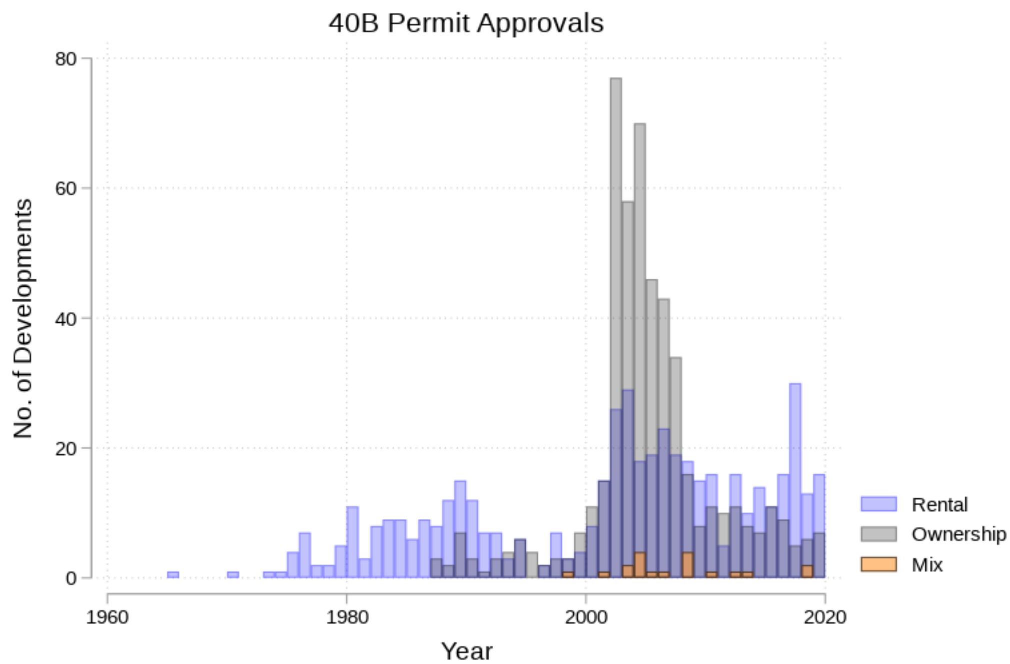


Figure 11: Treatment Timing

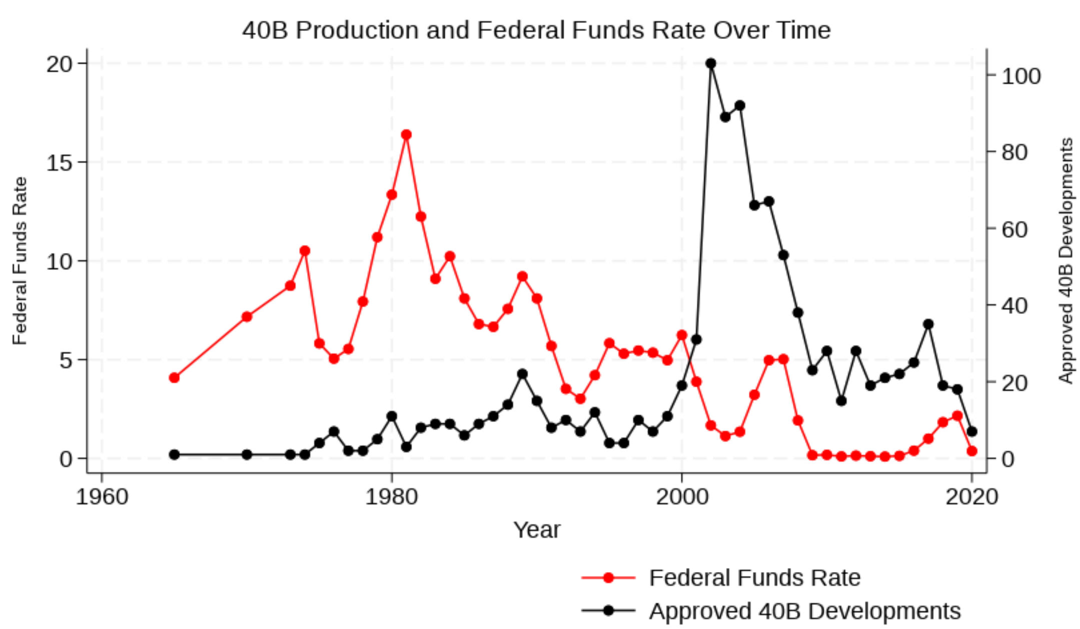


Figure 12: 40B Approvals v. Federal Funds Rate

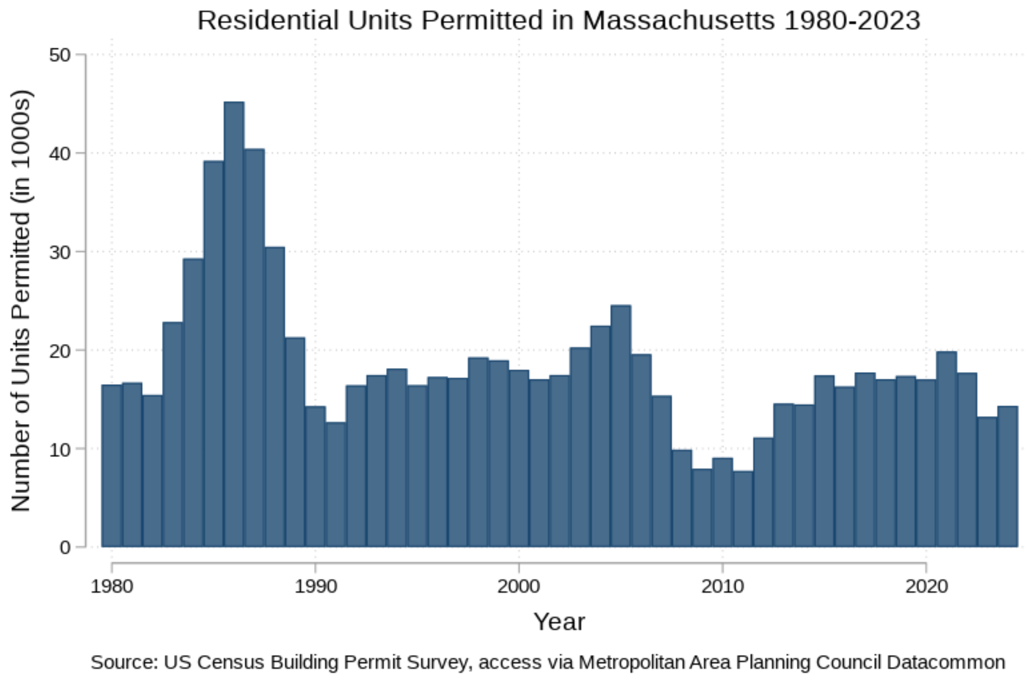


Figure 13: MA Permitting 1980-2023

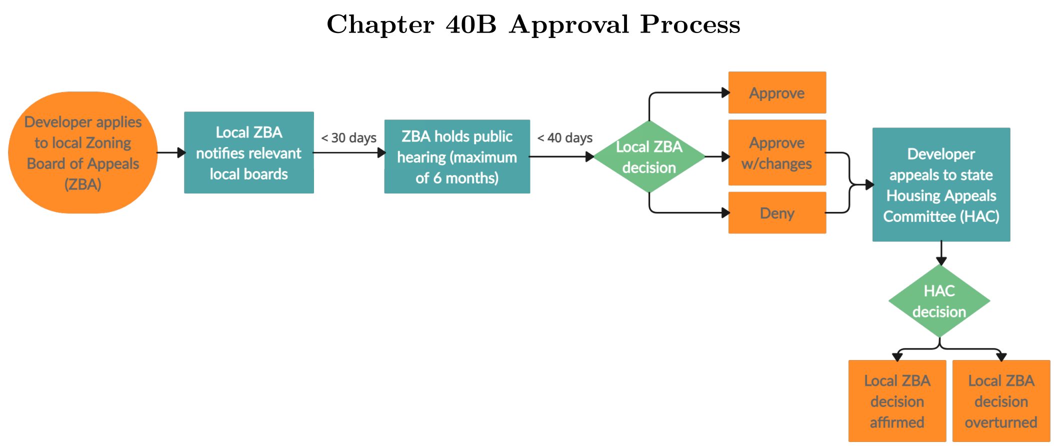


Figure 14: Chapter 40B Approval Process

| Zoning Use Code | All | 40B Single Family | 40B Multifamily | Non-40B Single Family | Non-40B Multifamily |
|---|---|---|---|---|---|
| Residential | 0.83 | 0.95 | 0.72 | 0.94 | 0.72 |
| Non-residential | 0.04 | 0.04 | 0.15 | 0.01 | 0.06 |
| Mixed Use | 0.11 | 0.01 | 0.10 | 0.04 | 0.21 |
| Other | 0.01 | 0.00 | 0.01 | 0.00 | 0.02 |
| **Multifamily (2)** | | | | | |
| No | 0.42 | 0.32 | 0.55 | 0.55 | 0.22 |
| By special permit | 0.20 | 0.39 | 0.28 | 0.24 | 0.14 |
| As of right | 0.38 | 0.29 | 0.16 | 0.21 | 0.64 |
| **Multifamily (3-4)** | | | | | |
| No | 0.43 | 0.32 | 0.55 | 0.55 | 0.23 |
| By special permit | 0.34 | 0.64 | 0.35 | 0.35 | 0.33 |
| As of right | 0.23 | 0.04 | 0.10 | 0.09 | 0.44 |
| **Multifamily (5-19)** | | | | | |
| No | 0.43 | 0.32 | 0.55 | 0.55 | 0.23 |
| By special permit | 0.38 | 0.64 | 0.43 | 0.38 | 0.40 |
| As of right | 0.19 | 0.04 | 0.02 | 0.07 | 0.36 |
| **Multifamily (20+)** | | | | | |
| No | 0.43 | 0.32 | 0.55 | 0.55 | 0.23 |
| By special permit | 0.39 | 0.64 | 0.43 | 0.38 | 0.42 |
| As of right | 0.18 | 0.04 | 0.02 | 0.07 | 0.35 |

Table 3: Summary of share of different types of housing (i.e.,, single v. multifamily; 40B v. non-40B ) that is constructed in areas with different categories of zoning allowances (i.e., multifamily allowed v. not allowed.). When compared to non-Chapter 40B multifamily units, Chapter 40B multifamily units are more likely to be built in areas where multifamily housing is not otherwise allowed as-of-right.

## 8.1 Heterogeneity

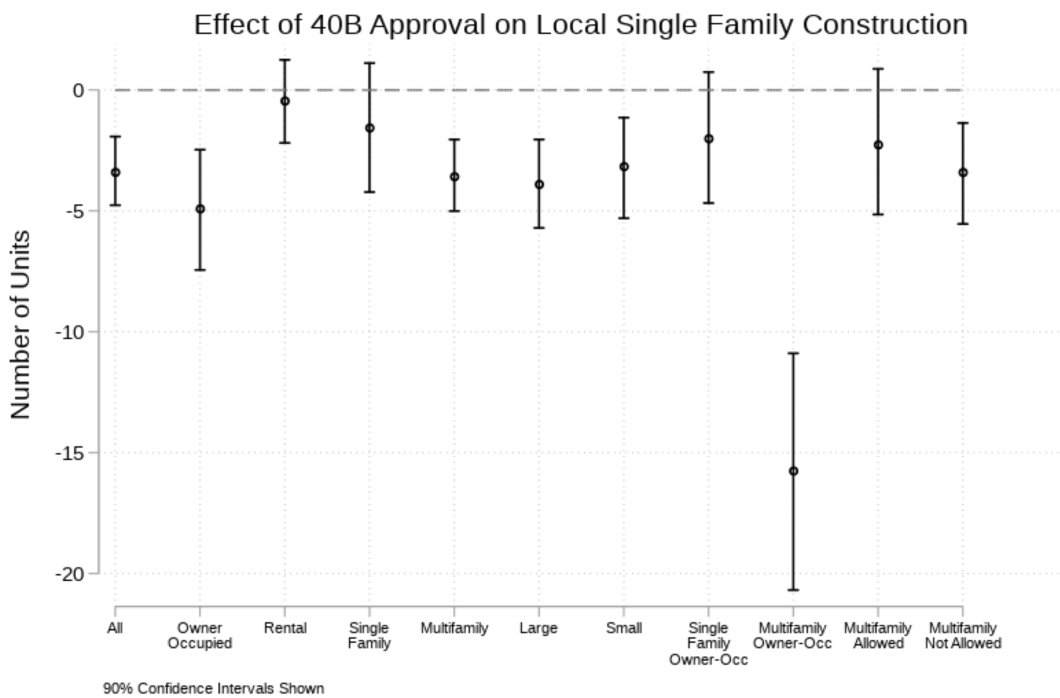


Figure 15: Cumulative 10-year treatment effect of 40B developments on local single family construction for specific subgroups of 40B developments, including owner-occupied 40Bs, rental 40Bs, single family 40Bs, multifamily 40Bs, large 40Bs (i.e., above median number of units), small 40Bs (i.e., below median number of units), single family owner-occupied, and multifamily owner-occupied..

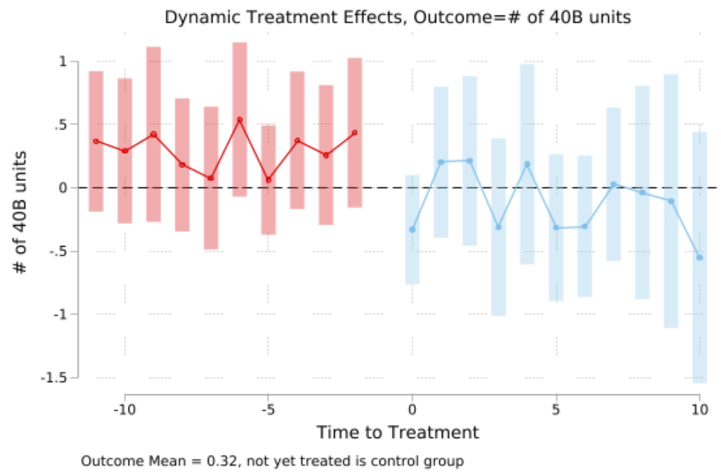


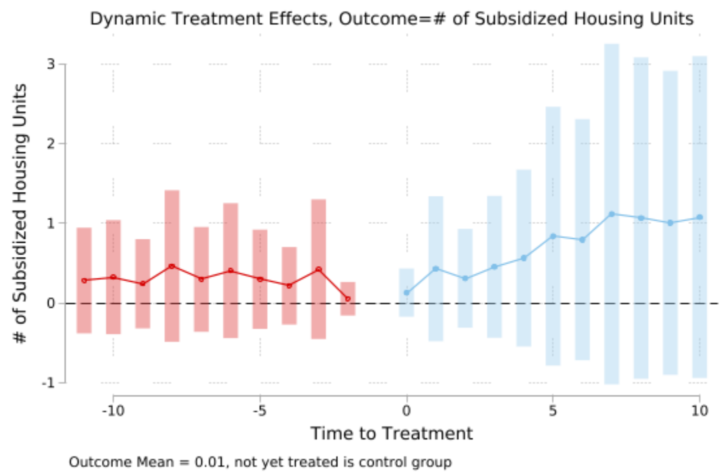


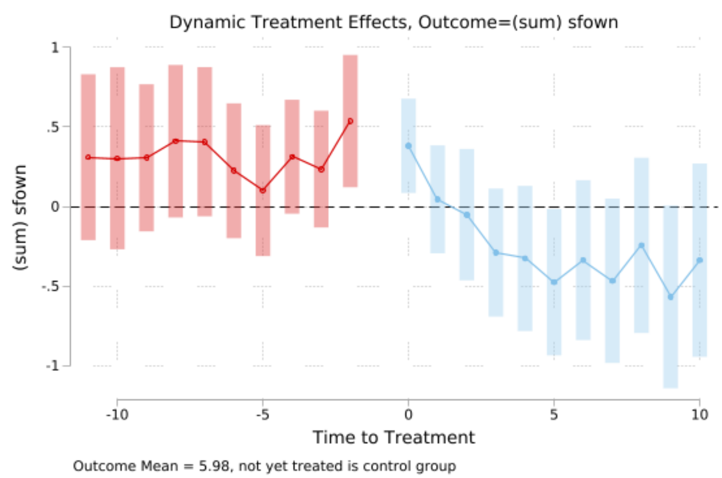


Figure 16: Effect of 40B approvals on local construction of other 40B units (top), other types of subsidized housing (middle), and owner-occupied single family units (bottom).

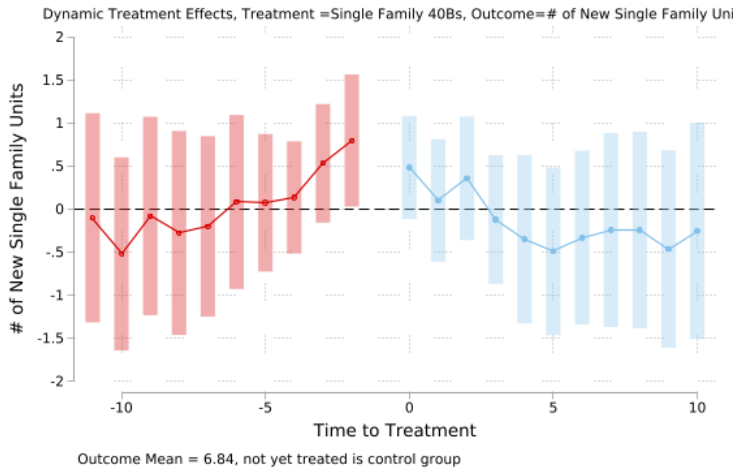


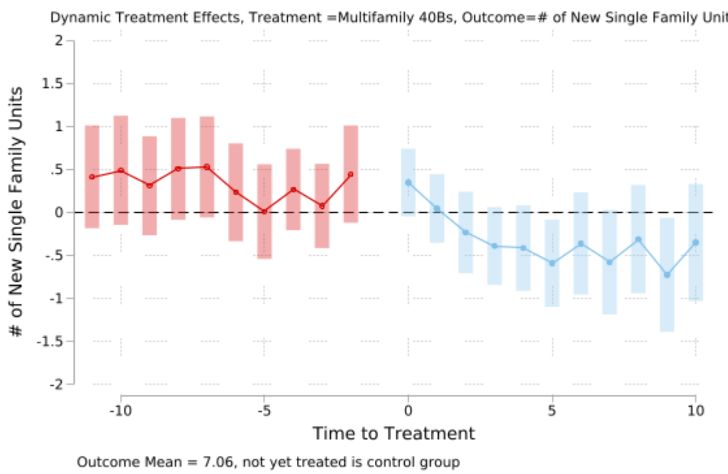


Figure 17: Effect of single family 40Bs (top) and multifamily 40Bs (bottom) on local single-family housing production. Both types of 40Bs have a similar effect on local single family housing production.

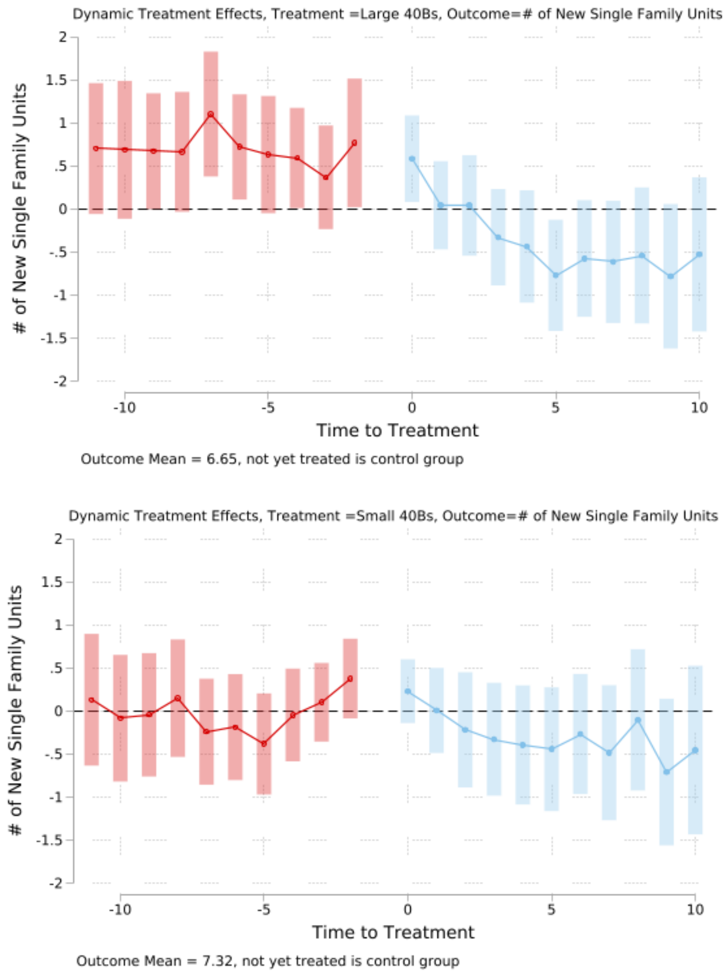


Figure 18: Effect of large 40B developments (top) and small 40B developments (bottom) on local single-family housing production. Large developments are the 50% of developments in the data with the highest numbers of housing units. Small developments are in the bottom 50% in terms of numbers of units. Both types of 40Bs have a similar effect on local single family housing production.

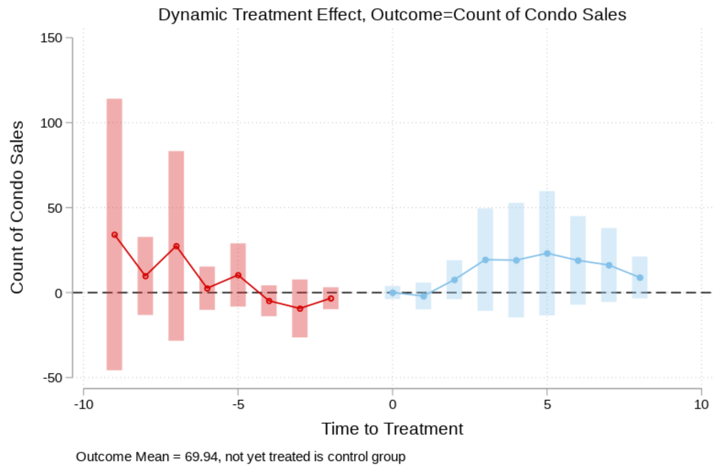


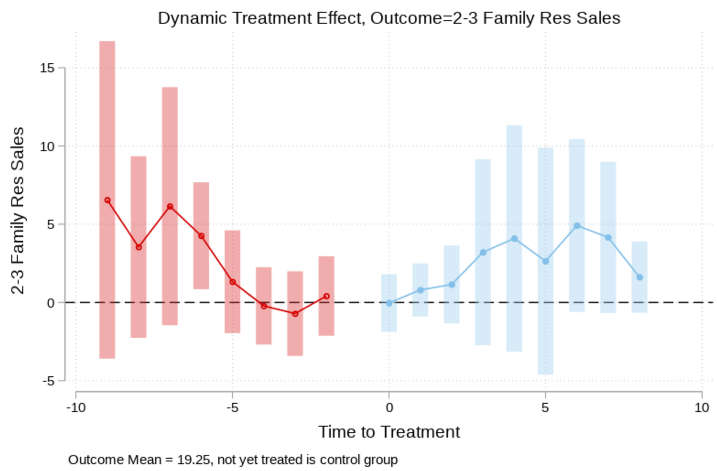


Figure 19: Effect of 40B approval on the number of sales of local condonimiums (top) and other units with 2-3 unit buildings (bottom).

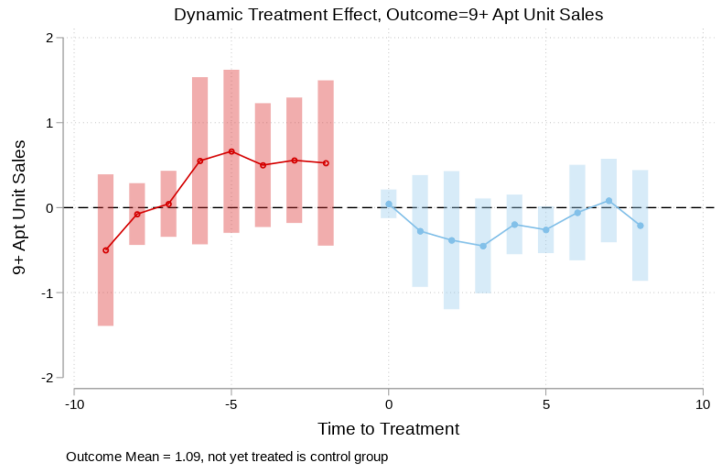


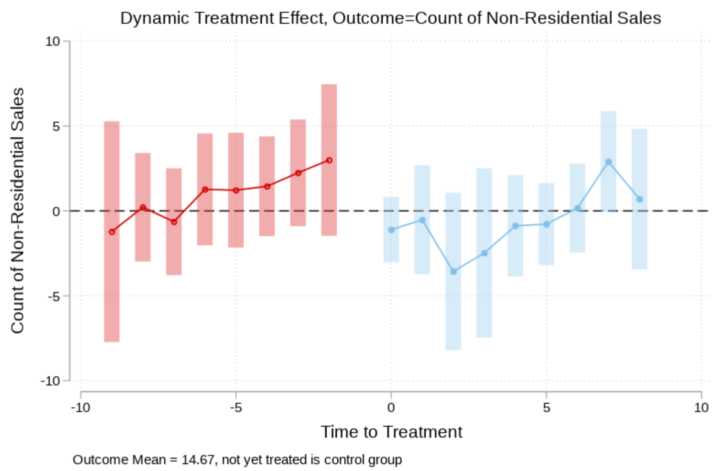


Figure 20: Effect of 40B approval on the number of sales of local 9+ unit apartment buildings (top) and number of non-residential transactions (bottom). 40B approvals largely do not affect these outcomes.

## 8.2 Alternate DiD Specifications

### Placebo Timing

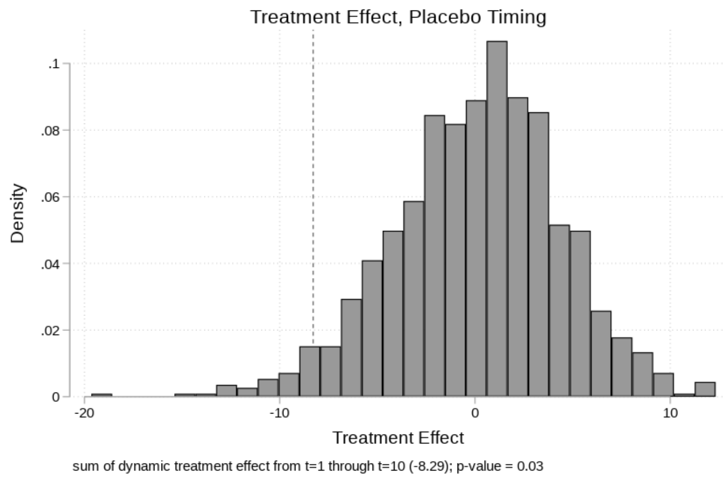


Figure 21: Staggered Differences-in Differences with Placebo Timing. Chapter 40B units were randomly reassigned permit approval years (i.e., randomly reassigned treatment timing), and a staggered differences-in-differences was restimated. This procedure was repeated 1,000 times. Results of this simulation are shown here. Treatment effect estimated with observed data is shown by the dashed line.

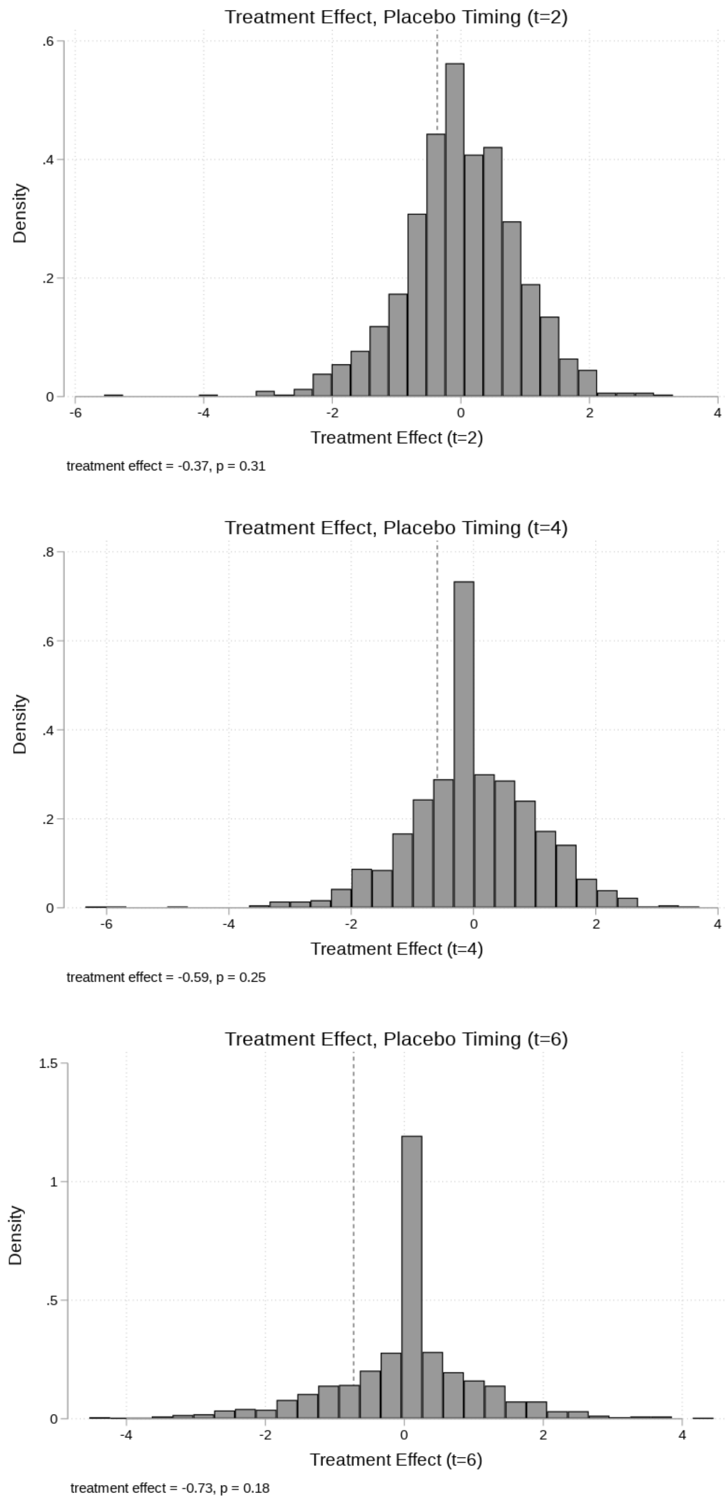


Figure 22: Staggered Differences-in Differences with Placebo Timing

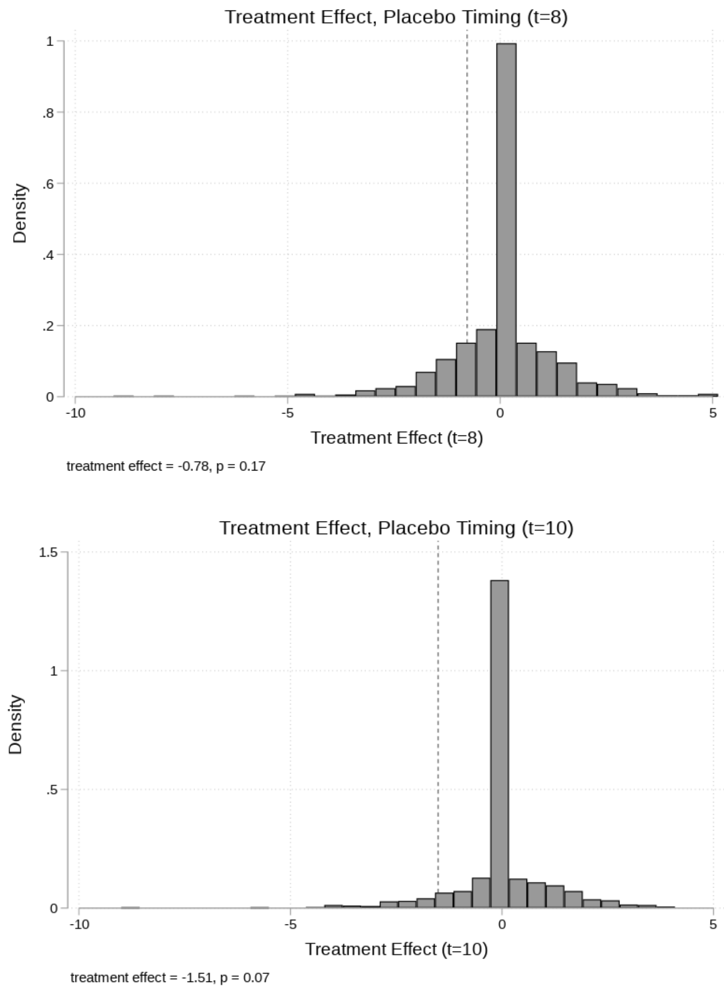


Figure 23: Staggered Differences-in Differences with Placebo Timing

**Constant Control**

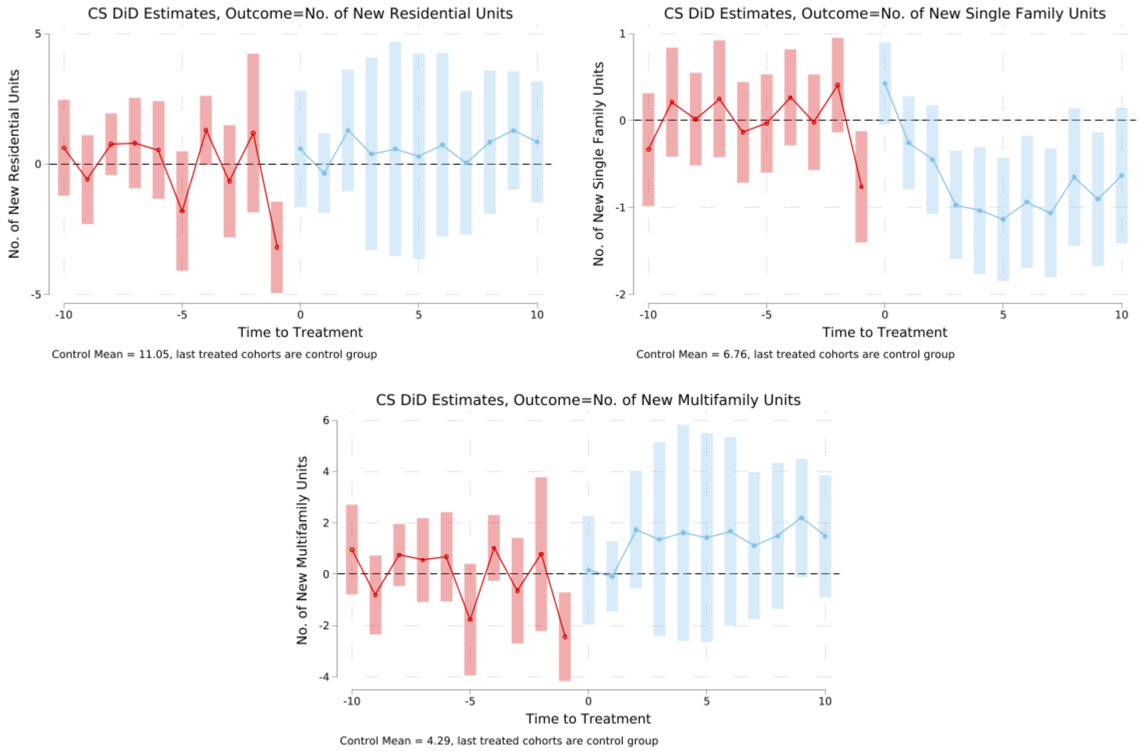


Figure 24: Staggered Differences-in-Differences with Constant Control Group

Figure 25: Dynamic treatment effect estimation resulting from a specification with a constant control group. Treatment cohorts are treated between years 1965 and 2006. The control group for each cohort includes development areas treated in years 2017-2020.

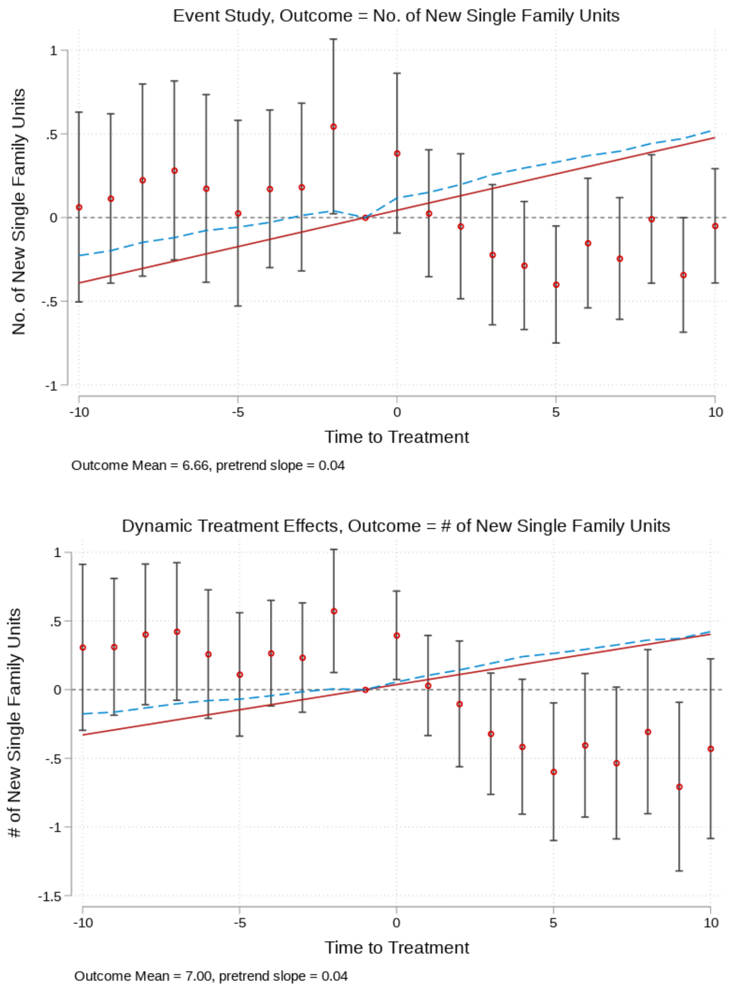


Figure 26: Event Study results (top) and Dynamic Treatment Effect Estimates (bottom) with hypothesized pretrends illustrated.

## 8.3 Treatment v. Control

Each treatment cohort's average outcome is plotted against an average outcome for the same calendar year. In each plot, control cohorts are all cohorts treated 12 or more years after the treatment cohort.

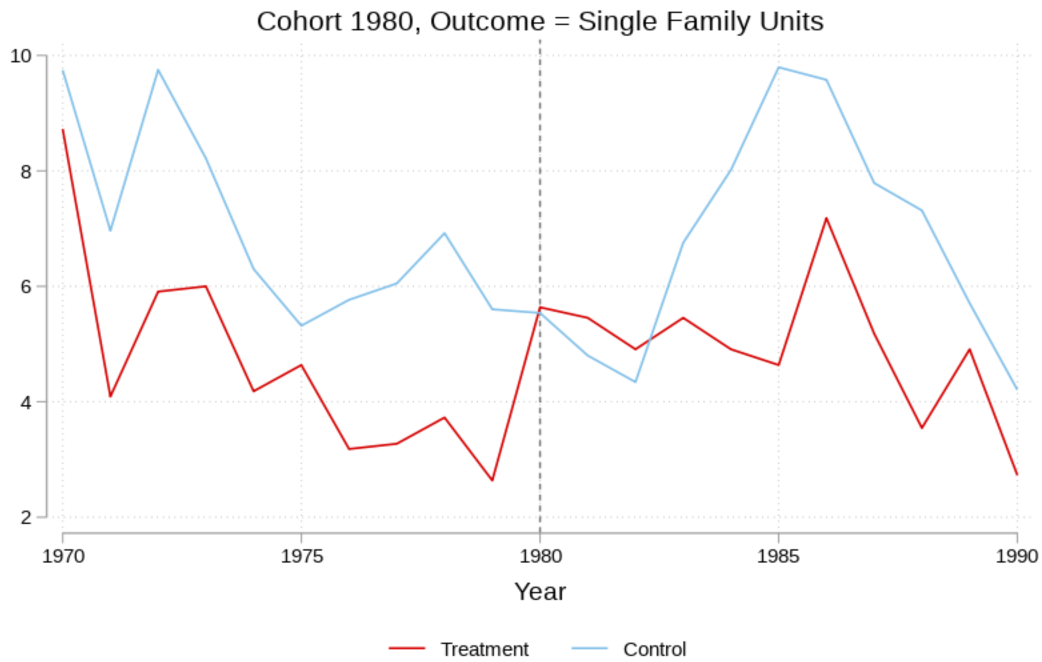


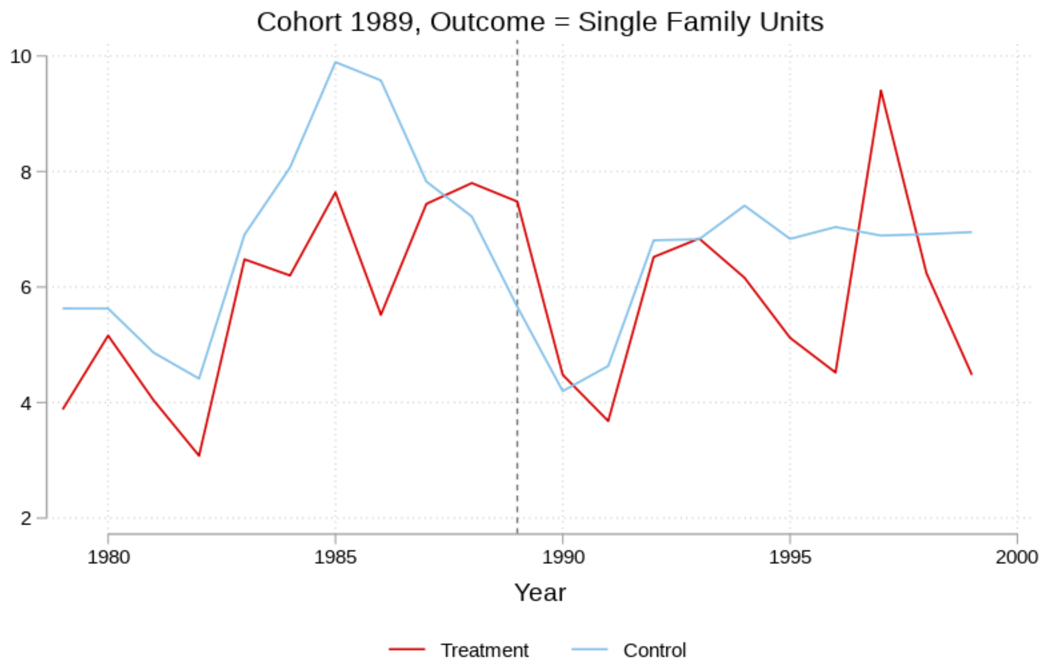

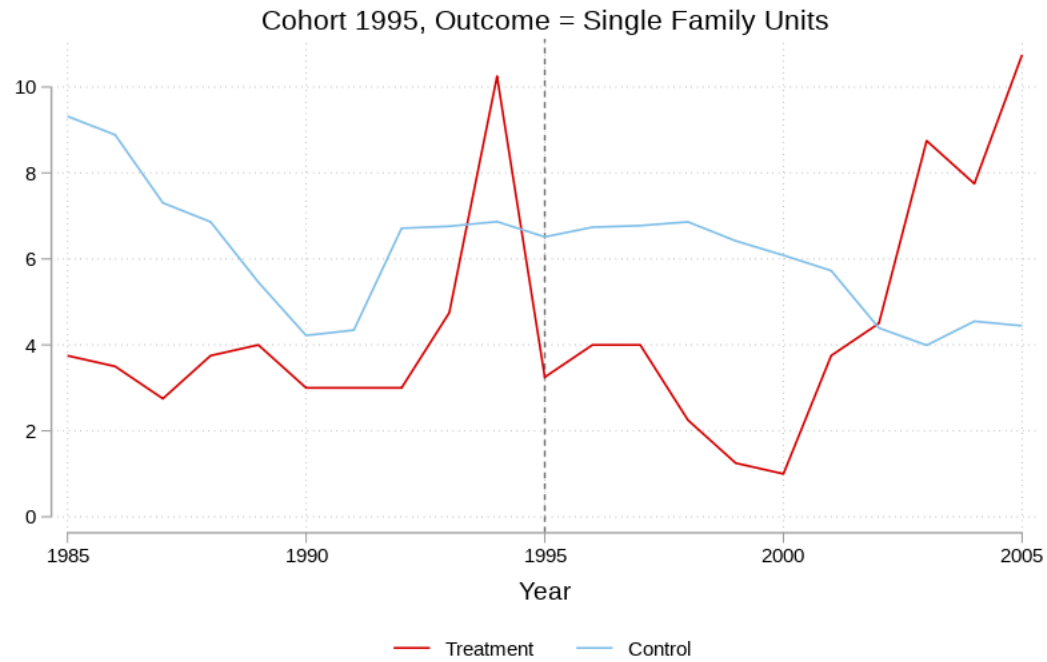

Cohort 1995, Outcome = Single Family Units
10
8
6
4
2
0
1985
1990
1995
2000
2005
Year
Treatment
Control


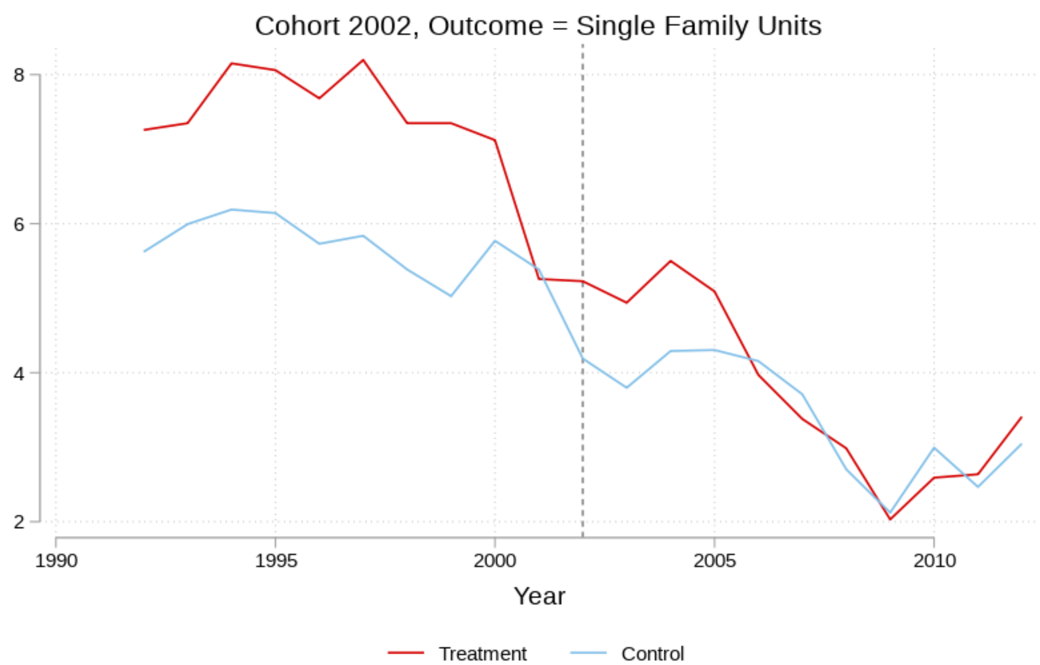

Cohort 2002, Outcome = Single Family Units
8
6
4
2
1990
1995
2000
2005
2010
Year
Treatment
Control

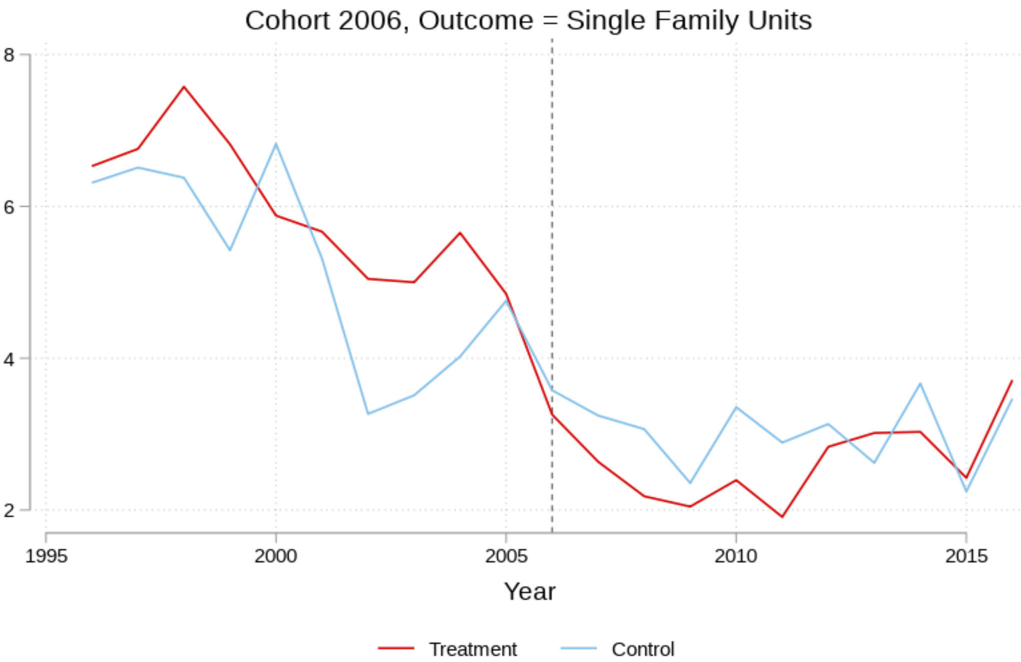


## 8.4 Treatment and Control Group Sizes

| Relative Year | Treatment Buffers | Control Buffers (Unique) | Control Buffers (Repeated) |
|---|---|---|---|
| 0 | 1,000 | 1,006 | 1,010,721 |
| 1 | 982 | 1,006 | 962,526 |
| 2 | 963 | 1,006 | 915,380 |
| 3 | 928 | 1,006 | 869,367 |
| 4 | 903 | 1,006 | 824,656 |
| 5 | 881 | 1,005 | 780,842 |
| 6 | 860 | 1,005 | 738,029 |
| 7 | 841 | 1,005 | 696,628 |
| 8 | 813 | 1,004 | 656,597 |
| 9 | 798 | 1,003 | 617,538 |
| 10 | 767 | 999 | 579,390 |

Table 4: Number of Buffer Areas used for Estimation of Dynamic Treatment Effect at each Relative Year after Treatment

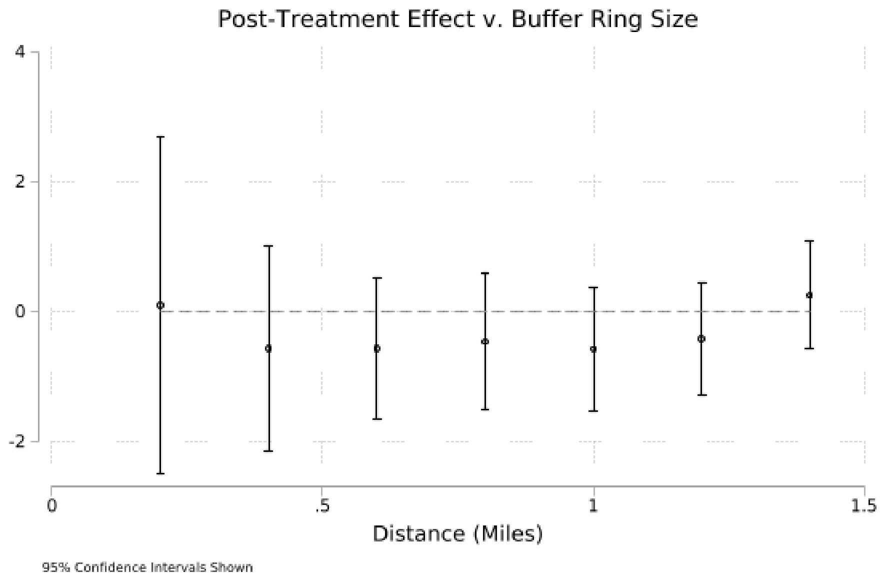


Figure 27: Showing decay of treatment effect at various distances from the 40B development. For each development area ring (i.e., 0-0.2mi, 0.2-0.4mi, 0.4-0.6mi, up to 1.4 mi max), a dynamic treatment effect was estimated on outcomes within that ring. Secondly, so that rings of different sizes are comparable, the outcomes are normalized according to the outcome mean for each ring size.