

MIT Open Access Articles

The Price of Inclusion: Evidence from Housing Developer Behavior

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: 2022. "The Price of Inclusion: Evidence from Housing Developer Behavior." Review of Economics and Statistics.

As Published: https://doi.org/10.1162/rest_a_01231

Persistent URL: <https://hdl.handle.net/1721.1/147008>

Version: Final published version: final published article, as it appeared in a journal, conference proceedings, or other formally published context

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



The Price of Inclusion: Evidence from Housing Developer Behavior

Evan J. Soltas*

June 2, 2022

Abstract

In many cities, incentives and regulations lead developers to integrate low-income housing into market-rate buildings. How cost-effective are these policies? I study take-up of a tax incentive in New York City using a model in which developers trade off between tax savings and pre-tax income. Estimating the model using policy variation and microdata on development from 2003 to 2015, I find a citywide marginal fiscal cost of \$1.6 million per low-income unit. Differences in neighborhoods, not developer incidence, explain the cost premium over other housing programs. Weighing costs against estimates of neighborhood effects, I conclude middle-class neighborhoods offer “opportunity bargains.”

JEL Codes: H21, H22, H32, H71, R28, R31, R38

*MIT Economics. Email: esoltas@mit.edu. I am indebted to David Autor and Jim Poterba for their support and advice. I thank Daron Acemoglu, Alain Bertaud, Ingrid Gould Ellen, Harris Eppsteiner, Amy Finkelstein, Ed Glaeser, Jon Gruber, Andrew Kaplan, Stephen Morris, Dev Patel, Charlie Rafkin, Anna Russo, Tobias Salz, Ben Sprung-Keyser, Bill Wheaton, and conference participants at the NYU Furman Center, St. Anselm College, the 2020 Virtual UEA Conference and PhD Workshop, and YES–Penn for many helpful suggestions. I gratefully acknowledge support from the National Science Foundation Graduate Research Fellowship (Grant No. 1122374) and the Lincoln Institute of Land Policy’s C. Lowell Harriss Dissertation Fellowship.

1 Introduction

Inclusionary housing programs, which encourage or require developers to include low-income units in new market-rate multifamily housing construction, are an increasingly common type of urban policy in the United States. Of the top 100 U.S. cities by population, 33 had an inclusionary housing program as of 2020, 18 of which did not in 2010. A key motive for these policies is social integration: Recipient households usually live in buildings with higher-income tenants and in higher-income neighborhoods than recipients of other forms of housing assistance. Yet the benefits of such policies come at a price. When mandatory, inclusionary housing is an implicit tax on market-rate housing; when voluntary, it requires subsidies or valuable exceptions to zoning regulations to generate participation. How costly is it to induce developers to provide inclusionary housing? What explains cost differences between it and other housing programs? And are the better neighborhoods it obtains for recipients worth their costs?

I build and estimate a model of housing developer behavior to answer these questions. In my model, developers choose whether to enter a new building into a voluntary inclusionary housing program, in which they receive tax benefits for reserving a fraction of units in the building for low-income tenants who pay regulated rents. Developers compare their potential tax savings and forgone rental income, participating if the former exceeds the latter in present value. To estimate the model, I use several sources of policy variation that cause developers of similar buildings to be offered tax benefits of greatly different value. As the tax benefit a developer receives is a function of potentially endogenous building characteristics, such as its number of units, I apply a simulated-instruments approach, as in Currie and Gruber (1996). This approach exploits both policy variation in the tax-benefit formula as well as a subset of lot characteristics, such as lot area and zoning, that are plausibly exogenous to tax policy but that greatly affect building choices. From my estimates, I obtain the distribution of developers' breakeven thresholds, the supply of inclusionary units, and the average marginal fiscal cost per unit.

In my empirical analysis, I evaluate New York City's 421-a property tax exemption for inclusionary housing. With an annual fiscal cost of \$1.6 billion as of 2020, 421-a is New York City's

largest residential property tax expenditure and the largest inclusionary housing program of any U.S. city. From 2003 to 2015, the period I study, about 12,000 onsite inclusionary units were built under 421-a, or 4.7 percent of all exemption-eligible new residential units. Developers of eligible buildings face the same choice as in my model: They may accept the exemption in compensation for setting aside one in five units for low-income tenants, or they may decline the exemption and charge market rents on all units. I estimate the model using the participation decisions of all eligible developments and data on their lot and block characteristics. I use my estimates to characterize developer behavior in the face of 421-a, both on citywide average and for 179 neighborhoods. I also evaluate the cost-effectiveness of 421-a relative to Section 8 vouchers and the Low-Income Housing Tax Credit (LIHTC), analyzing in particular the roles of neighborhood differences and incidence on developers in determining program cost differences. Finally, I use external estimates of neighborhood effects to weigh the fiscal costs of 421-a against the long-run benefits of moving households to higher-opportunity neighborhoods.

There are five main conclusions of my empirical analysis. First, developers respond to incentives. I estimate that increasing the 421-a incentive by one percentage point of building value would increase the take-up rate by 0.59 percentage points. Second, 421-a is costly by comparison to other forms of housing assistance in New York City. On citywide average, I estimate that the fiscal cost of the marginal inclusionary unit is about \$1.6 million, which is about six times the citywide per-unit cost of Section 8 vouchers or the LIHTC. Third, there is immense variation across neighborhoods in developers' breakevens, supply responses, and average marginal fiscal costs per inclusionary unit. For instance, in some Manhattan neighborhoods, the average marginal fiscal cost of an inclusionary unit is as high as \$2 million. By contrast, in many neighborhoods in the Bronx, Queens, and Staten Island, the average marginal fiscal cost is less than \$150,000.

These disparities across neighborhoods also suggest that cost differences between 421-a and other housing assistance programs may simply reflect differences in the geographic distributions of their units. Whereas Section 8 voucher and LIHTC units are concentrated in low-income neighborhoods, inclusionary units are concentrated in higher-income ones. In my fourth empirical contribu-

tion, I adjust for neighborhood differences to conduct a closer comparison of programs. Applying the reweighting method of DiNardo et al. (1996), I estimate how much the average inclusionary unit would cost New York City if, instead, such units had the same geographic distribution as Section 8 and LIHTC units. By this counterfactual, I find that the average cost difference between 421-a and these programs entirely reflects differences in neighborhoods. I also rule out the possibility that differential incidence on developers explains an economically significant share of the 421-a cost premium. Between-program cost differences thus primarily reflect the fundamental trade-off in housing policy between quality and cost per unit.

Should New York City's government be willing to pay such a premium for inclusionary units in higher-rent neighborhoods? In my fifth and final contribution, I estimate the marginal value of public funds (MVPF) of using 421-a to add inclusionary units in each neighborhood. To do so, I combine my neighborhood-specific estimates of the marginal fiscal cost of inclusionary units with estimates of the long-run effects of neighborhoods on children's adult incomes from Opportunity Insights (Chetty et al., 2018). As in Hendren and Sprung-Keyser (2020), the MVPF of 421-a weighs its present-value fiscal cost, net of fiscal externalities, against households' willingness to pay for inclusionary units and the long-run neighborhood effects on their children's after-tax incomes. I detect "opportunity bargains" in many middle-class neighborhoods. Other neighborhoods have low MVPFs, due to high fiscal costs in the highest-income neighborhoods and negligible benefits in the lowest-income neighborhoods. An across-the-board increase in the 421-a incentive would be poorly targeted, producing more inclusionary housing in low-MVPF neighborhoods than in high-MVPF ones. Overall, 421-a appears a mistargeted, but not inefficient, policy to encourage mixed-income housing development.

My results inform debates over inclusionary housing policies in New York City and other cities. 421-a's future is now in question amid legislative gridlock over state property-tax reform, and much controversy surrounds a Mandatory Inclusionary Housing (MIH) policy in select neighborhoods. Proposals for inclusionary housing policies are also under consideration in several mid-size U.S. cities. Meanwhile, several U.S. states have preemption laws that forbid their cities from estab-

lishing inclusionary housing programs. Such policy activity belies the near-total lack of evidence on the impacts of inclusionary housing. Interestingly, the key questions in these debates relate to housing supply under inclusionary policies. For example, the Association for Neighborhood and Housing Development, an advocacy group for housing nonprofit organizations, writes that 421-a “forfeits billions of dollars in public money for minimal public benefit in return,” describing it as “a windfall for real estate developers, with little return for communities.” On the other hand, a former senior official of the Department of City Planning criticizes New York City’s MIH program for undercompensating developers: “MIH reflects a fundamental failure to recognize that inclusionary zoning is always voluntary. No development occurs without the expectation of a threshold rate of return on investment.”¹ My results also inform debates over MIH policies and, in particular, suggest that MIH’s implicit surtax is large relative to typical property tax rates.

This paper contributes to several literatures in public finance and urban economics. It is one of the first to conduct a microeconomic evaluation of developer participation in any type of housing program. DiPasquale (1999) argues that “our understanding of the micro foundations of housing supply” would be best advanced by “bringing new data to bear on the decision-making processes” of developers. Credible evidence on cost-effectiveness is particularly scarce in housing policy: Olsen and Zabel (2015) argue that such analyses should be “the highest priority for housing policy research.” Scholars have previously studied market-level equilibrium impacts of Section 8, the LIHTC, and other housing subsidies (Susin, 2002; Sinai and Waldfogel, 2005; Gibbons and Manning, 2006; Baum-Snow and Marion, 2009; Eriksen and Rosenthal, 2010; Diamond and McQuade, 2019). This paper is also related to research evaluating the impacts of housing policies, especially ones aimed at inclusion and desegregation (Wong, 2013; Chetty et al., 2016; Chyn, 2018; Collinson and Ganong, 2018; Diamond et al., 2019; Favilukis et al., 2019; Van Dijk, 2019; Bergman et al., 2020; Davis et al., 2020). The central role of developer and landlord participation

¹Association for Neighborhood and Housing Development, “421A Developer’s Tax Break,” 2 January 2014. Eric Kober, “De Blasio’s Mandatory Inclusionary Housing Program: What Is Wrong, and How it Can Be Made Right,” Manhattan Institute, January 2020.

in U.S. housing policy since the 1970s suggests many applications of my approach beyond inclusionary housing. Previous research on inclusionary housing is limited to city-level case studies (e.g., Schuetz et al., 2011). A recent paper, Singh (2020), uses reforms to 421-a as a source of variation to study the impact of new housing development on gentrification. Methodologically, my work takes a revealed-preference approach to study selection into voluntary regulation, similar to Anderson and Sallee (2011) on fuel-economy standards, Benzarti (2020) on tax itemization, Kisin and Manela (2016) on bank capital requirements, and Einav et al. (2020) on provider choice of payment systems in public health insurance.

This paper proceeds as follows. Section 2 explains New York City’s 421-a exemption and the broader context. Section 3 explains how I measure the 421-a tax incentive. Section 4 presents a model of developer behavior under voluntary inclusionary housing. Section 5 introduces the data. Section 6 presents simple graphical evidence for developer responses to 421-a. Section 7 shows how I estimate my model of developer behavior, and Section 8 reviews the estimation results. Section 9 evaluates the cost-effectiveness of 421-a. Section 10 concludes.

2 Inclusionary Housing in New York City and the U.S.

In this section, I first provide an overview of inclusionary housing in the context of U.S. housing policy. I then explain relevant details about the housing market in New York City. Finally, I introduce the key features of New York City’s 421-a exemption.

2.1 Inclusionary Housing and U.S. Housing Policy

One third of the 100 most populous U.S. cities had inclusionary housing policies in 2020, as shown in Appendix Figure A1. There is, however, much variation in form across cities. First, inclusionary housing policies may be mandatory or voluntary. Of these 100 cities, 24 have mandatory programs, and 12 have voluntary programs. Second, among voluntary policies, cities differ in the type of incentives used to attract developer participation. Most commonly, cities grant “density bonuses,” which are development rights to exceed allowable floor area on a given lot under zoning regulations, or expedite the permitting process. Third, policies vary in the covered share of new buildings

and in the set-aside share of units inside covered buildings.² The long history of New York City’s 421-a program, its considerable scale, and the explicit cash value of the incentive make it a natural starting point for economic research on inclusionary housing.

There are major differences between inclusionary housing and the two main forms of low-income housing policy in the U.S. (Collinson et al., 2015), Section 8 vouchers and the LIHTC. These differences lie the nature of the buildings and the neighborhoods in which recipient households live. Table 1 shows these differences in New York City by combining several sources of microdata from the U.S. Department of Housing and Urban Development (HUD, see Appendix C). In LIHTC buildings in New York City, nine in ten units are for low-income tenants. By comparison, one in five units are for low-income tenants in 421-a buildings. The surrounding neighborhoods also differ. The median household in an inclusionary unit lives in a Census block group with a median annual income around \$100,000, more than twice that of Section 8 voucher or LIHTC households. The educational and demographic composition of neighborhoods with 421-a, Section 8, and LIHTC units also differs markedly. At both building and neighborhood levels, inclusionary housing is unusual among U.S. housing policies in deconcentrating poverty.

2.2 The Housing Market in New York City

New York City’s housing market is atypical among U.S. cities in several respects that are relevant to my analysis. First, in the extent of government intervention: Of rental units, about 14 percent are social housing (i.e., public or subsidized), and about half are rent-regulated, both much more than in other U.S. cities (Metcalf, 2018). It is thus possible to conduct sensible comparisons within New York City between housing programs. Second, zoning and other land-use regulations are set such that, at the lot level, they are usually a binding constraint on housing supply (Glaeser et al., 2005). As development is almost predetermined, there is little space for responses to tax incentives along margins such as the number of units or type of land use.

²Other margins of variation include the existence of options to build inclusionary units offsite or to pay “in-lieu fees” to be discharged of obligations, the comparability of market-rate and inclusionary units, allowable rents on inclusionary units, and the income range of inclusionary tenants.

Third, New York City’s property tax code exhibits very high statutory tax rates (above 10 percent, see Appendix Figure A4) but on assessed values that are usually small fractions of properties’ market values. On average, effective property tax rates in New York City (expressed as a share of market value) are below rates in most U.S. cities for owner-occupied housing but are near the U.S. average for rental housing (Lincoln Institute of Land Policy and Minnesota Center for Fiscal Excellence, 2019). The existence of high statutory rates and often-arbitrary assessment practices, however, generates variation between buildings in effective tax rates and sustains demand from developers for tax relief. In particular, 421-a is one of several tax incentives for property investment in New York City. An implication of these alternatives is that the budgetary cost of 421-a overstates the additional revenue New York City would raise if 421-a were abolished, as some developers who take up 421-a would take up other exemptions in its absence. The existence of alternative tax incentives also creates challenges for my analysis that I discuss in Sections 3 and 8.

2.3 The 421-a Exemption

Section 421-a of the New York State Real Property Tax Law codifies an “exemption of new multiple dwellings from local taxation.” Here I discuss its key features and history from its introduction in 1971, through the period I study (2003–2015, coinciding approximately with Michael Bloomberg’s three terms as mayor), up to its latest iteration following a reform in 2015. In Section 3, I describe in detail the sources of policy variation I use in estimation. In both sections, I draw primarily on the text of the statute as well as on administrative rules.

421-a is a multi-year partial exemption from property taxation that is primarily for multifamily residential developments with “inclusionary” units—that is, units for low-income tenants who pay regulated below-market rents. The exemption applies only to the building (i.e., non-land) component of property value, thereby transforming the property tax into a de-facto land value tax. It may last between 10 and 25 years after construction. In some neighborhoods, residential buildings are eligible for a shorter “as-of-right” exemption even if they do not provide any inclusionary housing. In 1985, New York City introduced its requirement that 421-a buildings include low-income units, a stipulation at first applying only to midtown Manhattan.

The 421-a exemption is the largest residential property tax expenditure in New York City. Since the 1980s, about one in three new units received a 421-a exemption, as shown in Appendix Figure A2. The fiscal cost of 421-a is substantial: In total, 421-a exempts from taxation about 15 percent of the assessed value of all multifamily residential property in New York City. With the continued rise of property values in New York City over time, the tax expenditure on 421-a has grown about twenty-fold in inflation-adjusted terms since 1999. Appendix Figure A17 shows that 421-a's cost greatly exceeds that of other tax expenditures on residential property in New York City.

To receive the most generous exemption, a developer must reserve one in five units for inclusionary tenants. This fraction is fixed; developers cannot reserve more or fewer units than one in five to adjust the value of the exemption. In each year, the maximum rent on inclusionary units is fixed citywide relative to the Area Median Income (AMI) for New York City, as determined by HUD. Inclusionary unit rents must be less than 30 percent of income to households making 60 percent of the New York City AMI. In 2019, the AMI for a family of three was \$96,100, and so their maximum monthly rent was \$1,442. Whereas these rent caps apply uniformly across New York City, market rents vary markedly, creating variation in the “replacement rate” of the regulated rent for the market rent. Inclusionary tenants are selected through a non-market allocation process jointly run by the city government and the property manager.

A new multifamily dwelling is eligible for 421-a if it has a sufficient number of units and is built on an eligible lot. Before December 2007, the minimum number of units was three, after which the minimum rose to four units. A lot is eligible if it was vacant or deemed “underutilized” for three years before construction. Residential buildings are eligible whether their market-rate units are rentals or owner-occupied as a cooperative or condominium.

Until December 2007, developers could choose whether to build inclusionary units onsite or offsite. If offsite, they received a shorter tax exemption. From 2008 onward, the offsite option was abolished. The focus of this paper is onsite inclusionary housing, but I analyze the offsite option in Appendix D. Developers have an obvious incentive to minimize the value of onsite units they set aside for low-income tenants. The comparability of onsite inclusionary units to market-rate units

in the same building is, accordingly, a focus of regulation and enforcement.³

In January 2016, beyond my data, the 421-a program temporarily closed to new development. It returned in April 2017 as the Affordable New York Housing Program with several notable reforms that applied retroactively to developments that filed an initial permit after December 2015.⁴ Under the new policy, developers may select among various combinations of a set-aside share and tenant income bands for inclusionary units.

3 Measuring the 421-a Incentive at the Building Level

This section sets out my definition of the value of tax savings from 421-a. It then discusses the key sources of policy variation in developers' incentives to participate in 421-a, as determined by a microsimulation model of New York City's property tax code (see Appendix C).

3.1 Definition of Tax Savings

The ideal economic measure of a developer's incentive to provide onsite inclusionary housing under 421-a is the difference in the present values of its expected after-tax profits between accepting 421-a and its next-best alternative option. There are two conceptual challenges in implementing this measure in the data. First, I do not know the developer's forecast for their building's market value over time, which determines assessed values and thereby tax payments. Second, I do not know the developer's entire choice set, including all combinations of tax exemptions and buildings that could have been built on their lot, from which it finds its next-best alternative. To make

³Since July 2008, developers are required to have a "proportional" mix of inclusionary and market-rate units with respect to the number of bedrooms, spread across floors, with equal access to building amenities. Many unit characteristics, from views to appliances, are also regulated. Developers have nevertheless tried to game these rules, infamously in the 2014 "poor door" scandal.

⁴The reform also allowed buildings where construction began before December 2015, but which had not yet received any 421-a benefit, to opt into a more generous exemption. The developers who accepted this option appear to have already decided to accept 421-a in its prior form. I treat the 16 affected buildings as accepting 25-year exemptions.

progress, I specify developers' expectations and choice sets.

For each building i , I define the 421-a tax savings as the difference in its present values of tax liabilities between providing and not providing onsite inclusionary housing, expressed as a share of i 's total market value and holding fixed all other building characteristics. Let $\tau_{i,s}^1$ and $\tau_{i,s}^0$ be i 's potential tax rates in year s , respectively conditional upon participating and not participating. The difference in present values is approximately

$$\Delta\text{PDV}_{i,0} \approx E_0 \left[\sum_{s=0}^{\infty} \beta^s (\tau_{i,s}^0 - \tau_{i,s}^1) V_{i,s} \right],$$

where $E_s[\cdot]$ is expectation in year s and $V_{i,s}$ is i 's market value in s .⁵ To simplify this expression, I make three further assumptions. First, market values reflect expected paths of future rents: $V_{i,t} = E_t [\sum_{s=t}^{\infty} \beta^s (1 - \tau_s) R_{i,s}]$. Second, the expected path of future rents is $E_t [R_{i,t}] = R_{i,t} (1 + r)^{s-t}$. Dividing $\Delta\text{PDV}_{i,0}$ by $V_{i,0}$, substituting for $V_{i,s}$, and letting $\rho = [\beta(1 + r)]^{-1} - 1$ be the capitalization rate, I obtain my measure of tax savings:

$$\Delta\tau_i = \frac{\Delta\text{PDV}_{i,0}}{V_{i,0}} = E_0 \left[\sum_{s=0}^{\infty} \frac{\tau_{i,s}^0 - \tau_{i,s}^1}{(1 + \rho)^s} \right],$$

which is the sum over time of differences in property tax rates between participating and not participating in 421-a, discounting by the capitalization rate. I set $\rho = 0.05$, consistent with industry surveys (e.g., by the CBRE Group) of multifamily housing developers in New York City.

3.2 Sources of Variation in the 421-a Exemption

In Appendix C, I explain how New York City determines property tax bills, and I replicate their calculations in a microsimulation model of the property tax code. Here I review the policy variation in 421-a, as embedded in my microsimulation model, that occurs due to features of 421-a itself and indirectly due to interactions of 421-a with the rest of the property tax code.

The duration in years of the 421-a exemption depends upon whether the development provides inclusionary housing, whether the inclusionary housing is onsite or offsite, the year of filing the

⁵This approximation ignores that a lower tax rate may raise the building's market value, indirectly raising tax payments.

initial construction permit, and the region of New York City in which the development is located (see Appendix Table A2). In summary, developments receive longer exemptions when they provide onsite inclusionary housing, and the incentive to provide inclusionary housing is considerably stronger in some neighborhoods of New York than in others.

Several sources of geographic variation exist in the incentive for inclusionary units. First, there is the “Geographic Exclusion Area” (GEA), which initially covered midtown Manhattan. Development in the GEA faces a larger incentive for inclusionary units than development outside the GEA. The GEA has also changed its boundary over time. In May 2005, some of the Greenpoint–Williamsburg neighborhood of Brooklyn was added to the GEA as part of a rezoning. In July 2008, reforms to 421-a expanded the GEA region to cover all Manhattan as well as several other neighborhoods, primarily parts of Brooklyn. Appendix Figure A5 maps the GEA over time.

Second, as part of the Neighborhood Preservation Program (NPP)—a now-obscure policy established in 1973 to prevent urban decay—buildings in areas of Brooklyn, the Bronx, and Queens were eligible until 2008 for the longest-duration 421-a exemption “as of right,” meaning they faced no tax incentive to provide inclusionary housing. Appendix Figure A6 maps the NPP areas.

The same 2008 law also limited as-of-right 421-a exemptions citywide in buildings with more than four units to the first \$65,000 of assessed value per unit, a reform known as the “AV cap.” The introduction of the AV cap therefore strengthened incentives to provide inclusionary housing for more-expensive buildings relative to less-expensive buildings.

Policy variation in 421-a also emerges indirectly through assessment practices. Exemptions are less valuable to developers who anticipate the assessed value of their building will be lower relative to the building’s true market value. In particular, a bias in the valuation of condominiums and limits on the growth of assessed values, discussed in Appendix C, mean that 421-a predictably reduces some buildings’ effective tax rates, defined with respect to market value, more than those of other buildings. Furthermore, the AV cap is less likely to bind for underassessed buildings, further reducing their 421-a incentive.

To rule out the possibility my results are driven by other policies in New York City, not 421-a,

my analysis also controls for several smaller-scale housing policies of which I am aware and that plausibly interact with 421-a. I discuss these in Appendix C.

4 Tax Incentives and Housing Developer Behavior

This section develops a model of a profit-maximizing developer who chooses whether to participate in a voluntary inclusionary housing program. In particular, I show their participation choice is governed by a trade-off between the present values of tax savings and pre-tax rental income.

In return for setting aside a share λ of units for low-income tenants who pay below-market rents, participating developers pay a reduced property tax rate. I let $\Delta\tau_i \geq 0$ denote the present value of tax savings, expressed as a share of building i 's value. I take the building as fixed and assume away any intensive margin response in housing supply, motivated by the stringency of zoning regulations in much of New York City. I also treat the housing market as perfectly competitive, so that conditional on the set-aside share λ , developers take rents as given. Appendix B presents a model that relaxes both assumptions. Taxes are levied in proportion to rental income.

Inclusionary tenants pay rents to the developer that are fixed at r_t citywide in year t . In market-rate units, the rent $m_{i,t}(\lambda)$ is a function of building characteristics and the set-aside share. If $\partial m_{i,t}(\lambda)/\partial\lambda < 0$, then setting aside units for inclusionary tenants reduces the willingness to pay of market-rate tenants. Putting these together, the building average rent in year t is:

$$p_{i,t}(\lambda) = \lambda r_t + (1 - \lambda)m_{i,t}(\lambda).$$

Letting building i 's pre-tax rental income in year t be $R_{i,t}(\lambda) = N_i p_{i,t}(\lambda)$, where N_i is its number of residential units, the present value of expected after-tax rental income is

$$\pi_i(\lambda) = E_0 \left[\sum_{s=0}^{\infty} \beta^s (1 - \tau_s(\lambda)) R_{i,s}(\lambda) \right].$$

To analyze the impact of the inclusionary units, it will be convenient to define the log-difference in the present-value average rent between the participation and nonparticipation options:

$$\Delta \log p_i(\lambda) \approx \sum_{s=0}^S \frac{p_{i,s}(\lambda) - p_{i,s}(0)}{(1 + \rho)^s p_{i,s}(0)} = -[\lambda \mu_i + (1 - \lambda) \delta_i(\lambda)], \quad (1)$$

where I further define

$$\mu_i = \sum_{s=0}^{\infty} \frac{m_{i,s}(0) - r_s}{(1 + \rho)^s m_{i,s}(0)} \quad \text{and} \quad \delta_i(\lambda) = \sum_{s=0}^{\infty} \frac{m_{i,s}(0) - m_{i,s}(\lambda)}{(1 + \rho)^s m_{i,s}(0)}.$$

The term μ_i reflects the inclusionary discount: It is the difference in rent between an inclusionary unit and a market-rate unit in the nonparticipation counterfactual. The term $\delta_i(\lambda)$ reflects the disamenity: It is the discount on rent that the developer offers on market-rate units to compensate for the presence of inclusionary units. Weighted by the set-aside share, these terms yield the forgone rental income, expressed as a share of potential rental income.

The developer participates in the inclusionary housing program if its potential after-tax rental income under participation exceeds its potential after-tax rental income under nonparticipation:

$$D_i = 1[\Delta\pi_i(\lambda) \geq 0],$$

where $\Delta\pi_i(\lambda) = \pi_i(\lambda) - \pi_i(0)$ denotes the difference in after-tax rental income for i . As market rents are fixed conditional on λ and the developer has no intensive supply margin, the log-difference in after-tax rental income is

$$\Delta \log \pi_i(\lambda) = \Delta \tau_i + \Delta \log p_i(\lambda), \tag{2}$$

reflecting the offsetting forces of a lower property tax liability and a lower pre-tax rental income. As $\Delta\pi_i(\lambda) \geq 0$ if and only if $\Delta \log \pi_i(\lambda) \geq 0$, I obtain the program participation condition:

$$\Delta \tau_i \geq \lambda \mu_i + (1 - \lambda) \delta_i(\lambda), \tag{3}$$

using the expression for $\Delta \log p_i(\lambda)$ in Equation 13. In deciding whether to participate in an inclusionary housing program, a developer compares the tax savings (the left-hand side of Equation 3) and the impact of the lower average rent on its pre-tax rental income (the right-hand side of Equation 3). Forgone rental income is greater when the set-aside share λ , the inclusionary discount μ_i , or the disamenity $\delta_i(\lambda)$ is larger.

5 Data

I use several public data sources on property taxation and land use in New York City. Together, my data cover the universe of new development that was eligible for 421-a and allow me to calculate, for each such development, the potential property tax savings under 421-a from building onsite inclusionary units, whether or not the developer in fact chose to participate. Observing both the potential tax savings at the individual building level and developers' actual decisions allows me to estimate the supply of inclusionary units. For technical details about the data, see Appendix C.

5.1 Sources

Building Permits. I identify lots with new development from 2003 to 2015 using the Department of Buildings permit database, in which each construction project is assigned a Building Identification Number (BIN). To define the sample period, I use the issuance date of the initial permit for a new building at each BIN, which is also the date at which 421-a eligibility is determined.

Lot and Building Characteristics. I use the Department of City Planning's Primary Land Use Tax Lot Output (PLUTO) database, which contains detailed information on lot and building characteristics. All rental buildings, as well as all individual owner-occupied units, are uniquely identified in New York City records by a Borough–Block–Lot (BBL) code. Since developers decide whether to take 421-a in common across units in a building, not independently by unit, I collapse owner-occupied unit data up to the building level.

Property Tax Assessments. I use two datasets on assessments. First, I scraped assessment documents, in PDFs formatted exactly as sent to taxpayers, from the Department of Finance website. I parsed these assessment PDFs to obtain administrative codes which identify the type of any 421-a tax exemptions. This information allows me to determine whether the developer built inclusionary housing offsite, onsite, or not at all. See Appendix Figure A7 for an example. Second, I use the 2010–2019 Real Property Assessment Databases of the Department of Finance, which contain the universe of assessed values, to forecast the dynamic impacts of assessment growth caps.

Block Characteristics. From public-use tabulations of 2010 U.S. Census and the 2013–2017 American Community Survey (ACS), I collect characteristics of the immediate area surrounding the development at the block or block-group level. From the Census at the block level, these are the median age of the resident population, shares of the resident population by race/ethnicity (non-Hispanic white, non-Hispanic black, Hispanic, Asian), the share of households that rent rather than own their homes, and the share of housing units that are vacant. From the ACS at the block-group level, I observe median household income, the poverty rate, educational attainment shares (less than high school, high school graduate, some college, four-year college graduate, more than four years of college), and commuting mode shares (car, bus, subway, walk, other). For blocks with missing data, I impute values from either the block-group or tract level.

Throughout this paper, I use a statistical definition of neighborhoods from the Department of City Planning: the Neighborhood Tabulation Area (NTA), which aggregates Census tracts to 195 distinct areas, 179 of which contain development in my sample. These NTAs are intended to reflect coherent neighborhoods. For example, separate NTAs exist for the East Village, the Lower East Side, and Chinatown. The NTA is always the level at which I cluster standard errors.

5.2 Summary Statistics

Table 2 reports that, from 2003 to 2015, New York City added 261,572 housing units in 12,146 421-a eligible new buildings. In 2019, the total market value of these buildings, as estimated by the Department of Finance, amounted to \$210 billion. Of these, 581 buildings, representing 4.7 percent of buildings and 22.7 percent of residential units, chose to provide onsite inclusionary housing. These figures imply that around 11,878 onsite inclusionary units were developed under 421-a, at an annualized loss of tax revenue of \$439 million, or \$36,920 per unit per year.

Developers who faced stronger tax incentives to provide onsite inclusionary housing were more likely to provide it. On average, developers who chose to provide onsite inclusionary housing received a tax exemption worth 33 percent of building value for doing so, as compared to an exemption worth 5 percent of building value that developers who declined to provide such housing would have received. On a per-inclusionary-unit basis, developers who did not provide onsite inclusion-

ary housing forewent an average tax savings of \$126,610, whereas developers who provided such housing saved \$443,782 on average. In present value, 421-a buildings pay an average tax rate about 0.16 percentage points below that of buildings who decline 421-a. The buildings where onsite inclusionary housing was provided differ in their observable characteristics from eligible buildings where developers chose not to provide onsite inclusionary housing.

6 Simple Evidence of Developer Responses to 421-a

This section presents graphical evidence that developers are more likely to build onsite inclusionary units if offered more valuable tax exemptions for doing so. I focus on two reforms to 421-a that I describe in Section 3: the GEA expansion and the end of as-of-right 421-a in NPP areas. In later sections, I use my measure of the 421-a incentive, $\Delta\tau_i$, to estimate a model of developer behavior.

In Figure 1, I show how these reforms changed 421-a incentives in different regions of New York City and, side-by-side, the attendant changes in participation rates over time in each region. The panels in the first row focus on the GEA expansion. The left panel shows that, among developments in the year they were issued their initial permit, the average 421-a incentive offered in the GEA expansion region increases sharply in 2008, when this reform occurs. By contrast, increases in the incentive in 2008 are smaller for two comparison regions: lots always and never in the GEA. The right panel shows that, in the GEA expansion region, the 421-a participation rate rises relative to both comparison regions after 2008.

In the second row, I offer the same presentation for the NPP reform. I compare lots in NPP areas to lots that are not themselves in an NPP area but are in Census tracts that partially overlap with an NPP area. The left panel shows that, before 2008, developers in NPP areas have no incentive to provide onsite inclusionary units, whereas developers just outside of an NPP area do have incentives before 2008. After 2008, developers in NPP areas face similar incentives to those just outside of an NPP area. Upon the reform, the 421-a participation rate in NPP areas rises from zero to about 15 percent, the participation rate on lots just outside of an NPP area.

The figure can be used to estimate the developer supply response to 421-a by taking the ratio of a participation-rate change and an incentive change. Relative to the never-GEA region, the

average 421-a incentive and developer participation rate respectively increase by about 50 percent of building value and 30 p.p. in the GEA expansion region. This would imply a local supply response of about 0.6 ($= 0.3/0.5$). Similarly, relative to non-NPP areas in the same Census tract, the average incentive and participation rate in NPP areas rise by about 10 percent of building value and 5 p.p. respectively, also implying a local supply response of about 0.5. Reassuringly, both numbers are close to my estimate of the developer supply response in Section 8. In Appendix D, I redo these analyses as event-study regressions and find similar results.

7 Estimating a Model of 421-a Participation Choice

This section introduces my approach to estimating the model of 421-a participation. First, I present a baseline that takes building, lot, and neighborhood characteristics as exogenous to the participation decision. Second, given the potential endogeneity of building characteristics to the value of 421-a, I also present a simulated-instruments approach that requires that some subset of lot and block characteristics are exogenous to tax policy but determine the building characteristics that determine the value of 421-a. Third, I report estimates under both approaches.

7.1 Exogenous Building Characteristics

Following Equation 3, I assume that a building i 's share of rent that it would forego if it provides onsite inclusionary units is well approximated by observable characteristics x_i of the lot and block. I also allow for unexplained cost components $e_{i,1}$ and $e_{i,0}$ of providing and not providing inclusionary units, which I assume are both distributed i.i.d. Type I extreme value with dispersion parameter σ . Writing the difference of these components as $\Delta e_i = e_{i,1} - e_{i,0}$, the log-difference in developer i 's after-tax rental income $\Delta \log \pi_i$ between providing onsite inclusionary housing under 421-a and not participating in 421-a is

$$\Delta \log \pi_i = \Delta \tau_i + x_i \beta + \Delta e_i. \quad (4)$$

Under the distributional assumption on the error terms, the probability that building i participates in 421-a, conditional on its characteristics x_i , is a logit:

$$\Pr(D_i = 1 | x_i) = \frac{\exp[(\Delta\tau_i + x_i\beta)/\sigma]}{1 + \exp[(\Delta\tau_i + x_i\beta)/\sigma]}, \quad (5)$$

suppressing the conditioning on the tax incentive $\Delta\tau_i$. I estimate the logit parameters (σ, β) by maximum likelihood, clustering standard errors by Neighborhood Tabulation Area.

7.2 Endogenous Building Characteristics

Developers may choose what type of building to build with the possibility of receiving the 421-a exemption in mind. Such simultaneity is likely to bias upward my estimates of the responsiveness of participation to the incentive. For example, consider a lot which is relatively overassessed, making 421-a relatively more attractive. All else equal, a developer should be relatively more likely to build a large rental-unit building on such a lot. This is because such buildings typically face higher tax rates than smaller rental buildings or condominiums, and so the high value of 421-a on this lot gives such a building a comparative advantage relative to other possible buildings on this lot or similar buildings on other lots. The value of 421-a will thus be larger precisely because the developer already intended to take 421-a when designing the building.

I address this endogeneity problem using an extension of the simulated-instruments approach of Currie and Gruber (1996). In my application, I simulate potential buildings on each lot by sampling from the empirical distribution of buildings on other lots, calculating the value of 421-a in each simulation, and instrumenting for the actual value of 421-a using the mean simulated value of 421-a. To increase statistical power, I restrict the set of simulated potential buildings on a given lot to those in the same borough and which would be “feasible” on that lot. Here as in its original application, the simulated-instruments approach can be viewed as a way of parametrizing exogenous policy variation while purging endogenous characteristics.

Formally, I estimate the logit specification in Equation 5 by generalized method of moments.

My moment conditions take the form:

$$\mathbb{E} \left[\left\{ D_i - \Lambda \left(\frac{\Delta\tau_i + x_i\beta}{\sigma} \right) \right\} z_i \right] = 0, \quad (6)$$

where $\Lambda(\cdot) = \exp(\cdot)/[1 + \exp(\cdot)]$ and the instrument vector is $z_i = (\widetilde{\Delta\tau}_i, x_i)$. The simulated instrument $\widetilde{\Delta\tau}_i$ is defined as the unweighted mean of the value of 421-a for lot i over all simulated buildings j in \mathcal{F}_i , the set of all other buildings that would have been feasible on i 's lot:

$$\widetilde{\Delta\tau}_i = \frac{1}{|\mathcal{F}_i|} \sum_{j \in \mathcal{F}_i} \Delta\tau(b_j, x_i),$$

where $|\mathcal{F}_i|$ is the cardinality of i 's feasible set, b_j is j 's vector of endogenous building characteristics, and $\Delta\tau(b_j, x_i)$ is obtained using my tax microsimulation model from Section 3. As above, I cluster standard errors by Neighborhood Tabulation Area.

I say a building j is “feasible” on a lot i , as denoted by $j \in \mathcal{F}_i$, if the pair (i, j) passes two tests. First, j 's lot area must be within 20 percent of i 's lot area.⁶ Second, I restrict to lots in the same zoning class as i . In New York City, zoning classes are primarily defined with respect to the maximum allowable floor area ratio—that is, floor area divided by lot area—and range from areas that only allow detached single-family residences to areas that allow skyscrapers. For a typical lot, these restrictions eliminate about 80 percent of all other lots in the same borough. For further detail on the simulated instrument, including a case study showing how it purges endogenous variation in building choice, see Appendix C.

Is the simulated instrument $\widetilde{\Delta\tau}_i$ plausibly valid and informative for $\Delta\tau_i$? Its validity rests on two assumptions: Neither the 421-a exemption formula itself, nor the subset of lot and block characteristics used to predict building characteristics, is endogenous to the participation decision. I argue that both assumptions are reasonable. As I document in Sections 2 and 3, the evolution of 421-a, along with that of the broader property tax code, reflects citywide political and fiscal considerations, not efforts to rig the formula to benefit specific developments. Furthermore, the dimensions and zoning of the lot are essentially predetermined at the time of construction.⁷ In Section 8, I

⁶Appendix Table A13 shows that my results are robust to changing this 20-percent threshold.

⁷No change in zoning occurred between 2002 and 2019 for 59 percent of lots in the data.

show that $\widetilde{\Delta\tau}_i$ is highly informative for $\Delta\tau_i$. This should be unsurprising. First, there is substantial variation in the tax-rate formula over time and geography. Second, the physical dimensions of a lot, along with zoning regulations, constrain the feasible set of buildings. These constraints typically bind. Insofar as the 421-a exemption formula varies with building characteristics that are likely constrained by lot size and zoning—such as the number of units—such restrictions make the $\Delta\tau_i$ of buildings on similar lots a highly relevant predictor of a building’s own $\Delta\tau_i$.

7.3 Do Developers Respond to 421-a?

Table 3 presents estimates of Equation 5 with various sets of controls, as well as with fixed effects for boroughs or neighborhoods. The fixed effects and controls help to isolate policy variation by removing static differences among locations that may otherwise be a source of confounding relationships. As Section 6 documents in two examples, there is considerable policy variation in 421-a tax incentives across locations over time. The controls are introduced in Section 5. I report the estimated fixed effects and control-variable coefficients in Appendix Tables A7 and A8.

The results suggest that, holding lot and block characteristics fixed, developers are more likely to participate in 421-a when offered more in tax savings. In particular, an increase in the 421-a tax savings of one percent of building value would increase a building’s participation probability by 0.18 to 0.26 percentage points. For comparison, I also report linear probability model estimates in Appendix Table A4. In the rest of the paper, I take the specification in Column 4 of Table 3, with fixed effects for borough and year as well as lot and block controls, as the baseline specification.

These reduced-form estimates also have a structural interpretation. The coefficient in Table 3 is not only a behavioral semi-elasticity but also $1/\sigma$, a measure of the dispersion of unobservable profit shocks Δe_i in my model.⁸ My results imply that the standard deviation of unobservable profit shocks Δe_i is on the order of 30 percent of market value. Furthermore, the estimated coefficients

Almost all changes in zoning are part of neighborhood-wide rezonings. Individual-lot rezonings are rare due to New York City’s laborious review process for any change in zoning.

⁸The parameter σ determines the shape of the Type I extreme value distribution for e_i . Using distributional properties, the standard deviation of Δe_i is $\sigma\pi/\sqrt{3}$.

$\hat{\beta}$ in Appendix Tables A7 and A8 can be interpreted as descriptively characterizing variation in forgone rental income. I also check the goodness-of-fit of my model using the two policy reforms examined in Section 6. In particular, Appendix Figure A8 confirms that the predicted probabilities of 421-a participation evolve similarly to the actual probabilities around these reforms. The reforms occur in the sample but do not constitute all of the variation in the incentive $\Delta\tau_i$.

7.4 Allowing for Endogenous Building Characteristics

In Figure 2, I present binned scatterplots for the “zeroth” and first stages of the simulated instruments GMM approach I develop in Section 7.2. In the zeroth stage, I confirm that buildings resemble, on endogenous building characteristics, the buildings to which they are matched using the subset of lot characteristics. In particular, the matching approach finds comparison lots whose buildings have quite similar numbers of residential units, floor area, floor area ratios, condominium shares of residential units, and total assessed values. Developers appear to have relatively little scope to respond to 421-a on these margins of building characteristics.

The rightmost lower panel of Figure 2 is a binned scatterplot of the first stage. The 421-a tax savings from the matched buildings is lower than the savings from the actual building, especially at the high end of tax savings. What explains this discrepancy? Tax exemptions are substitutes. If a developer does not accept 421-a, it has a financial incentive to take up another exemption. This will reduce both its tax liability and its potential savings if it also were to pursue 421-a. For instance, a developer could instead use the building for university faculty housing, which is also tax-exempt, leaving no further savings from providing onsite inclusionary housing. Tax planning, rather than construction, seems to be the primary margin of adjustment to 421-a.

In my simulated-instrument GMM approach, I treat other exemptions as endogenous building attributes. Appendix Table A10 repeats the format of Table 3, but using simulated-instrument GMM, and finds similar coefficient estimates. However, this approach does not address mismeasurement of the value of the tax incentive among participants due to alternative tax exemptions. To evaluate the likely direction of bias, I incrementally scale down the value of the 421-a incentive among participants as a way to approximate substitution to these alternatives. These results sug-

gest that this bias causes me to overstate developer responsiveness to 421-a. By implication, the measurement bias is toward understating, rather than overstating, 421-a’s fiscal cost per unit.

8 The Supply and Fiscal Cost of Inclusionary Housing

In this section, I define and estimate the key economic objects of interest. These are the distribution of developer breakevens, supply responses to changes in the 421-a incentive, and fiscal costs per inclusionary unit, on the margin and on average.⁹

8.1 Definitions

Breakevens. A developer’s breakeven $\Delta\tau_i^*$ is the amount of tax savings per inclusionary unit that would make it indifferent in expectation between having building i participate versus not participate in 421-a. From Equation 4, this breakeven occurs at

$$E[\Delta\pi_i|x_i] = \Delta\tau_i + x_i\widehat{\beta} = 0 \implies \Delta\tau_i^*(x_i) = -x_i\widehat{\beta}.$$

To obtain breakevens in dollar terms, I multiply by the market value per inclusionary unit:

$$\text{Breakeven}_i = -(v_i/\lambda) \cdot x_i\widehat{\beta}, \quad (7)$$

where λ is the set-aside share and v_i is the average market value per unit in building i .

Supply Response. From the logit functional form, I obtain an expression for the supply response of inclusionary units in a neighborhood n to a change in the tax incentive $\Delta\tau_i$. This expression is

$$\eta_n = \int \frac{\partial \Pr(D_i = 1|x_i)}{\partial(\Delta\tau_i)} dF_n(x_i) = \frac{1}{\sigma} \int \Pr(D_i = 1|x_i) [1 - \Pr(D_i = 1|x_i)] dF_n(x_i), \quad (8)$$

where $F_n(x_i)$ is the distribution of buildings with characteristics x_i in n . Neighborhood-level supply responses crucially reflect both dispersion σ in building-level unobservables as well as observable

⁹Appendix D includes supplementary empirical analyses: (1) I estimate the entire supply curve for inclusionary housing; (2) I estimate the effects of changing the set-aside share on the total production of inclusionary housing; (3) I investigate developer heterogeneity and sorting; and (4) I estimate a model with developer choice between no, offsite, and onsite inclusionary housing.

heterogeneity among buildings, as implied by variation in $\Pr(D_i = 1|x_i)$. In implementation, I weight by the number of residential units in each building.

Marginal and Average Fiscal Costs. For a neighborhood n , I measure the average marginal fiscal cost of inclusionary units, given a small change in the 421-a incentive $\Delta\tau_i$, as the ratio of the resulting total change in the 421-a tax expenditure in n to the resulting total change in the number of inclusionary units in n :

$$\begin{aligned} \text{MFC}_n &= \frac{1}{\lambda} \int \frac{\partial [v_i \Delta\tau_i \Pr(D_i = 1|x_i)]}{\partial (\Delta\tau_i)} dF_n(x_i) \bigg/ \int \frac{\partial \Pr(D_i = 1|x_i)}{\partial (\Delta\tau_i)} dF_n(x_i) \\ &= \frac{1}{\lambda} \int \left[\frac{v_i \Pr(D_i = 1|x_i)}{\eta_n} + \omega_i v_i \Delta\tau_i \right] dF_n(x_i) \end{aligned} \quad (9)$$

The second line of Equation 9 shows that the average marginal fiscal cost MFC_n is the sum of two terms, which respectively reflect the contributions of inframarginal and marginal units. First, the contribution of inframarginal units to cost depends inversely upon the supply response, η_n . As η_n increases, $\Delta\tau_i$ must increase less to obtain one additional inclusionary unit, and so less tax revenue is lost to inframarginal units. Second, the contribution of marginal units to cost is a weighted average of the 421-a tax savings of all buildings in neighborhood n , with weights ω_i proportional to building-level supply responses.¹⁰ Intuitively, the average marginal fiscal cost up-weights buildings on the margin of 421-a participation.

The average fiscal cost per inclusionary unit in neighborhood n is, quite similarly,

$$\text{AFC}_n = \frac{1}{\lambda} \int v_i \Delta\tau_i \Pr(D_i = 1|x_i) dF_n(x_i) \bigg/ \int \Pr(D_i = 1|x_i) dF_n(x_i). \quad (10)$$

To compute supply responses as well as marginal and average fiscal costs, I increment all tax differentials $\Delta\tau_i$ in neighborhood n by a small $\varepsilon > 0$. I then simulate the predicted probability of 421-a participation for each building i in neighborhood n under the new and old $\{\Delta\tau_i\}$. The average difference in these probabilities, weighting by the number of residential units and dividing by ε , yields the estimated supply response. With the set-aside share and the actual assessed value of each building i , I also use the predicted probabilities to simulate both the total change in tax

¹⁰ $\omega_i = \Pr(D_i = 1|x_i) [1 - \Pr(D_i = 1|x_i)] / \int \Pr(D_i = 1|x_i) [1 - \Pr(D_i = 1|x_i)] dF_n(x_i)$.

revenue from increased 421-a participation and the total change in the number of inclusionary units. Dividing the former by the latter yields the marginal fiscal cost. I compute standard errors for both supply responses and marginal fiscal costs by the cluster bootstrap.

8.2 Distribution of Developer Breakevens

Figure 3 plots a histogram of developer breakevens, expressed in dollars per inclusionary unit, split by 421-a status. Citywide, the median building's breakeven is about \$810,000. There is much variation in breakevens: The 25th and 75th percentiles are respectively around \$500,000 and \$1,600,000. Participating buildings are visibly selected towards lower breakevens, both in dollar terms and, as I show in Appendix Figure A13, in terms of shares of building value.

At most levels of breakeven cost, there are both buildings that accept and reject 421-a. This is for two reasons. First, both the 421-a benefit and its opportunity cost are larger in dollar terms when buildings are more expensive. This explains why a \$100,000-per-unit benefit could be accepted by a low-rent building, but a \$500,000-per-unit benefit could be rejected by a high-rent building. Second, I cannot perfectly explain the participation decision with observables. Estimated breakevens thus will not perfectly divide buildings into those which accept and those which reject.

What descriptively explains this variation in breakevens? In Appendix Figure A10, I present a binned scatterplot of buildings' breakevens versus the average capitalized rent of units in the same block group as the building. There is a robust positive relationship between breakevens and block-group rents. However, breakevens are higher than typical capitalized rents in the same block.

8.3 Citywide Supply Response and Marginal Fiscal Cost

Table 4 reports my estimates of the citywide supply response, marginal fiscal cost, and average fiscal cost of inclusionary units, as derived respectively in Equations 8, 9, and 10. Column 1 contains the baseline estimates that treat building characteristics as exogenous, while Column 2 presents the simulated-instrument GMM results.

I find a citywide 421-a supply response η of 0.59. This estimate implies that an increase in the 421-a tax savings of one percentage point of building value would increase the number of

inclusionary units by 0.59 percent. I also find a citywide marginal fiscal cost of about \$1.6 million per unit. The average fiscal cost is lower at about \$640,000 per unit. Accounting for endogenous building characteristics leads to a slightly larger supply elasticity and a slightly lower cost per inclusionary unit. For the rest of this paper, all results take building characteristics as exogenous.

In Appendix Table A11, I estimate the supply response and marginal fiscal cost under alternative parametric specifications: a logit model with neighborhood random effects, a probit model, and a linear probability model. These specifications reassuringly suggest my estimates are not highly sensitive to functional form. In Appendix Figure A25, I show how the supply response and average marginal fiscal cost would change under alternative capitalization rate assumptions. Supply responses are increasing, and marginal fiscal costs are decreasing, in the capitalization rate. These shifts occur because the present value of the 421-a incentive decreases in the capitalization rate. If, for instance, the capitalization rate were instead 2.5 percent rather than 5 percent, I would obtain a marginal fiscal cost around \$1.8 million and a supply response of about 0.5.

8.4 Estimates by Neighborhood

How do the supply responses and average marginal fiscal costs of inclusionary units vary among neighborhoods? In this subsection, I compute these objects by Equations 8 and 9 for each of the 179 Neighborhood Tabulation Areas with eligible development from 2003 to 2015. In all specifications, I include lot- and block-level controls as well as borough fixed effects.¹¹

Panel A of Figure 4 maps estimated supply responses η_n for each neighborhood. They vary dramatically: The minimum is less than 0.01 and the maximum is 1.22. Supply responses are highest in Manhattan and the Bronx and are lowest in Queens and Staten Island. Variation in supply responses between neighborhoods primarily reflects variation in take-up rates rather than between-neighborhood variation in the extent of within-neighborhood building heterogeneity.

¹¹To estimate the logit with NTA fixed effects, I must drop NTAs in which all buildings either uniformly accept or reject 421-a. Such NTA fixed effects are otherwise perfect predictors of 421-a participation. For results with NTA fixed effects, see Appendix Figure A15. In all my results, I cannot estimate quantities of interest in neighborhoods without eligible new development.

Panel B of Figure 4 maps neighborhood average marginal fiscal costs per inclusionary unit. The citywide average masks substantial variation by neighborhood. In high-rent Manhattan neighborhoods such as the West Village and the Upper East Side, the fiscal cost of the marginal inclusionary unit is around \$2 million. Such high estimates are sensible for these neighborhoods, where two-bedroom condominiums regularly traded at prices above \$1 million during my period. By contrast, in typical neighborhoods in the Bronx, Queens, and Staten Island, the fiscal cost of the marginal inclusionary unit is less than \$150,000. Such low fiscal costs are consistent with market rents close to the regulated maximum rent on inclusionary units.

9 The Cost-Effectiveness of Inclusionary Housing

In this section, I evaluate the cost-effectiveness of 421-a in three parts. First, I benchmark 421-a's fiscal cost per unit against Section 8 vouchers and the LIHTC. I focus here on average, rather than marginal, fiscal costs due to a lack of marginal-cost estimates for Section 8 and the LIHTC. Second, I explain these cost differences, focusing on the roles of neighborhoods and incidence on developer profits. Third, I calculate the marginal value of public funds spent on 421-a in each neighborhood, combining my cost estimates with external estimates of the long-run benefits of moving households to higher-opportunity neighborhoods.

9.1 Cost Differences Between Programs

I construct estimates of the average fiscal cost of Section 8 voucher and LIHTC units in New York City for comparison to my estimates for 421-a inclusionary units. For Section 8 vouchers, I consult the HUD Picture of Subsidized Households database (see Appendix C), which aggregates household-level administrative data from local housing authorities. On average in 2015, the fiscal cost per Section 8 voucher unit in New York City was \$245,800. For the LIHTC, I consult four external sources of estimates, as reviewed in Appendix C. These estimates range from \$197,400 to \$234,580, suggesting there is relatively little uncertainty about the per-unit cost of the LIHTC. I take \$220,000 as a reasonable midpoint estimate. In New York City, Section 8 voucher and LIHTC units therefore have similar average fiscal costs. Figure A19 shows, however, that average costs for

421-a units are much costlier than either program. In particular, my estimate of the average fiscal cost per inclusionary unit is about three times higher than that of Section 8 or the LIHTC.

There are several potential explanations for the higher average cost of 421-a inclusionary units. These explanations can be divided into between- and within-neighborhood factors. The obvious between-neighborhood explanation is that 421-a units are in costlier neighborhoods, consistent with Table 2, implying higher breakevens. However, it is also possible that 421-a inclusionary units are costlier than Section 8 or LIHTC units within the same neighborhood. Potential explanations of such within-neighborhood differences include differences between programs in administrative costs, building amenities, and incidence on developer profits.

Here I examine the extent to which between-neighborhood differences can explain the 421-a cost premium and the empirical importance of one salient within-neighborhood explanation: developer incidence. One virtue of such a decomposition is that, insofar as governments have some willingness to pay to move households to “better” neighborhoods, it has welfare implications. That is, the 421-a cost premium may be justified as the price of better neighborhoods, whereas within-neighborhood factors (especially incidence) seem less conducive to such justifications.

Neighborhoods. To assess the role of differences in neighborhoods to cost differences, I attempt to make the neighborhoods of 421-a units comparable to those of Section 8 voucher and LIHTC units. I do so following DiNardo et al. (1996), which reweights a sample so that its distribution over a set of variables resembles the distribution of a target sample on those variables.

For a Census block b , let $N_{1,b}$ be the number of 421-a units, and let $N_{0,b}$ be the number of pooled Section 8 and LIHTC units. I then construct a block’s 421-a share of units $s_{1,b} = N_{1,b}/(N_{0,b} + N_{1,b})$ and estimate the fractional logistic regression $s_{1,b} = \exp(z_b\gamma)/[1 + \exp(z_b\gamma)]$, which describes the 421-a share of units as a function of block characteristics z_b . Applying Bayes’ rule, DiNardo et al. (1996) obtain the reweighting factor

$$\psi(z_b) = \frac{dF(z_b | \text{Unit in Section 8 or LIHTC})}{dF(z_b | \text{Unit in 421-a})} \propto 1 - 1/\widehat{s}_{1,b},$$

where $dF(z_b | \cdot)$ is a joint density function over block characteristics, conditional on units being in

either 421-a or pooled Section 8 and LIHTC, and $\widehat{s}_{1,b}$ is the fitted 421-a share from the fractional logistic regression. The reweighting factor $\psi(z_b)$ thus makes the 421-a units resemble the Section 8 and LIHTC units on block characteristics. In particular, $\psi(z_b)$ tends to up-weight 421-a buildings in low-cost neighborhoods, given that Section 8 and LIHTC units are concentrated in low-cost neighborhoods relative to 421-a units, as shown in Appendix Figure A12.

Introducing these weights into Equation 10, I obtain the reweighted citywide average fiscal cost under the counterfactual in which 421-a units are geographically distributed as if they are Section 8 and LIHTC units:

$$\widetilde{\text{AFC}} = \frac{1}{\lambda} \int \psi(z_i) v_i \Delta \tau_i \Pr(D_i = 1 | x_i) dF(x_i) / \int \psi(z_i) \Pr(D_i = 1 | x_i) dF(x_i).$$

Estimation of reweighted average cost follows the same procedures in Section 7 for actual marginal and average cost. Appendix Figure A19 displays the results of this reweighting. While there are large unadjusted cost differences between 421-a and Section 8 or between 421-a and the LIHTC, these cost differences are entirely eliminated once the 421-a units are reweighted. The large differences in neighborhood characteristics between 421-a and other housing programs, as documented in Table 1, are therefore key to the cost premium of 421-a.

Developer Incidence. How much of every dollar spent on 421-a do developers capture? I define the incidence of 421-a as the share of the fiscal cost that accrues to developer profits. By the distributional assumption on Δe_i in Section 7, developer incidence has a convenient expression: a log-sum of exponentials, akin to the measure of consumer surplus in Small and Rosen (1981):

$$\frac{E[\Delta \log \pi_i | D_i = 1]}{E[\Delta \tau_i | D_i = 1]} = \frac{\sigma \log \int \left[1 + \exp\left(\frac{\Delta \tau_i + x_i \beta}{\sigma}\right) \right] dF_n(x_i | D_i = 1)}{\int \Delta \tau_i dF_n(x_i | D_i = 1)}, \quad (11)$$

where $F_n(x_i | D_i = 1)$ is the distribution of characteristics among 421-a participants in a neighborhood n . To measure incidence with respect to the average dollar spent on 421-a, I weight by market value. I report estimates and cluster-bootstrap standard errors in Appendix Table A9.

I estimate that, citywide, developers capture in additional profits about \$0.46 of every \$1 spent

on 421-a.¹² This degree of developer incidence appears similar to that of Section 8 and the LIHTC: Collinson and Ganong (2018), for example, find developers also capture \$0.46 cents of every \$1 in increased Section 8 voucher generosity. Consequently, higher developer incidence of 421-a versus other housing programs is unlikely to be an important source of the 421-a cost premium. This conclusion contradicts critics of inclusionary housing programs quoted in Section 1.

Is there “excess” developer incidence of 421-a? That is, could a tax reform designed to “squeeze” developers meaningfully reduce the 421-a tax expenditure without reducing take-up? I argue that the answer is likely not, and that potential savings are considerably less than the estimated developer incidence of 46 percent. As developer breakevens are private information, any tax incentive that achieves positive take-up must pay an information rent and therefore cannot achieve zero developer incidence.¹³ It can be shown that the incidence-minimizing feasible policy that attains a developer participation rate of p is

$$\Delta\tau^*(x_i, p) = \log\left(\frac{p}{1-p}\right) - x_i\beta.$$

Intuitively, these tax incentives minimize developer incidence because the incidence share in Equation 11 is convex in $\Delta\tau_i$, which implies a necessary condition for incidence minimization is equalization of take-up probabilities. This $\Delta\tau^*(x_i, p)$ uniquely equalizes them at p .

I then simulate incidence under this counterfactual policy. To do so, I set p to the true take-up probability and use my estimates $\hat{\beta}$. I find that, under this incidence-minimizing counterfactual, developers capture about 36 percent of the total fiscal cost of 421-a, 10 percentage points less than actual incidence. Excess developer incidence appears to represent a small share of the total fiscal cost of 421-a and is unimportant to explaining cost differences between housing programs.¹⁴

¹²For incidence by Neighborhood Tabulation Area, see Appendix Figure A18.

¹³Developers only accept if $\Delta\pi_i \geq 0$. Idiosyncratic costs Δe_i are unobservable, uncorrelated with x_i , and have positive variance. Thus, it is infeasible to offer a set of incentives $\{\Delta\tau_i\}$ that achieves $\Delta\pi_i = 0$ for all i . A positive mass of developers must therefore be offered $\Delta\pi_i > 0$ if the participation rate is positive. This implies developer incidence.

¹⁴If developers earn profits under nonparticipation, then breakevens include a profit component.

9.2 Valuing the Benefits of Inclusionary Housing

Should governments be willing to pay the 421-a cost premium to move households into such neighborhoods—and, if so, which neighborhoods are best? Should governments move a few households to the highest-income and highest-cost neighborhoods, or should they move more households to middle-income neighborhoods that are less costly? Here I compare the neighborhood-level benefits and costs of 421-a. Such an exercise entails substantial further assumptions, as I explain below, relative to my analysis thus far. The intent of this section is therefore to illustrate how my cost estimates could be combined with estimates of benefits to conduct a cost-benefit analysis.

To value the benefits of 421-a, I focus on the long-run effects of such neighborhoods on children's future income and tax payments as adults.¹⁵ In particular, I follow Bergman et al. (2020) in using external estimates of neighborhood effects from the Opportunity Atlas (Chetty et al., 2018). In combination with my cost estimates, I use these estimated benefits to conduct an analysis of the marginal value of public funds (MVPF) of 421-a by neighborhood. Defined in Hendren and Sprung-Keyser (2020) as the ratio of households' WTP for a policy and its cost to the government net of fiscal externalities, a policy's MVPF measures the shadow price to the government of raising a household's utility via the policy. While a direct non-distortionary transfer has an MVPF of one—a household values \$1 at exactly the \$1 it costs the government—policies may have MVPFs above or below one depending on willingness to pay and fiscal externalities per dollar of program cost. The MVPF is therefore a useful summary measure of cost-effectiveness that enables comparisons, both between policies and places.

My calculation of MVPFs requires four present discounted values for each neighborhood: the

Then incidence is the incremental profit developers earn above their profit under nonparticipation.

¹⁵This definition of benefits omits several considerations. First, policy interventions to reduce segregation may have an efficiency motive due to non-fiscal externalities, e.g. via human capital spillovers (Bénabou, 1993). Second, I do not account for any impacts on WTP or fiscal externalities via crime or other channels beyond the tax-and-transfer system. Third, I ignore non-individualistic motives for inclusion. Fourth, I ignore all impacts on incumbent residents of neighborhoods.

WTP for housing, the marginal fiscal cost, the neighborhood effect on children’s after-tax incomes as adults, and the long-run fiscal externality on children—that is, the increase in their tax payments as adults. In particular, the MVPF of 421-a in neighborhood n is defined as

$$\text{MVPF}_n = \frac{\text{PDV}(\text{WTP for Housing}_n) + \text{PDV}(\text{Future After-Tax Income}_n)}{\text{PDV}(\text{Marginal Fiscal Cost}_n) - \text{PDV}(\text{Long-Run Fiscal Externality}_n)}. \quad (12)$$

If the impact of inclusionary housing on earnings and thus the long-run fiscal externality is zero, then the MVPF of 421-a in a neighborhood follows from developer incidence. To the extent that its intergenerational impacts are economically large, however, the MVPF will diverge from incidence.

I use my estimates to compute the PDVs of the WTP for housing and the marginal fiscal cost. For the first item, I compute “high-WTP” and “low-WTP” scenarios. In the high-WTP scenario, I assume that households value inclusionary units at exactly their breakeven cost. This represents an upper bound if households can rent units on the private market at their breakevens. In the low-WTP scenario, I assume that households value each \$1 of breakeven at \$1 to the extent it is less than their counterfactual housing expenditure and at zero for every \$1 above this counterfactual. This is a lower bound insofar as the marginal utility of housing consumption is weakly positive. Second, I use directly my estimates for neighborhood-specific costs.

To obtain the fiscal externality and the impact on after-tax income, I closely follow the procedure in Bergman et al. (2020) to predict the upward mobility impacts of moving households to specific neighborhoods. This procedure converts the Opportunity Atlas estimates of intergenerational income mobility by neighborhood into percentage changes in income using the citywide income distribution and an assumption that 62 percent of the variation in upward mobility across neighborhoods reflects their causal effects on adult income rank. The main departure from the Bergman et al. (2020) procedure is that, as I do not observe a household’s counterfactual neighborhood had the inclusionary unit not existed, I use a calibrated model of neighborhood choice to predict these counterfactual neighborhoods. Building upon the publicly-released code base of Hendren and Sprung-Keyser (2020), I compute the implications of 421-a for children’s lifetime after-tax income and fiscal externalities. This calculation also requires information on New York

City’s labor income tax schedule, age–earnings profile, and intergenerational rank–rank mobility function as well as several statistics on 421-a beneficiary households that I assemble from several sources of available data. As observed in Hendren and Sprung-Keyser (2020), the calculation of MVPFs involves numerous “judgment calls,” which I document in further detail in Appendix C.¹⁶ Appendix D contains several extensions and robustness checks of the MVPF analysis.

Figure 5 maps the MVPF of 421-a by Neighborhood Tabulation Area in the low-WTP scenario. MVPFs vary immensely between neighborhoods: Whereas 421-a produces near-zero benefit per dollar of net fiscal cost in some neighborhoods ($MVPF \approx 0$), the benefits of a better neighborhood come very cheaply in others ($MVPF > 5$). Neighborhoods broadly divide into three groups with respect to their MVPFs. First, some have low fiscal costs but such poor outcomes for children that their MVPFs are near zero. Most strikingly, this condition characterizes the South Bronx. Second, some neighborhoods have good outcomes for children but are quite costly, as in Manhattan, generating MVPFs below one in the low-WTP scenario. Third, some neighborhoods are “opportunity bargains.” These neighborhoods have low fiscal costs and yet children raised there achieve outcomes no worse than do children raised in high-cost areas. These high-MVPF neighborhoods are largely middle-class communities in Queens, Staten Island, and outlying areas of Brooklyn.

10 Conclusion

Residential segregation by household income is a fact of urban life, one that harms children who grow up in neighborhoods that offer little hope of upward mobility. Local governments have introduced housing policies that seek to alleviate segregation and expand economic opportunity via mandates and incentives for mixed-income housing development. How costly are these programs? How do their costs compare to housing programs that leave unaddressed the socioeconomic divisions between neighborhoods? Is paying more today for inclusionary housing a cost-effective way to improve the long-run outcomes of children born into low-income households?

In this paper, I introduce a microeconometric approach to evaluating a voluntary inclusionary

¹⁶One key assumption is that the average 421-a household has 1.44 children, which I calibrate using data on all housing-subsidy recipient households in New York City from 2003 to 2015.

housing program, and I apply it to the paradigmatic example of such policies, New York City's 421-a tax incentive for mixed-income development. My approach begins from the observation that these programs face a participation constraint: Developers participate only if it is profitable for them. From the distribution of these constraints emerges the supply of inclusionary units. Using developer revealed preference and variation in potential tax savings under 421-a, I estimate the marginal fiscal cost of inclusionary units, both citywide and neighborhood-by-neighborhood.

I find that, on citywide average, it would cost New York City \$1.6 million to add another inclusionary unit under 421-a, which greatly exceeds the per-unit costs of Section 8 vouchers and the LIHTC. However, I also find that the 421-a cost premium primarily reflects differences in the distribution of units across neighborhoods, rather than within-neighborhood cost differences. There is no evidence of differentially higher incidence of 421-a on developers. Whether 421-a is cost-effective thus boils down to whether governments should pay higher rents to move families on housing assistance into better neighborhoods. Weighing 421-a's costs against its potential long-run benefits, I find that some—but not all—neighborhoods are “opportunity bargains.”

References

Anderson, Soren T and James M Sallee, “Using Loopholes to Reveal the Marginal Cost of Regulation: The Case of Fuel-Economy Standards,” *American Economic Review*, 2011, 101 (4), 1375–1409.

Baum-Snow, Nathaniel and Justin Marion, “The Effects of Low Income Housing Tax Credit Developments on Neighborhoods,” *Journal of Public Economics*, 2009, 93 (5-6), 654–666.

Bénabou, Roland, “Workings of a City: Location, Education, and Production,” *The Quarterly Journal of Economics*, 1993, 108 (3), 619–652.

Benzarti, Youssef, “How Taxing is Tax Filing? Using Revealed Preferences to Estimate Compliance Costs,” *American Economic Journal: Economic Policy*, 2020, 12 (4), 38–57.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F Katz, and Christopher Palmer, “Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice,” Working Paper 26164, National Bureau of Economic Research 2020.

Chetty, Raj, John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter, “The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility,” Working Paper 25147, National Bureau of Economic Research 2018.

—, **Nathaniel Hendren, and Lawrence F Katz**, “The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment,” *American Economic Review*, 2016, *106* (4), 855–902.

Chyn, Eric, “Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children,” *American Economic Review*, 2018, *108* (10), 3028–56.

Collinson, Robert and Peter Ganong, “How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality?,” *American Economic Journal: Economic Policy*, 2018, *10* (2), 62–89.

—, **Ingrid Gould Ellen, and Jens Ludwig**, “Low-Income Housing Policy,” in “Economics of Means-Tested Transfer Programs in the United States, Volume 2,” University of Chicago Press, 2015, pp. 59–126.

Currie, Janet and Jonathan Gruber, “Health Insurance Eligibility, Utilization of Medical Care, and Child Health,” *The Quarterly Journal of Economics*, 1996, *111* (2), 431–466.

Davis, Morris, Jesse Gregory, Daniel Hartley, Kegen Tan et al., “Neighborhood Effects and Housing Vouchers,” Working Paper 2020.

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

—, —, and **Franklin Qian**, “The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco,” *American Economic Review*, 2019, 109 (9), 3365–94.

Dijk, Winnie Van, “The Socio-Economic Consequences of Housing Assistance,” Working Paper, University of Chicago 2019.

DiNardo, John, Nicole M Fortin, and Thomas Lemieux, “Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach,” *Econometrica*, 1996, 64 (5), 1001–1044.

DiPasquale, Denise, “Why Don’t We Know More About Housing Supply?,” *The Journal of Real Estate Finance and Economics*, 1999, 18 (1), 9–23.

Einav, Liran, Amy Finkelstein, Yunan Ji, and Neale Mahoney, “Voluntary Regulation: Evidence from Medicare Payment Reform,” Working Paper 27223, National Bureau of Economic Research 2020.

Eriksen, Michael D and Stuart S Rosenthal, “Crowd Out Effects of Place-Based Subsidized Rental Housing: New Evidence from the LIHTC Program,” *Journal of Public Economics*, 2010, 94 (11-12), 953–966.

Favilukis, Jack, Pierre Mabile, and Stijn Van Nieuwerburgh, “Affordable Housing and City Welfare,” Working Paper 25906, National Bureau of Economic Research 2019.

Gibbons, Stephen and Alan Manning, “The Incidence of UK Housing Benefit: Evidence from the 1990s Reforms,” *Journal of Public Economics*, 2006, 90 (4-5), 799–822.

Glaeser, Edward L, Joseph Gyourko, and Raven Saks, “Why is Manhattan So Expensive? Regulation and the Rise in Housing Prices,” *The Journal of Law and Economics*, 2005, 48 (2), 331–369.

- Hendren, Nathaniel and Ben Sprung-Keyser**, “A Unified Welfare Analysis of Government Policies,” *The Quarterly Journal of Economics*, 2020, 135 (3), 1209–1318.
- Kisin, Roni and Asaf Manela**, “The Shadow Cost of Bank Capital Requirements,” *The Review of Financial Studies*, 2016, 29 (7), 1780–1820.
- Lincoln Institute of Land Policy and Minnesota Center for Fiscal Excellence**, *50-State Property Tax Comparison Study: For Taxes Paid in 2018* 2019.
- Metcalf, Gabriel**, “Sand Castles Before the Tide? Affordable Housing in Expensive Cities,” *Journal of Economic Perspectives*, 2018, 32 (1), 59–80.
- Olsen, Edgar O and Jeffrey E Zabel**, “US Housing Policy,” in “Handbook of Regional and Urban Economics,” Vol. 5, Elsevier, 2015, pp. 887–986.
- Schuetz, Jenny, Rachel Meltzer, and Vicki Been**, “Silver Bullet or Trojan Horse? The Effects of Inclusionary Zoning on Local Housing Markets in the United States,” *Urban Studies*, 2011, 48 (2), 297–329.
- Sinai, Todd and Joel Waldfoegel**, “Do Low-Income Housing Subsidies Increase the Occupied Housing Stock?,” *Journal of Public Economics*, 2005, 89 (11-12), 2137–2164.
- Singh, Divya**, “Do Property Tax Incentives for New Construction Spur Gentrification? Evidence from New York City,” Working Paper, Columbia University 2020.
- Small, Kenneth A and Harvey S Rosen**, “Applied Welfare Economics with Discrete Choice Models,” *Econometrica*, 1981, pp. 105–130.
- Susin, Scott**, “Rent Vouchers and the Price of Low-Income Housing,” *Journal of Public Economics*, 2002, 83 (1), 109–152.
- Wong, Maisy**, “Estimating Ethnic Preferences Using Ethnic Housing Quotas in Singapore,” *Review of Economic Studies*, 2013, 80 (3), 1178–1214.

Table 1: How Do 421-a Onsite Inclusionary Units Compare to Other Social Housing?

	421-a	LIHTC	Tenant-Based Section 8	Project-Based Section 8	Public Housing	All Rental Units
<i>Panel A: Building-Level Characteristics</i>						
% Social Units	19.3	89.3	n.a.	94.0	100.0	n.a.
<i>Panel B: Block-Level Characteristics</i>						
Med. HH. Income	\$100,043	\$42,223	\$41,187	\$36,899	\$23,420	\$63,093
Med. Monthly Rent	\$2,163	\$1,128	\$1,189	\$964	\$553	\$1,452
% Poor	17.0	31.4	26.4	33.0	43.7	21.2
% Less than HS	10.8	25.7	26.3	27.6	34.2	19.8
% HS Graduate	13.8	24.9	27.3	25.1	30.9	23.0
% Some College	14.4	23.0	22.2	20.4	22.0	19.5
% College Graduate	33.5	16.7	15.3	17.3	9.4	22.1
% More than College	27.6	9.7	8.9	9.7	3.4	15.7
% Non-Hispanic White	44.8	13.0	19.4	17.7	4.9	32.9
% Non-Hispanic Black	16.9	40.5	30.9	32.0	42.8	23.0
% Hispanic	37.6	40.9	41.6	43.8	47.7	30.6
% Asian	14.6	5.1	5.8	5.9	4.4	11.6
Median Age	33.9	33.0	32.8	39.0	33.8	35.8
% Renters	83.8	90.8	83.3	91.1	97.7	78.6
<i>Panel C: Units by Borough</i>						
Manhattan	5,841	41,473	19,825	17,697	54,871	591,114
Bronx	1,166	47,778	49,851	17,354	40,249	392,117
Brooklyn	3,699	33,482	40,327	16,978	57,317	664,749
Queens	1,038	5,424	10,530	4,271	15,500	448,601
Staten Island	134	4,118	2,968	2,865	4,510	60,654

Notes: This table compares inclusionary units added under 421-a to other social housing units (LIHTC, Section 8, and public housing) as well as to all rental units in New York City on building- and block-level characteristics. All statistics in Panels A and B are means, weighted by counts of residential units. For data details, see Section 5 and Appendix C.

Table 2: Summary Statistics on 421-a Eligible Buildings, 2003–2015

	Does the Building Provide Onsite Inclusionary Housing?			
	Yes		No	
	Mean	Std. Dev.	Mean	Std. Dev.
Estimated Market Value (\$ Millions)				
Total	84.12	274.80	16.48	117.41
Land	9.82	39.41	2.05	20.05
Building	74.30	241.91	14.43	101.28
Total Per Unit	0.536	1.524	0.557	1.498
Tax Rate (p.p.)				
Level	0.31	0.32	0.47	0.81
Size of 421-a Incentive	32.95	28.23	4.74	13.07
Present Value of Tax Savings (\$)				
Total	16,107,285	40,975,046	1,238,956	8,486,184
Per Inclusionary Unit	443,782	549,262	126,610	523,299
Number of Units				
Total	105.89	178.21	19.95	85.74
Residential	102.23	175.66	17.48	69.45
Number of Floors	10.49	11.90	4.72	4.93
Rental (vs. Owner-Occupied, %)	0.670	0.471	0.762	0.426
Number of Buildings	581		11,565	
Number of Units	59,393		202,179	

Notes: This table reports means and standard deviations of financial and buildings characteristics of the dataset of developments built from 2003 to 2014 that were eligible for 421-a. See the text of the paper for variable definitions. Appendix Table A3 reports additional summary statistics.

Table 3: Developer Participation Responses to Variation in 421-a Incentive

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Fixed Effects:	None		Borough		Neighborhood	Tabulation Area	
Controls:	None	None	Lot	Lot & Block	None	Lot	Lot & Block
421-a Incentive	4.92*** (0.30)	5.18*** (0.54)	5.58*** (0.64)	5.99*** (0.69)	5.20*** (0.65)	5.87*** (0.76)	6.02*** (0.74)
<i>N</i>	11,669	11,669	11,647	11,640	7,465	7,450	7,445
Clusters	179	179	179	179	82	82	82
Marginal Effect	0.18*** (0.02)	0.18*** (0.03)	0.19*** (0.04)	0.19*** (0.04)	0.26*** (0.03)	0.26*** (0.03)	0.26*** (0.03)
Std. Dev. of Δe_i	0.368*** (0.023)	0.350*** (0.036)	0.325*** (0.038)	0.303*** (0.035)	0.349*** (0.044)	0.309*** (0.040)	0.301*** (0.037)

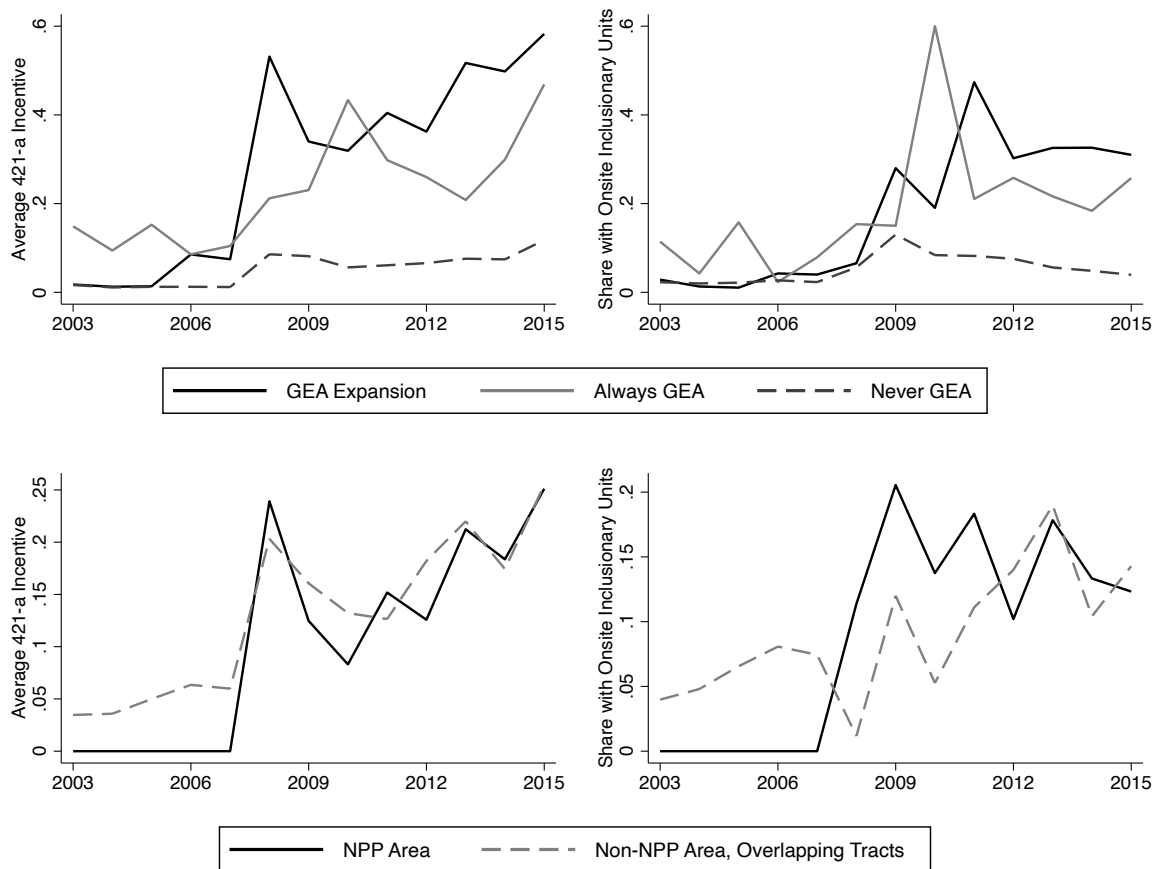
Notes: This table reports robustness checks of estimates of the participation effect of the present-value tax-rate differential $\Delta\tau_i$ from providing onsite inclusionary housing under the 421-a exemption. To the baseline specification in Equation 5, Columns 2–4 respectively add lot controls from PLUTO and Census block controls from the 2017 ACS and 2010 Census. In Columns 5–7, I replace the borough fixed effects with neighborhood fixed effects and reintroduce the lot and block controls. I always also include fixed effects for year of initial permit issuance. Marginal effects reflect the percentage-point increase in the onsite inclusionary housing participation rate per percentage-point increase in the tax incentive. Standard errors are clustered at the neighborhood level. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table 4: Bootstrap Estimates of Citywide Supply Response and Marginal Fiscal Cost

	Exogenous Construction	Endogenous Construction
	(1)	(2)
Supply Response	0.59*** (0.07)	0.60*** (0.10)
Marginal Fiscal Cost	\$1,593,037*** (326,216)	\$1,568,194*** (423,196)
Average Fiscal Cost	\$651,974*** (75,525)	\$650,714*** (76,108)

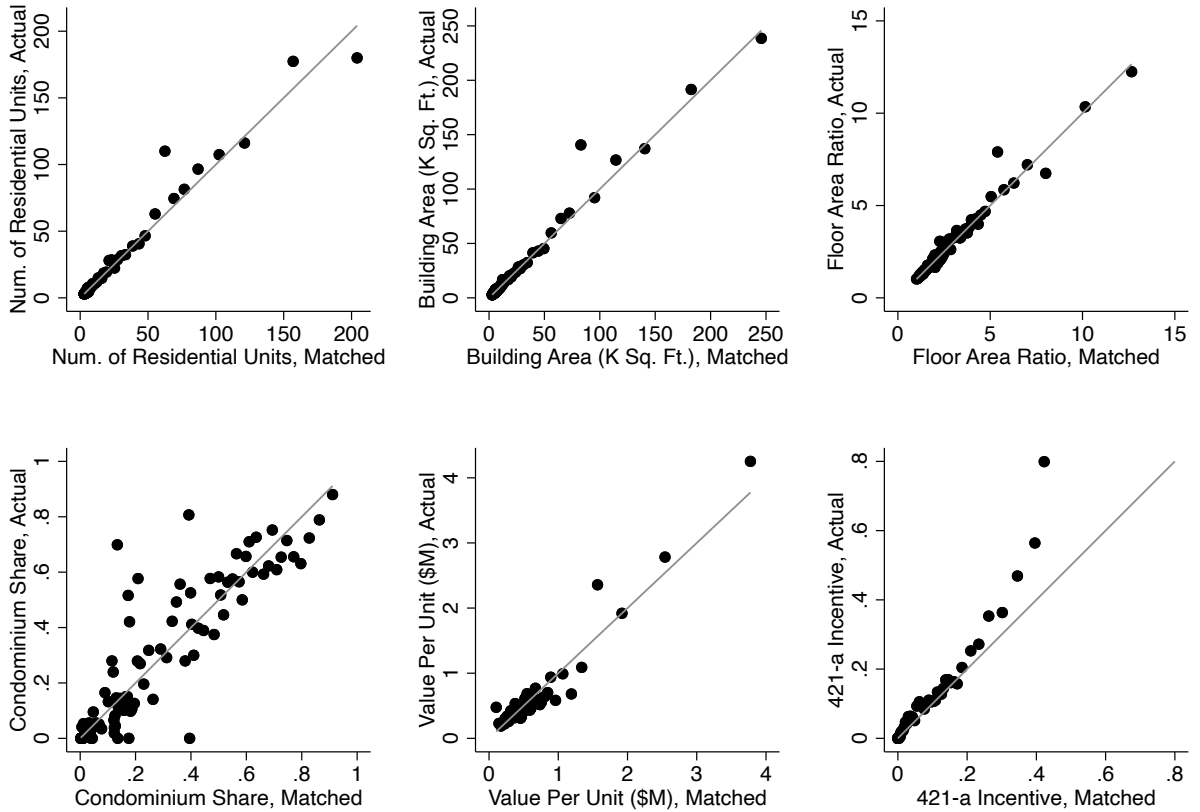
Notes: This table presents citywide supply responses and citywide average and marginal fiscal costs. In the first row, I report estimates of the citywide supply response of onsite inclusionary units to changes in buildings' 421-a incentive. In the second and third rows, I report marginal and average fiscal costs per inclusionary unit per year. In the left column, my estimates are from the baseline approach, which assumes building features are exogenously determined. In the right column, my estimates are from the simulated-instruments GMM approach. The original specification includes fixed effects for borough and year as well as lot and block controls. Standard errors are computed by a cluster-bootstrap at the level of Neighborhood Tabulation Areas. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Figure 1: Supply Responses to Two Changes in 421-a Incentives



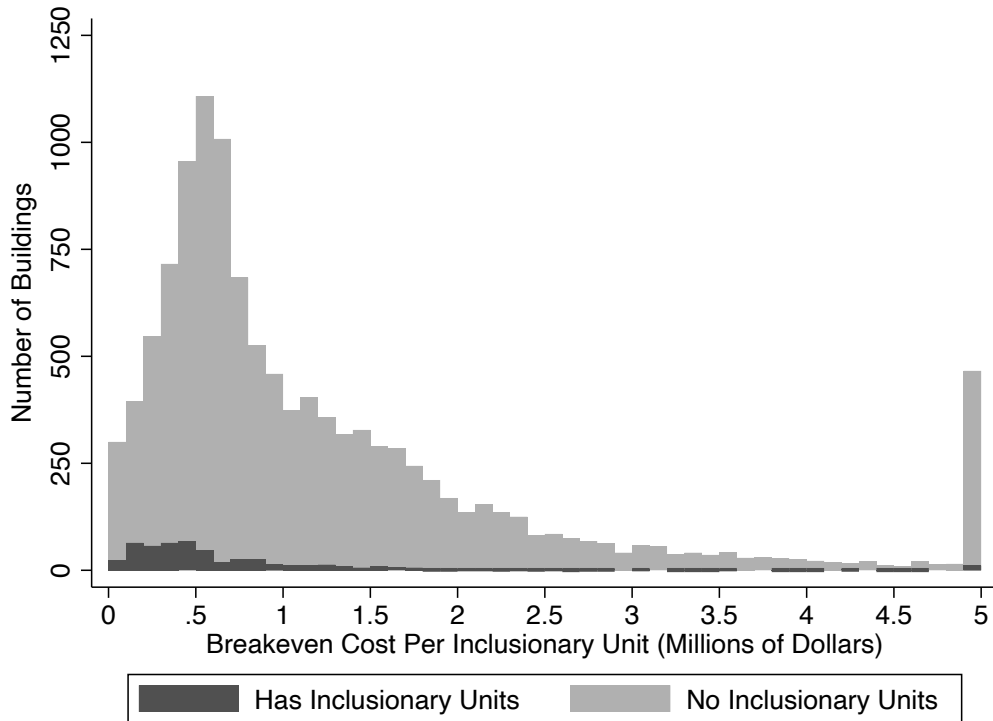
Notes: This figure plots time series of the average 421-a incentive (left panel) and share of buildings with onsite inclusionary units (right panel), split by regions defined as treatment and control groups. The first and second rows show respectively the GEA and NPP reforms in 2008. For further policy context on the GEA and NPP reforms, see Section 2. The incentive is measured by the present value of the 421-a tax savings as a share of market value.

Figure 2: Zeroth and First Stages of Simulated Instrument GMM



Notes: This figure presents binned scatterplots of several endogenous characteristics which compare each building to the average of buildings to which it is matched using lot characteristics. To aid visualization, I drop observations beyond the 1st and 99th percentiles of the simulated characteristic. See Section 7.2 for details on the matching procedure and simulated-instrument GMM approach. I refer to the first five plots as the “zeroth stage,” as they evaluate the performance of the matching procedure on physical building characteristics.


Figure 3: Distribution of Buildings' 421-a Participation Thresholds




Notes: This figure plots a histogram of estimated building-level participation thresholds, expressed in terms of the cost per inclusionary unit. The formula for this value is $-(v_i/\lambda) \cdot x_i \hat{\beta}$, as derived in Equation 7. To aid visualization, I winsorize the distribution at \$5 million per inclusionary unit. I also winsorize at zero eight observations with negative estimated breakevens. The specification includes fixed effects for borough and year as well as lot and block controls. For breakevens as a share of building market value, see Appendix Figure A13.

Figure 4: Estimates by Neighborhood

Panel A: Supply Response




Panel B: Marginal Fiscal Cost per Inclusionary Unit



Notes: These figures, in Panels A and B, respectively display estimates by neighborhood of the supply response of inclusionary units with respect to the tax-rate differential and the marginal fiscal cost per inclusionary unit. The units are percentage-point changes the participation rate per 1 p.p. increase in the 421-a incentive (in Panel A) and thousands of 2015 dollars (in Panel B). The estimated specification includes fixed effects for borough and year as well as lot and block controls. For models with neighborhood fixed effects, see Appendix Figure A15. I am unable to estimate supply responses and marginal fiscal costs neighborhoods with no 421-a eligible new development from 2003 to 2015. Cluster-bootstrap standard errors are displayed in Appendix Figure A14.

Figure 5: Marginal Value of Public Funds of 421-a by Neighborhood Tabulation Area



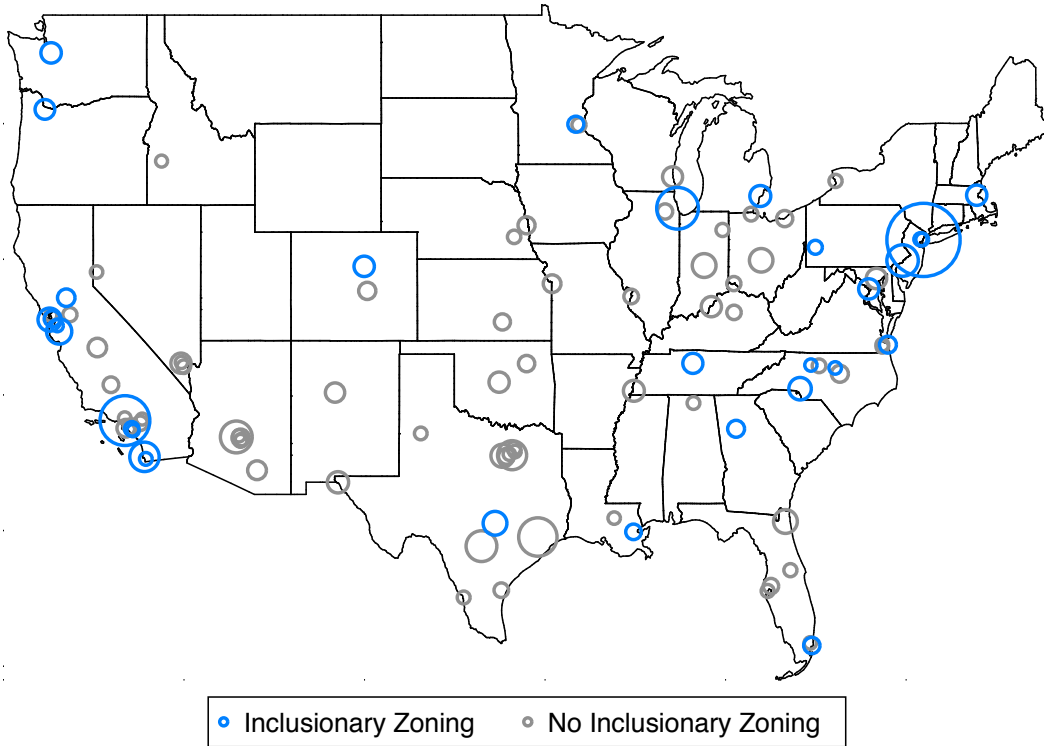
Notes: This figure depicts the marginal value of public funds (MVPF) for 421-a in each Neighborhood Tabulation Area. See Appendix Figure A20 for the estimates in the high-WTP scenario.

Appendices for Online Publication

A Additional Tables and Figures	47
B Model Appendix	89
C Data Appendix	93
D Additional Results	110

A Additional Tables and Figures


Figure A1: Inclusionary Housing in Top 100 U.S. Cities by Population




Notes: This figure maps inclusionary housing policies, either voluntary or mandatory, for the 100 most populous U.S. cities, ranked by their U.S. Census population estimate in 2013. Cities are defined according to municipal boundaries. Circle sizes reflect estimated populations. Data are from Thaden and Wang (2017) and are updated to October 2019 using news reports.

Figure A2: 421-a and the New York City Residential Property Market

Panel A: About One in Three New Units in New York City Are 421-a Participants

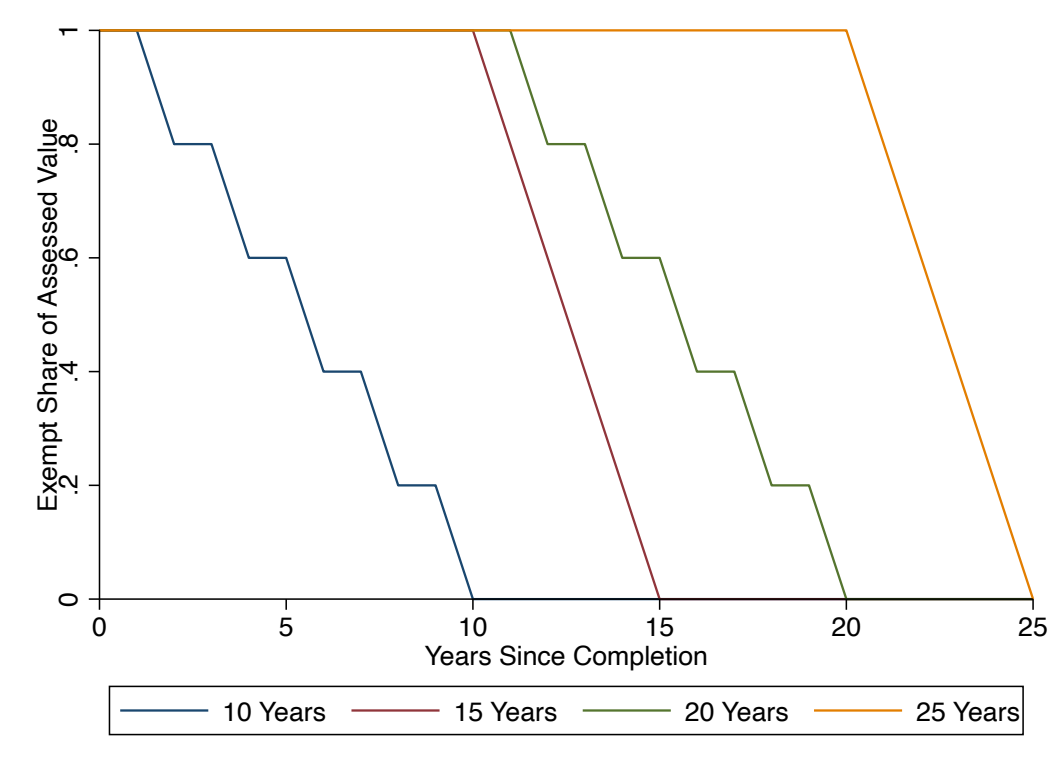


Panel B: 421-a Exemption Valued at 15 Percent of Multifamily Residential Tax Revenue




Notes: Panel A plots the number of new dwelling units issued a Final Certificate of Eligibility (FCE) under the 421-a exemption and the total number of new units completed in each year, as measured by the issuance of Certificates of Occupancy. An FCE is granted, with a regulatory delay, upon completion of construction and is the closest available approximation to the subtotal of completions for units that are 421-a participants. Panel B plots the exempt value of property under 421-a as a share of the total taxable value of all multifamily residential (Class 2) property. The gray bars highlight 2003–2015, the period I study. Sources: 2019 Housing Supply Report, New York City Rent Guidelines Board; 1990–1999 Annual Reports on Tax Expenditures and 2000–2019 Annual Reports of the New York City Property Tax, Division of Tax Policy, New York City Department of Finance; author’s calculations.

Figure A3: 421-a Exemption Schedules



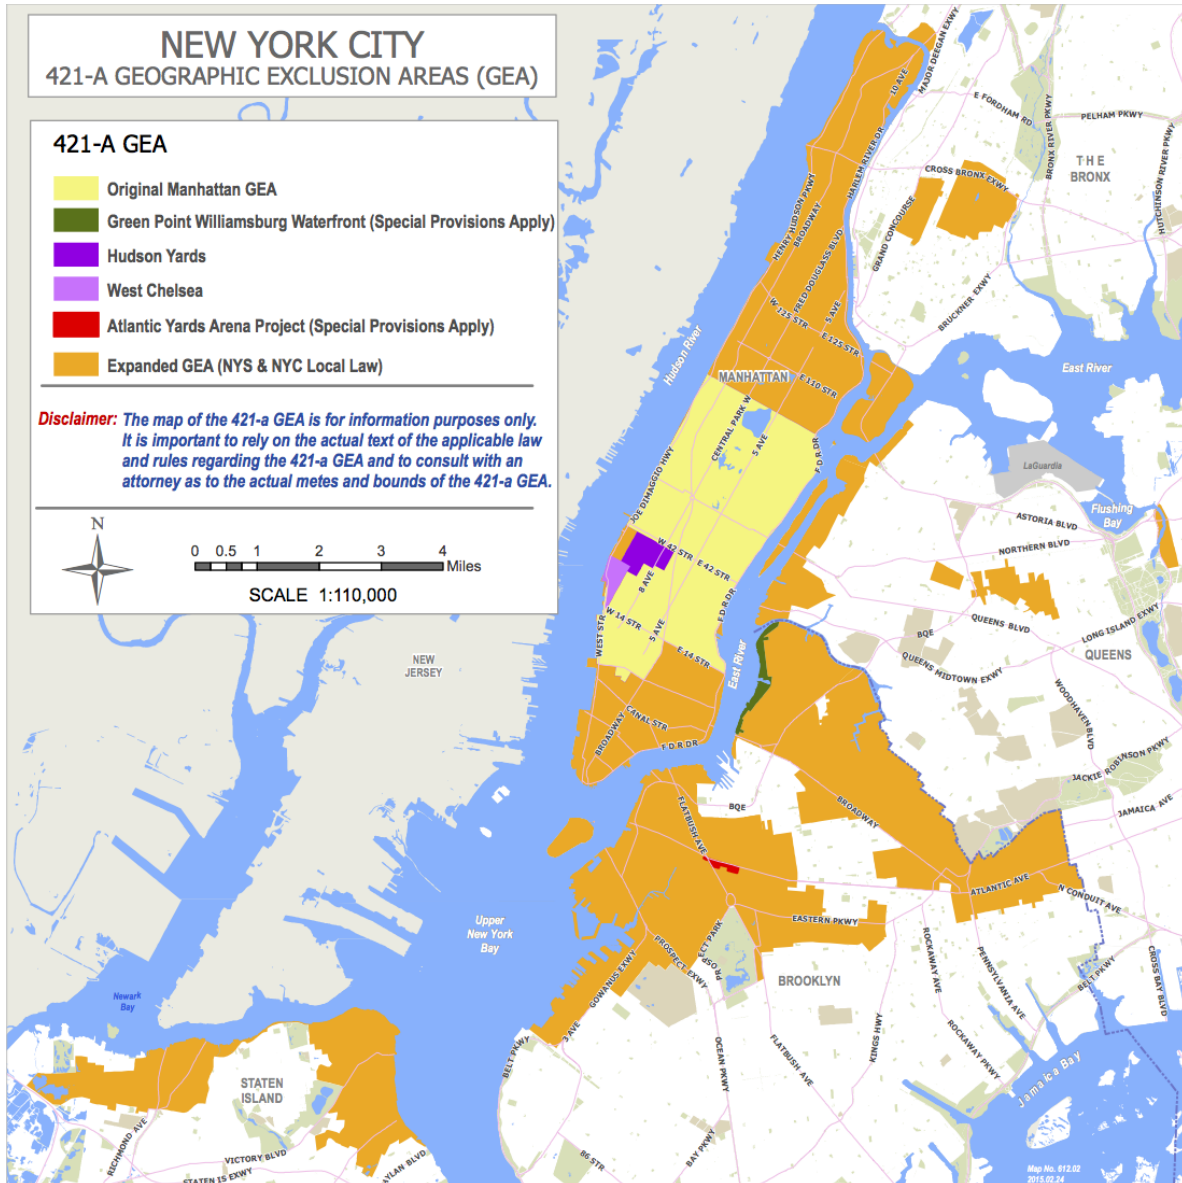
Notes: This figure displays, during each tax year and under each 421-a exemption schedule, the share of building value which is exempt from property taxation. Section 3 explains the applicability of the schedules.

Figure A4: Statutory Property Tax Rates in New York City Since 1981




Notes: This figure displays, from 1981 to 2019, the average statutory property tax rate in each fiscal year by building class. Rates are defined as a share of taxable (i.e., non-exempt) assessed value. In fiscal years with different tax rates by quarter, I take the simple average over quarters. The gray bars highlight 2003–2015, the period I study. Data are drawn from the 2000–2019 Annual Reports of the New York City Property Tax.

Figure A5: Map of Geographic Exclusion Area Over Time



Notes: This figure displays a map of the Geographic Exclusion Area (GEA) and its expansion from the “original” to the “expanded” area. For further details, see Section 3.

Figure A6: Map of Neighborhood Preservation Program (NPP) Areas



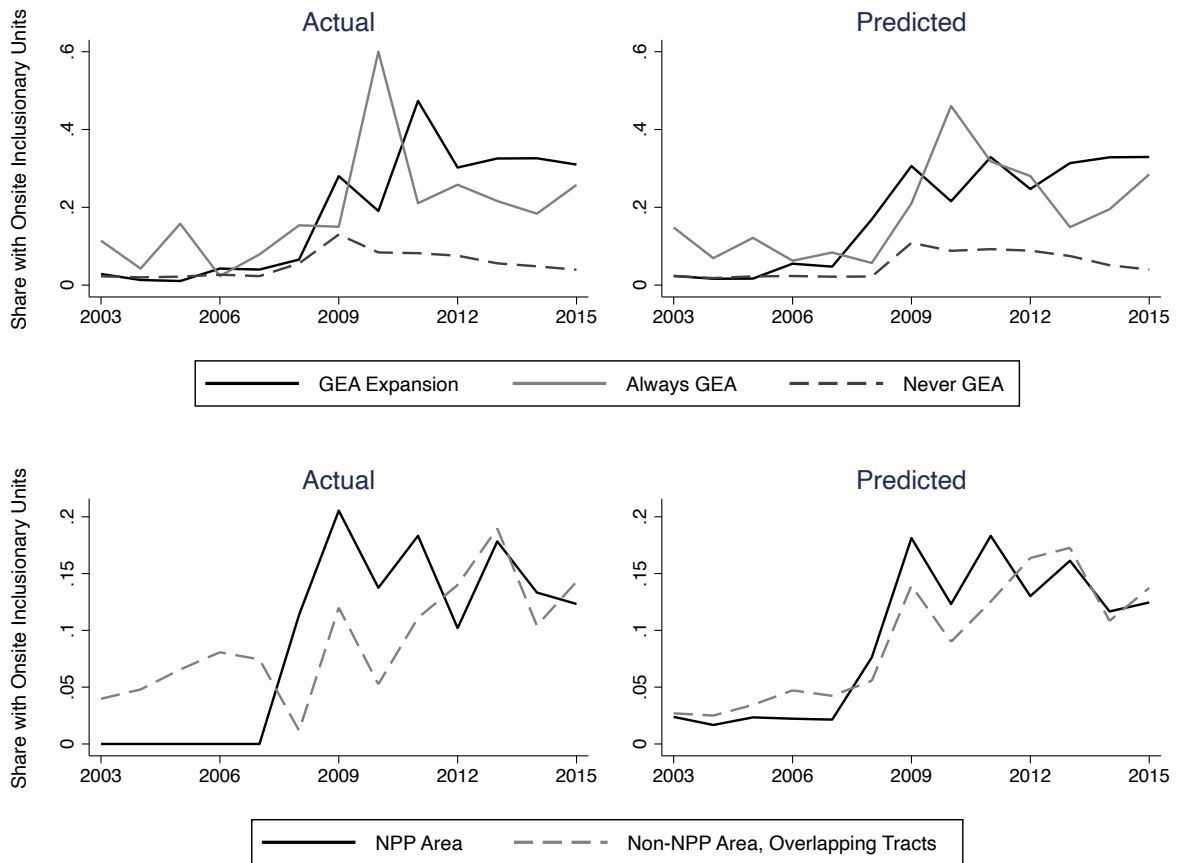
Notes: This figure displays a map of the Neighborhood Preservation Program (NPP) areas as of 1985. NPP areas are indicated in light blue. Source: Laws of the State of New York (208th Session), Volume IV, pp. 3924–3933. For further details, see Section 3.

Figure A7: Type of 421-a Exemption Is Reported on Assessment Roll

ASSESSMENT ROLL 2016-2017 City of New York		
Taxable Status Date: January 5, 2016		
Parcel Information		
Owner Name 110 CORONA LLC	Borough: Queens	Block: 2011
	Lot: 36	
Property Address and Zip Code 11017 CORONA AVENUE 11368		Tax Class: 2
Real Estate Billing Name and Address 110 CORONA LLC 110-17 CORONA AVENUE FLUSHING NY 11368		Building Class: D1
Land Information		
Lot Size	Irregular	Corner
72.37 FT x 112.44 FT	IRREG	
Building Information		
Number of Buildings:	1	
Assessment Information		
Description		
ESTIMATED MARKET		
ACTUAL AV		
ACTUAL EX AV		
TRANS AV		
TRANS EX AV		
Taxable/Billable Assets		
SUBJECT TO ADJUSTMENTS, YOUR	BASED ON	15,379
Exemption Information		
Code	Description	Exempt Value
• 5114-01 (48800)	421A (25 YR NOT CAP)	753,221

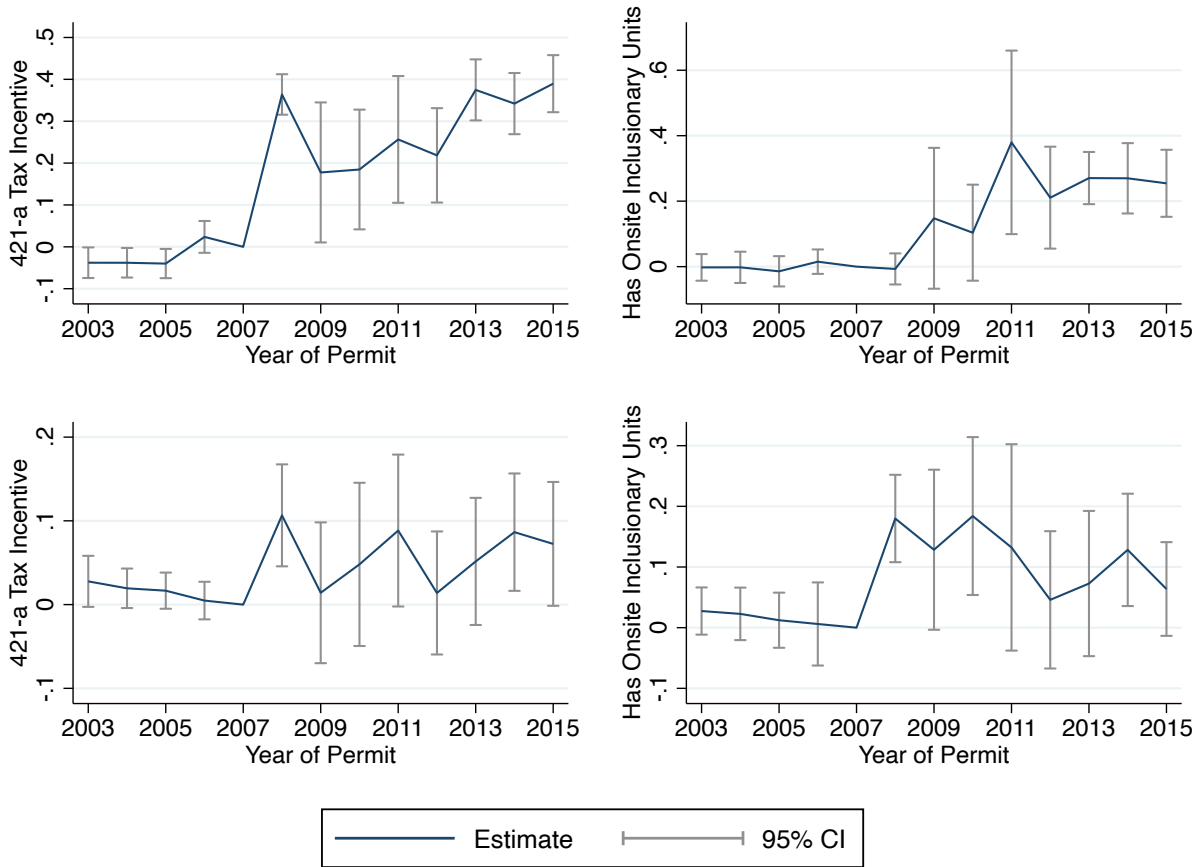
Notes: This figure displays a sample assessment PDF as scraped from the New York City Department of Finance website. In the pullout box, I highlight the “Exemption Information” area. The code 5114-01, along with the building’s BBL and the date on the initial permit for a new building, determines whether the building provided onsite inclusionary housing.

Figure A8: Goodness-of-Fit Check Using 421-a Reforms



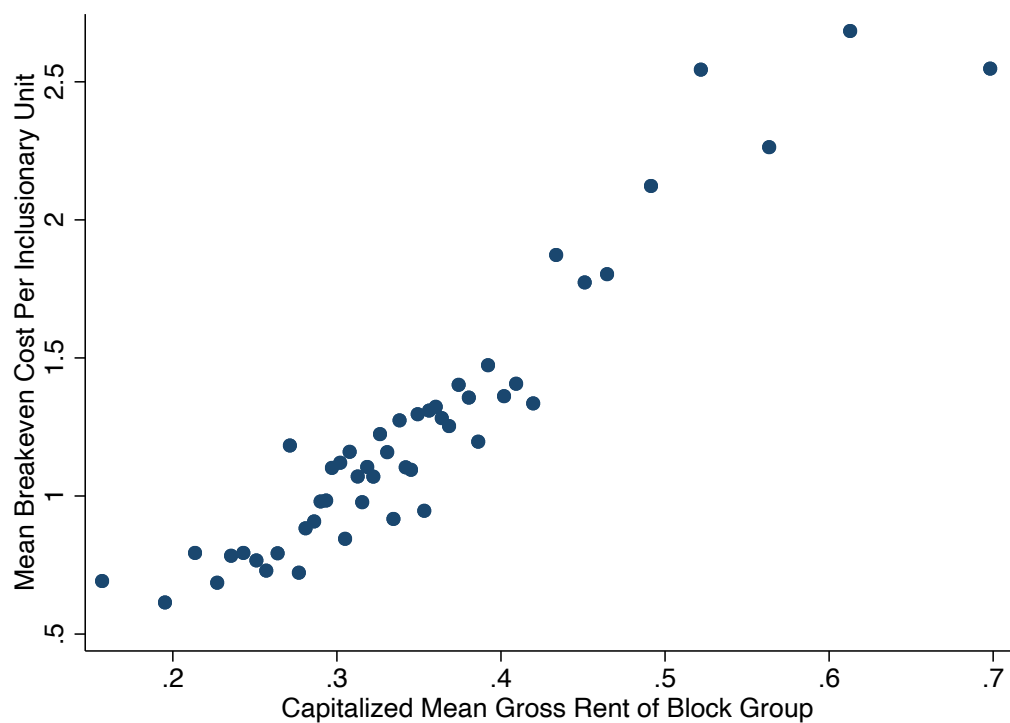
Notes: This figure plots time series of the actual and predicted shares of buildings with onsite inclusionary units, in the left and right panels respectively, split by regions defined as treatment and control groups. The first and second rows show respectively the GEA and NPP reforms in 2008. For further policy context on the GEA and NPP reforms, see Section 2. Figure 1 shows the evolution of the average 421-a incentives by region over time.

Figure A9: Event Studies of GEA and NPP Reforms



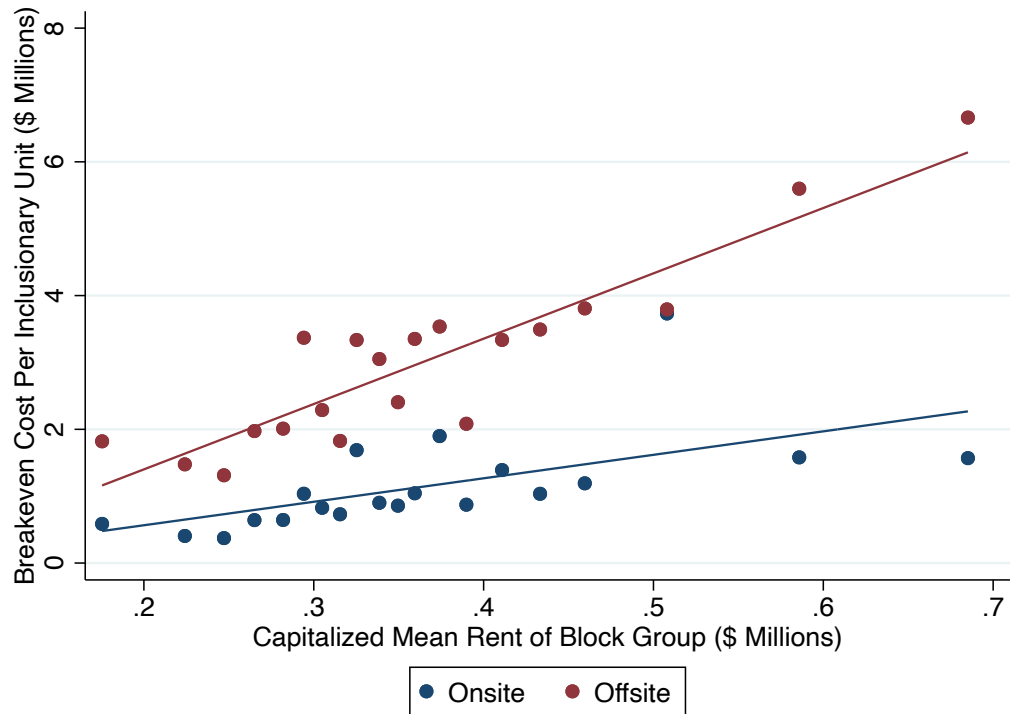
Notes: This figure plots event-study coefficient estimates of the effects of reforms to the Geographic Exclusion Area (GEA) and Neighborhood Preservation Program (NPP). The first and second rows show respectively the GEA and NPP reforms, whereas the left column shows the effect of each reform on the 421-a tax incentive and the right column shows its effect on the 421-a take-up rate. The base year is 2007, immediately both reforms went into effect. Standard errors are clustered by Neighborhood Tabulation Area. See Appendix D for details of the specification and Section 2 for further policy context on the GEA and NPP reforms. Figure 1 shows the evolution of the average 421-a incentives by region over time.

Figure A10: It Costs More to Put Inclusionary Units in Higher-Rent Neighborhoods



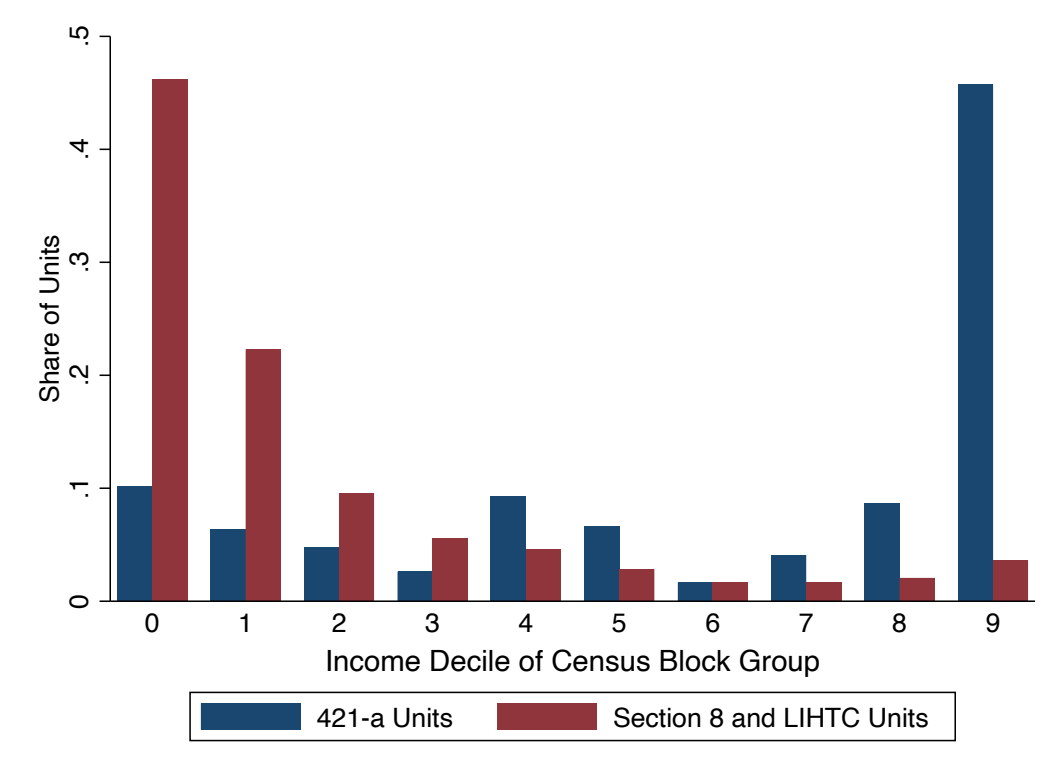
Notes: This figure is a binned scatterplot of the building-level estimated breakeven, as defined in Equation 7, against the capitalized mean gross rent paid in 2015 dollars by households in the same Census block group. I assume a 5-percent annual capitalization rate.

Figure A11: It Costs More to Put *Onsite* Inclusionary Units in Higher-Rent Neighborhoods



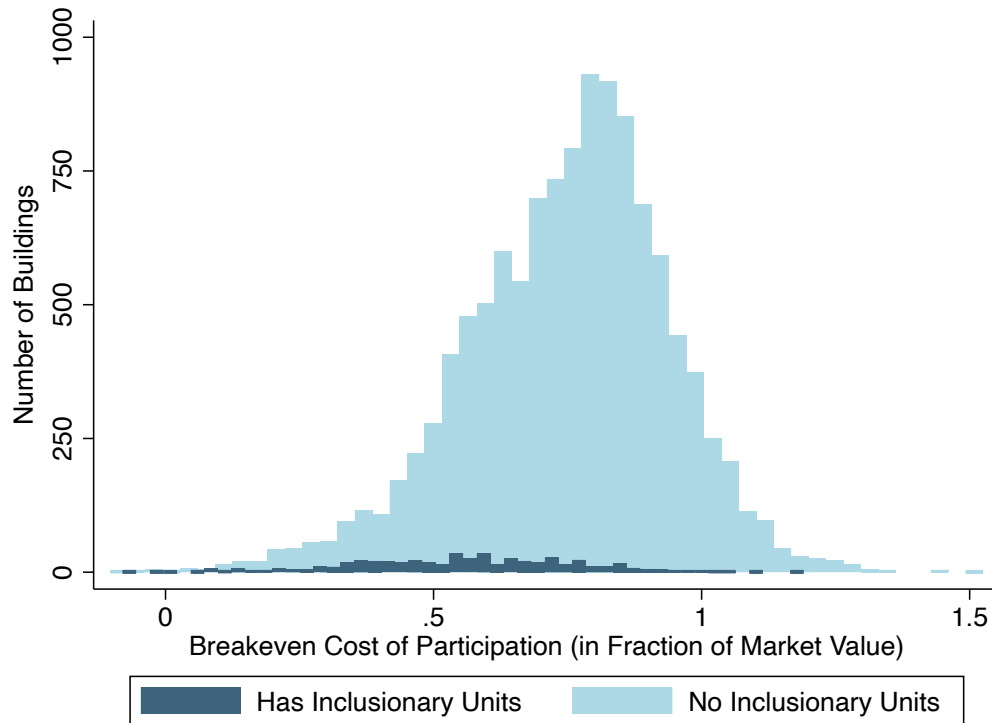
Notes: This figure is a binned scatterplot of the building-level estimated breakeven, as defined in Equation 7, against the mean annual gross rent paid in 2015 dollars by households in the same Census block group. I estimate breakevens for offsite and onsite units using the multinomial logit specification in Equation 20.

Figure A12: 421-a Units Are in Higher-Income Neighborhoods than Section 8 and LIHTC Units



Notes: This figure plots the shares of 421-a inclusionary units and pooled Section 8 voucher and LIHTC units in each deciles of the distribution of Census block-group median annual household income.

Figure A13: Distribution of Buildings' 421-a Participation Thresholds, as Fractions of Market Value





Notes: This figure plots a histogram of estimated building-level breakevens, defined as a fraction of market value. See Figure 3 for dollar breakevens. The specification includes fixed effects borough and year as well as lot and block controls. For more information, see Section 8.

Figure A14: Standard Errors by Neighborhood

Panel A: Supply Response

Panel B: Marginal Fiscal Cost per Inclusionary Unit

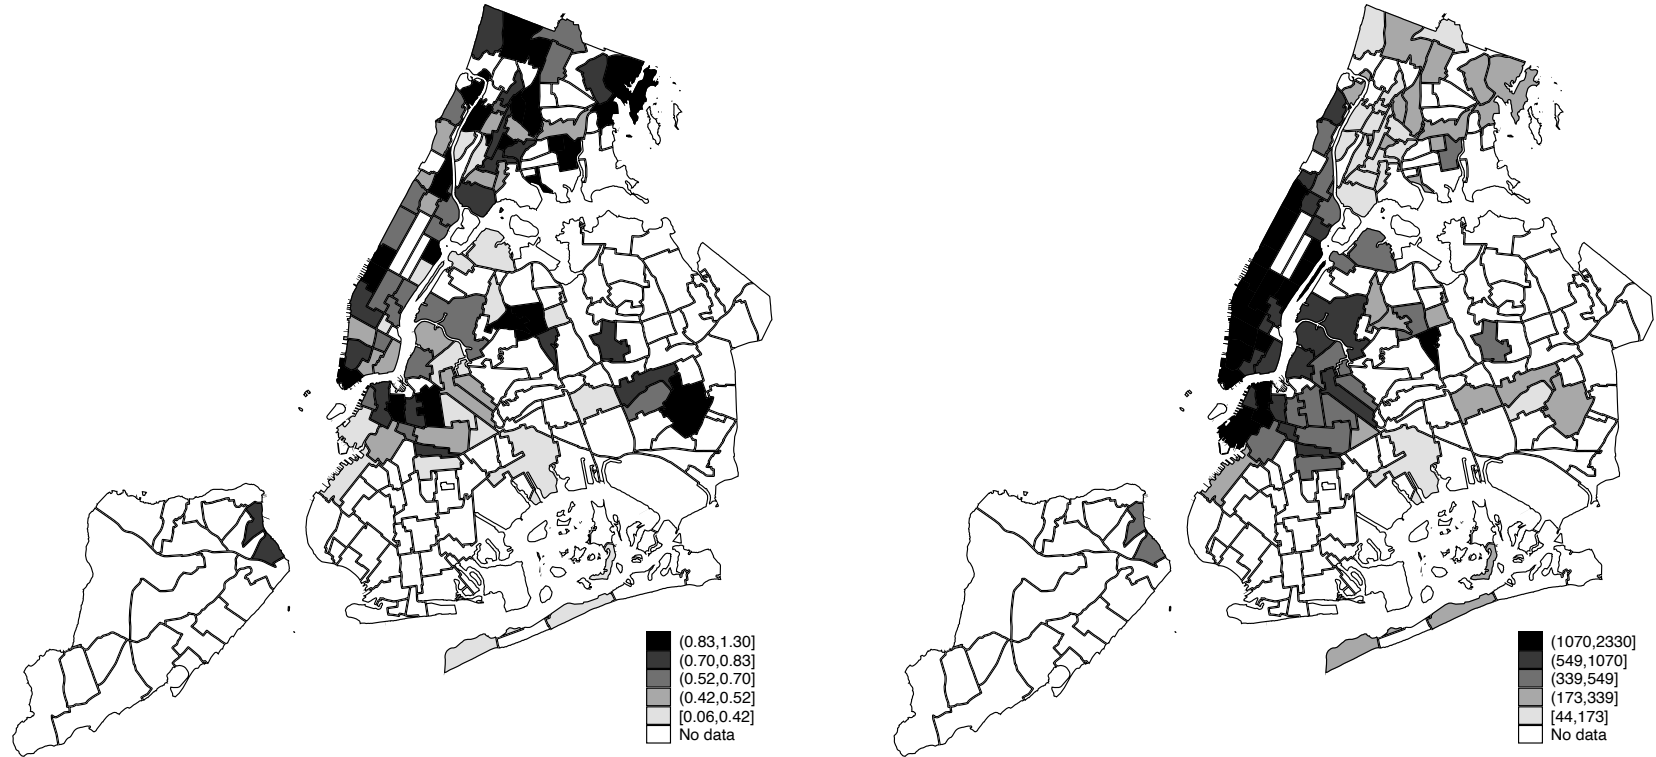


Notes: These figures, in Panels A and B, respectively display the standard errors for estimates by Neighborhood Tabulation Area of the supply response of inclusionary units with respect to the 421-a incentive and the marginal fiscal cost per inclusionary unit. Standard errors are computed by the cluster-bootstrap, clustering at the level of Neighborhood Tabulation Areas. See Figure 4.

Figure A15: Estimates by Neighborhood, with Neighborhood Fixed Effects

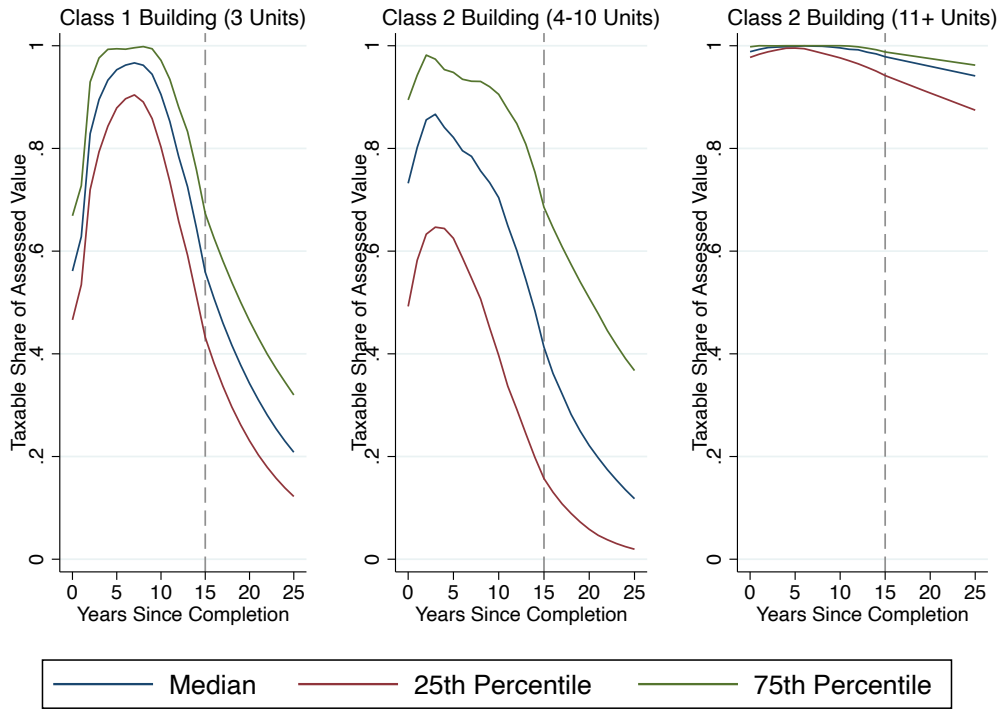
Panel A: Supply Response

Panel B: Marginal Fiscal Cost per Inclusionary Unit




Notes: These figures, in Panels A and B, respectively display estimates by Neighborhood Tabulation Area of the supply response of inclusionary units with respect to the 421-a incentive and the marginal fiscal cost per inclusionary unit. The units are percentage-point changes the participation rate per 1 p.p. increase in the 421-a incentive (in Panel A) and thousands of 2015 dollars (in Panel B). I estimate the models with lot- and block-level controls as well as Neighborhood Tabulation Area fixed effects, unlike in Figure 4.

Figure A16: Evolution of Estimated Taxable Shares over Building Lifecycle




Notes: This figure displays the evolution of the taxable share of the building’s assessed value. For each group of buildings, I display the median, 25th percentile, and 75th percentile share in each year since the building’s completion. Estimates of shares come from the statistical model of New York City assessment practices introduced in Appendix C. See Section 2 for background on the assessment process. The dashed vertical line at 15 years indicates the point beyond which values come purely from extrapolation.

Figure A17: Cost of Residential Property Tax Exemptions and Abatements in New York City



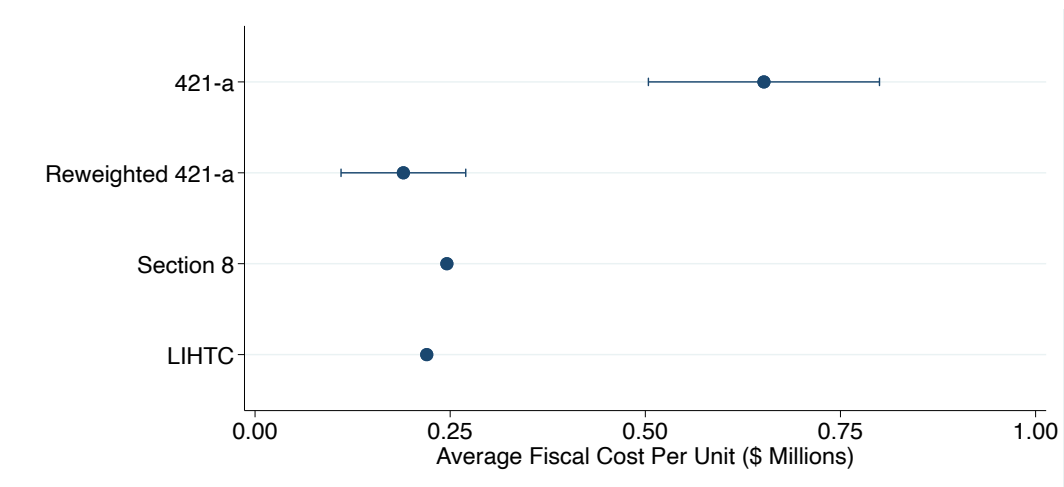
Notes: This figure compares the accounting cost of 421-a in fiscal year 2019 to all other residential and individual property tax exemptions and abatements in New York City. Source: Annual Reports of the New York City Property Tax, as published by the Division of Tax Policy, New York City Department of Finance.

Figure A18: Developer Profit Incidence Share by Neighborhood Tabulation Area




Notes: This figure displays the estimated share of the total fiscal cost of 421-a in each Neighborhood Tabulation Area that accrues to developer profit. Equation 11 formally defines this incidence share.

Figure A19: Do Differences in Neighborhood Characteristics Explain Cost Differences?




Notes: This figure plots actual and adjusted citywide average marginal fiscal costs of inclusionary units under 421-a, Section 8 voucher units, and LIHTC units. In the first two rows, I plot the actual and reweighted citywide average fiscal cost per inclusionary unit, following the procedure of DiNardo et al. (1996) to reweight units on their block characteristics. In the remaining rows, I plot the actual average fiscal costs for Section 8 and LIHTC units.

Figure A20: Marginal Value of Public Funds of 421-a by Neighborhood Tabulation Area (High-WTP Scenario)



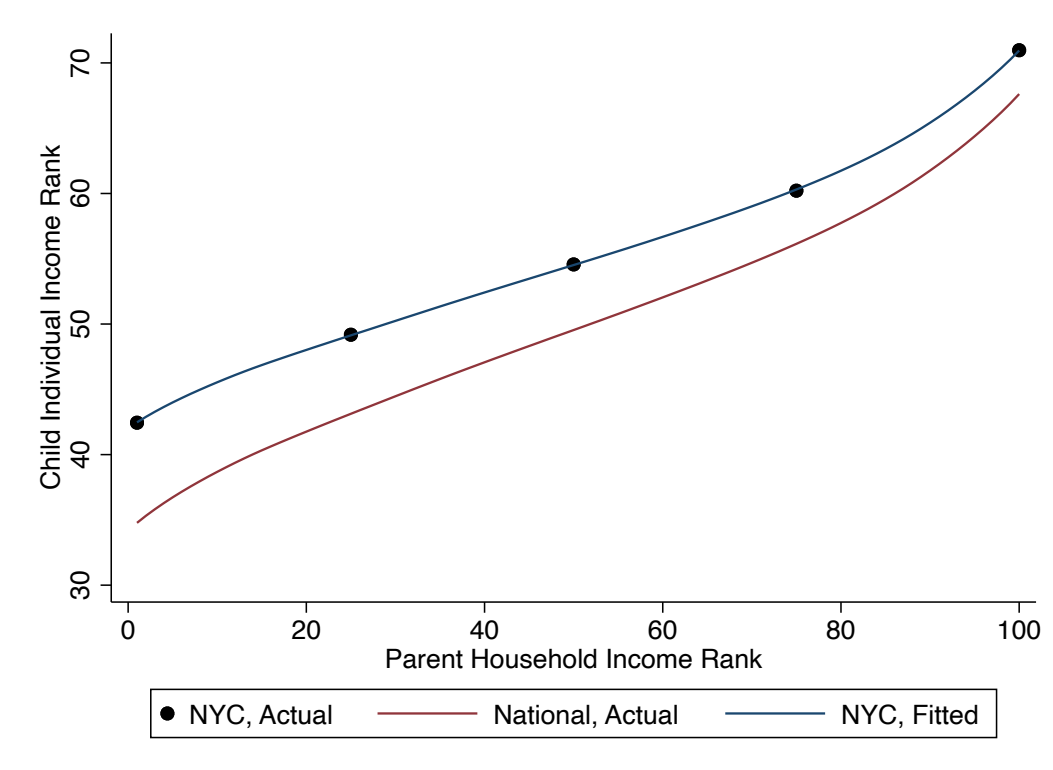
Notes: This figure depicts the marginal value of public funds (MVPF) for 421-a in each Neighborhood Tabulation Area. For the definition of the MVPF, see Section 9. In this high-WTP scenario, recipient households value inclusionary housing units at estimated developer breakevens. See Figure 5 for the estimates in the low-WTP scenario.

Figure A21: Marginal Value of Public Funds of 421-a by Neighborhood Tabulation Area (Close-Moves Scenario)



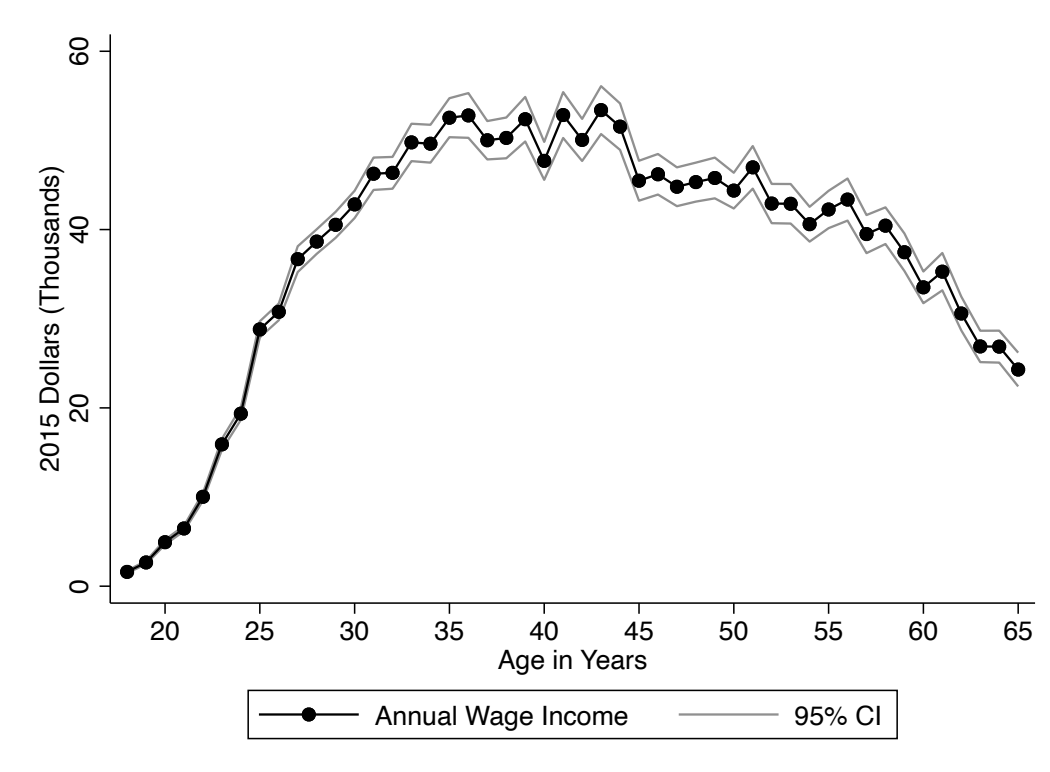
Notes: This figure depicts the marginal value of public funds (MVPF) for 421-a in each Neighborhood Tabulation Area. For the definition of the MVPF, see Section 9. In this close-moves scenario, I calibrate moving costs to match crowding-out estimates in Baum-Snow and Marion (2009) and Eriksen and Rosenthal (2010). This calibration yields an average move distance of 3.1 kilometers. I assume households value inclusionary housing as in the low-WTP scenario. See Figure 5 for the estimates in the main low-WTP scenario, which calibrates moving costs so that the average move distance is 10 kilometers.

Figure A22: Fitted Intergenerational Income Rank–Rank Function for New York City



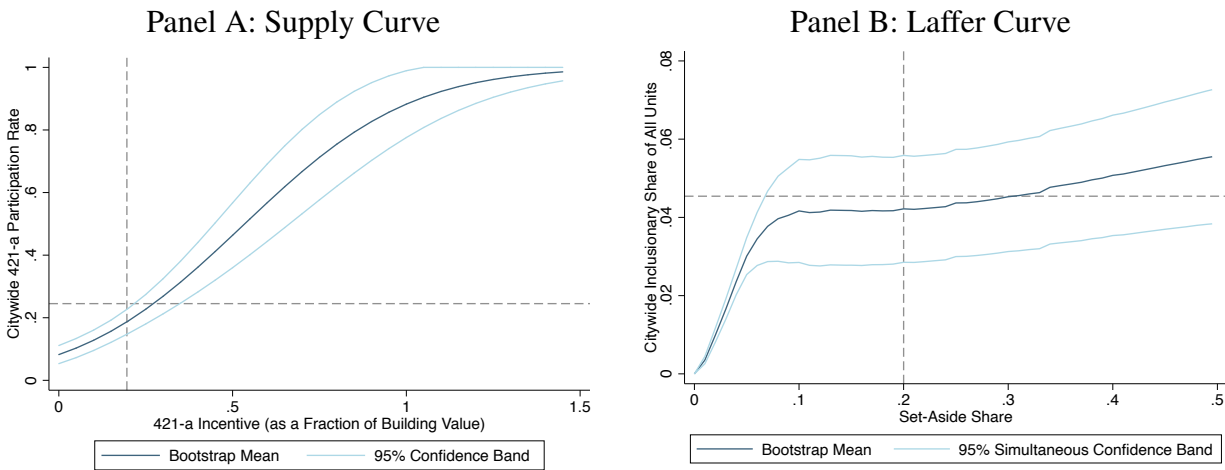
Notes: This figure displays, in solid points, the actual expected individual income ranks for children who grow up in New York City to native-born parents according to the income rank of their household. The red line shows the complete rank–rank function at the national level. The blue line presents the fitted values for the New York City rank–rank function. For further details, see Appendix C.

Figure A23: Age–Earnings Profile for New York City in 2015



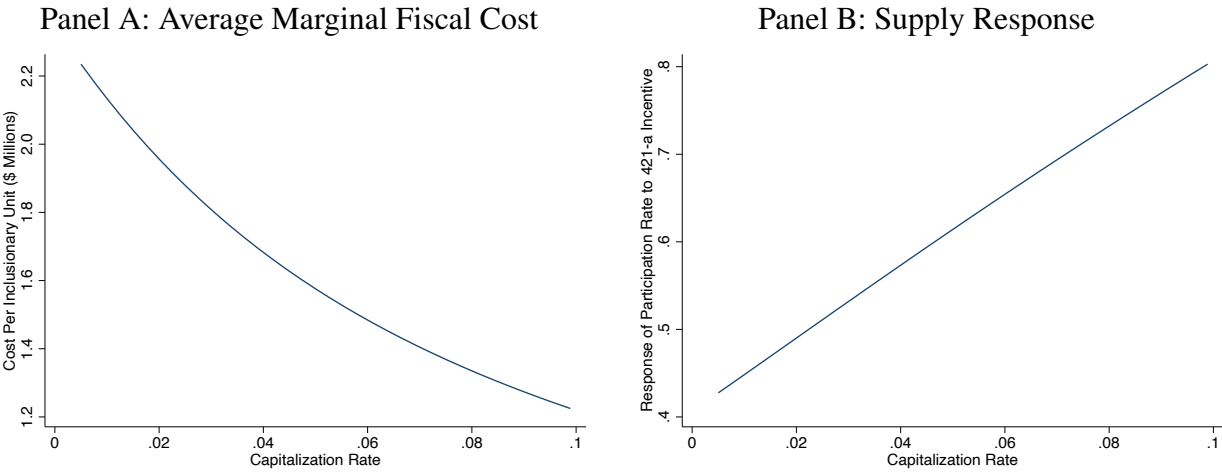
Notes: This figure displays the age–earnings profile for workers, ages 18 to 65, in New York City in the 5-year 2015 American Community Survey. All earnings are adjusted to constant 2015 dollars.

Figure A24: 421-a Participation Under Alternative Tax Incentives and Inclusionary Shares



Notes: In Panel A, I plot the estimated citywide share of units in buildings that would take up the 421-a exemption if offered a common incentive $\Delta\tau$. The dashed lines indicate the average incentive actually offered and the actual 421-a participation rate. In Panel B, I plot the estimated citywide share of units in all new residential buildings that are reserved for inclusionary tenants as a function of the set-aside share required of 421-a buildings. The dashed lines indicate the actual inclusionary share required of 421-a buildings and the actual inclusionary share of all units. In both panels, the light blue lines indicate bootstrapped 95-percent simultaneous confidence bands.

Figure A25: Citywide Estimates Under Alternative Discount Rate Assumptions

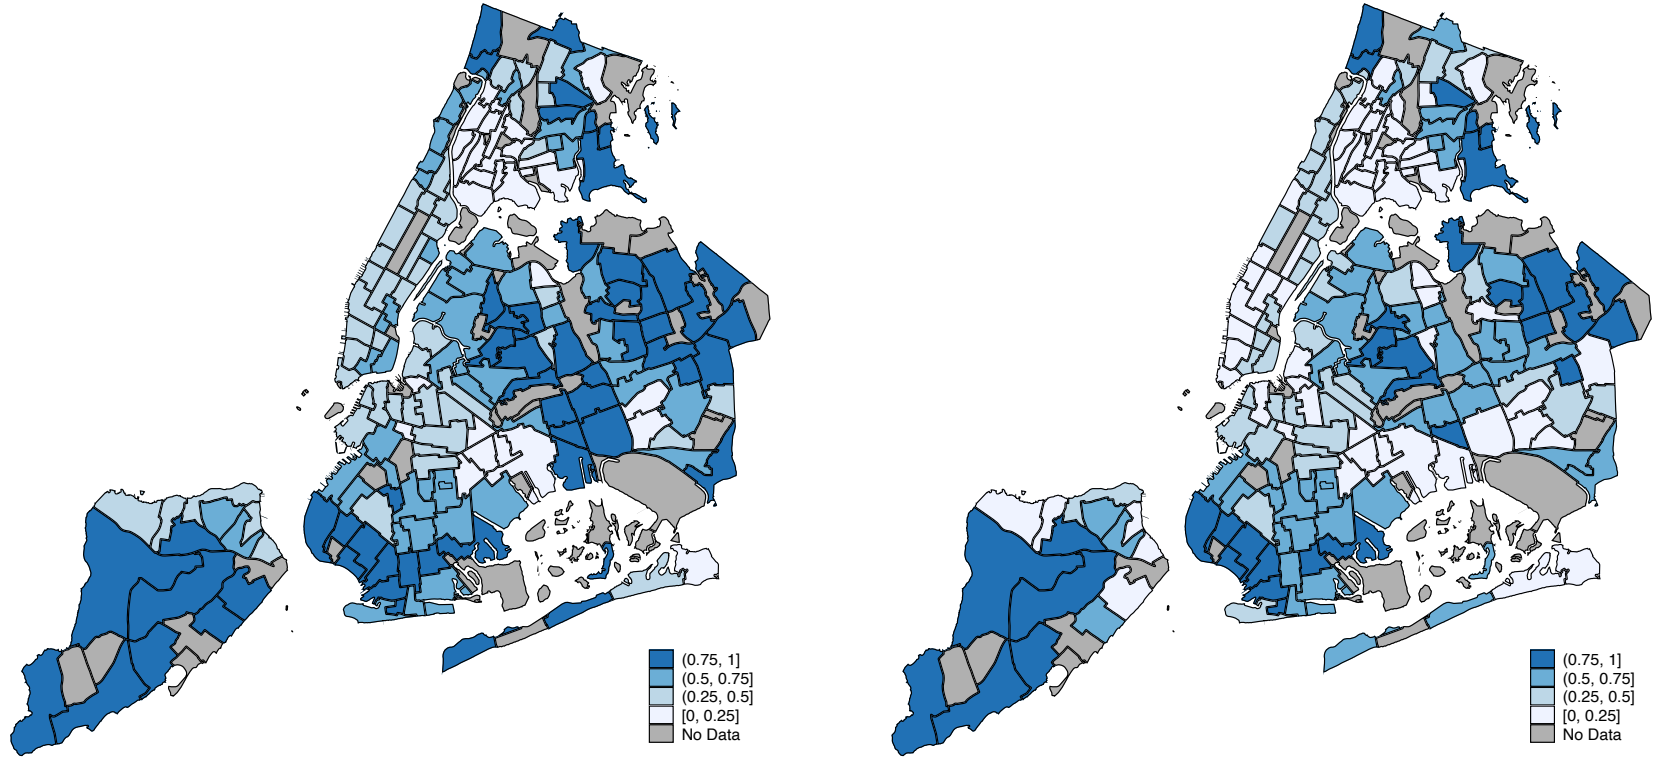


Notes: The annual capitalization rate is defined as the ratio annual after-tax rental income to building market value. The baseline calibration for ρ is 0.05. In Panel A, I plot the average marginal fiscal cost per inclusionary unit as a function of the annual discount rate. In Panel B, I plot the percentage-point increase in inclusionary units in response to a 1-p.p. increase in the average 421-a incentive as a function of the annual discount rate.

Figure A26: Ranking Neighborhoods by 421-a's Marginal Value of Public Funds

Panel A: Minimum Percentile Rank


Panel B: Maximum Percentile Rank



72

Notes: This figure presents the lower and upper bounds of the simultaneous confidence set on the percentile ranks of the neighborhood distribution of the marginal value of public funds (MVPF) spent on 421-a. I follow Mogstad et al. (2020) in constructing the simultaneous confidence set over ranks. To aid in comparison, the color scale is the same in Panels A and B.

Figure A27: Underassessment Ratio of Owner-Occupied Buildings



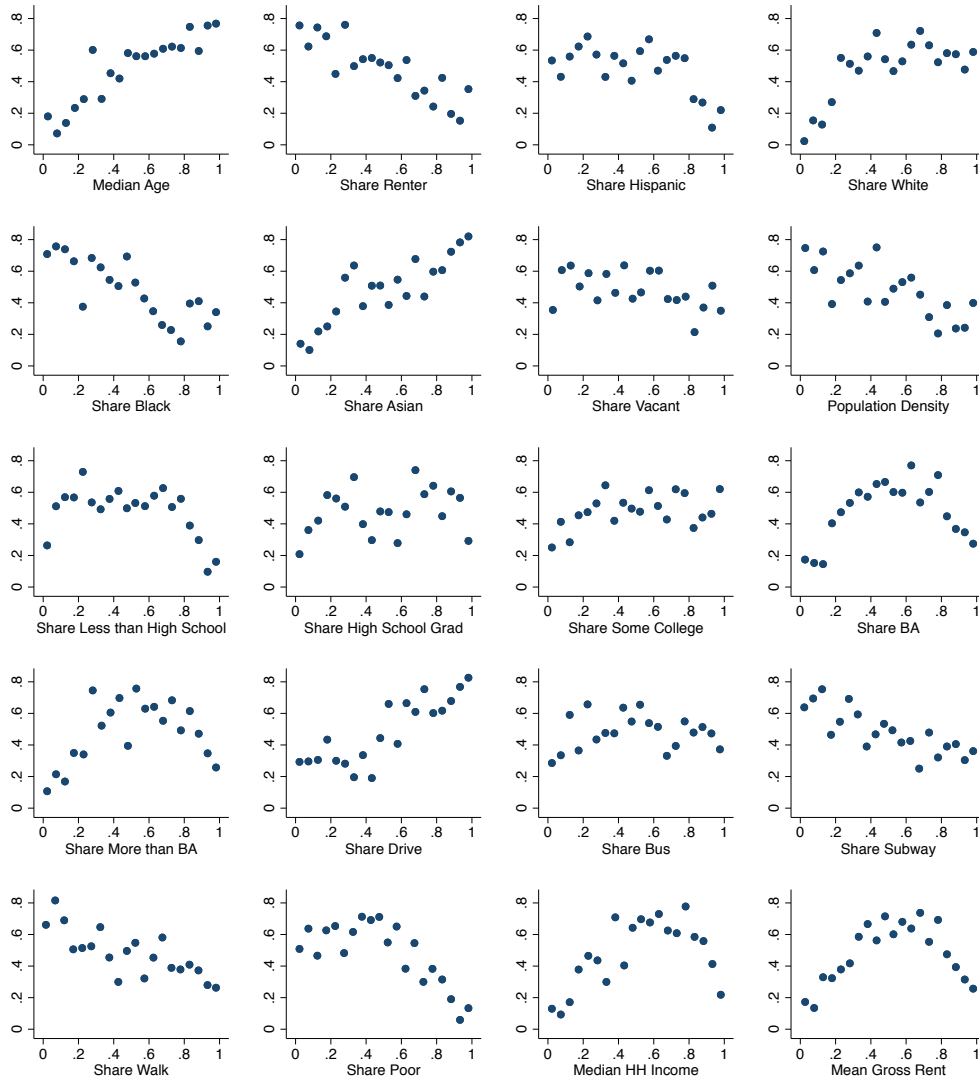
Notes: This figure is a histogram of the underassessment ratio, defined as the average assessed value per residential unit (as estimated by the New York City Department of Finance) divided by the average market value per residential (as estimated from market transactions). For more information, see Appendix C.

Figure A28: Case Study for Simulated Instrument



Notes: This figure maps the Manhattan neighborhood of East Harlem. It identifies the location of four buildings that were eligible for the 421-a tax exemption, along with local landmarks for geographic reference. Map data ©2021 Google.

Figure A29: Correlates of Neighborhood MVPF Percentile Rank



Notes: This figure presents binned scatterplots of the relationship between neighborhoods’ characteristics and their marginal values of public funds (MVPFs). Both neighborhood characteristics and MVPFs are transformed into percentile ranks.

Table A1: Calculating Tax Liabilities

Concept	Action
0. True Market Value	
1. Estimated Market Value	Estimate (0)
2. Actual Assessed Value	Apply assessment ratio to (1)
3. Transitional Assessed Value	Apply growth caps to (2)
4. Taxable Value	Apply exemptions to (3)
5. Tax Liability	Apply tax rate to (4)

Notes: This table outlines the steps performed by the New York City Department of Finance to calculate property tax liabilities and that I replicate in my calculator.

Table A2: Duration of the 421-a Exemption by Location, Year, and Inclusionary Housing Provision

Years	Location	Inclusionary Housing		
		None	Off-Site	On-Site
1985–2006	Manhattan GEA	0	10	20
	Manhattan non-GEA, South of 110th Street	10	10	20
	Neighborhood Preservation Program Areas	20	20	20
	All Other Areas	15	15	25
2006–2008	Manhattan GEA	0	10	20
	Manhattan non-GEA, South of 110th Street	10	10	20
	Greenpoint–Williamsburg GEA	0	15	20
	Neighborhood Preservation Program Areas	20	20	20
	All Other Areas	15	15	25
2008–2016	Expanded GEA	0	0	25
	All Other Areas	15	15	25

Notes: This table reports the duration in years of the 421-a exemption by the location of the development, the year of permit issuance, and whether the development includes offsite or onsite inclusionary housing. Appendix Figure A3 depicts the phase-out schedules for 421-a exemptions.

Table A3: Summary Statistics on 421-a Eligible Buildings, 2003–2015

	Does the Building Provide Onsite Inclusionary Housing?			
	Yes		No	
	Mean	Std. Dev.	Mean	Std. Dev.
<i>Panel A: Building Characteristics</i>				
Frontage (Feet)	70	84	35	91
Depth (Feet)	74	69	56	79
Floor Area (Sq. Ft.)	108,264	185,345	20,273	85,792
% Residential Area	91.0	12.5	94.6	11.9
% Condo	33.0	47.1	23.8	42.6
% In Historical District	0.7	8.3	0.6	7.9
Residential Max FAR	3.9	2.6	2.4	1.6
<i>Panel B: Area Characteristics</i>				
Educational Attainment				
% Less Than High School	19.2	12.7	22.6	13.8
% High School Graduate	21.3	11.2	26.2	12.1
% Some College	18.3	8.1	20.3	8.4
% College Graduate	25.3	13.7	19.7	13.1
% More than College	15.9	13.8	11.1	10.7
Commute Modes				
% Drive	15.3	10.5	23.7	15.2
% Bus	9.0	10.2	10.1	9.0
% Subway	52.9	18.9	49.2	18.7
% Walk	12.8	12.8	9.2	10.0
Demographic Composition				
% Non-Hispanic White	31.1	29.6	25.2	29.5
% Non-Hispanic Black	22.4	24.0	26.9	29.9
% Hispanic	39.1	27.3	34.6	26.3
% Asian	11.0	16.2	12.7	19.0
Median Age	31.2	7.8	32.4	6.5
Median Household Income	\$65,840	\$41,697	\$56,646	\$30,951
Mean Annual Gross Rent	\$19,568	\$6,768	\$17,332	\$5,014
% Renter (Units)	82.8	16.2	74.4	18.4
% Vacant (Units)	13.9	17.5	10.9	10.1

Notes: This table reports means and standard deviations of financial and buildings characteristics of the dataset of developments built from 2003 to 2015 that were eligible for 421-a. See the text of the paper for variable definitions.

Table A4: Developer Participation Responses to 421-a Incentive

	(1)	(2)	(3)	(4)
	LPM		Logit	
	Borough FE	NTA FE	Borough FE	NTA FE
421-a Incentive	0.557*** (0.055)	0.557*** (0.056)	5.184*** (0.538)	5.198*** (0.648)
<i>N</i>	11,669	11,666	11,669	7,465
Clusters	179	176	179	82
<i>R</i> ² (within)	0.167	0.222	n.a.	n.a.
Pseudo- <i>R</i> ²	n.a.	n.a.	0.233	0.292
Marginal Effect			0.183*** (0.033)	0.261*** (0.028)

Notes: This table reports estimates from Equation 5 of the participation effects of the present-value tax-rate differential $\Delta\tau_i$ from providing onsite inclusionary housing under the 421-a exemption. All specifications include fixed effects for the borough and year of initial permit issuance. Marginal effects reflect the percentage-point increase in the onsite affordable housing participation rate per percentage-point increase in the tax incentive. Standard errors are clustered at the neighborhood level. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A5: Impacts of Developer Specialization on 421-a Participation and Lot Ownership

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Fixed Effects:	None		Borough		Neighborhood	Tabulation Area	
Controls:	None	None	Lot	Lot & Block	None	Lot	Lot & Block
<i>Panel A: Does Developer 421-a Experience Affect Participation?</i>							
421-a Incentive	5.12*** (0.31)	5.27*** (0.56)	5.71*** (0.68)	6.22*** (0.73)	5.35*** (0.68)	6.09*** (0.79)	6.22*** (0.76)
Dev. % 421-a	2.02*** (0.35)	1.89*** (0.34)	1.56*** (0.36)	1.57*** (0.40)	1.88*** (0.41)	1.46*** (0.45)	1.41*** (0.45)
<i>N</i>	10,683	10,683	10,665	10,658	6,559	6,545	6,540
Clusters	178	178	178	178	79	79	79
Cost Savings	0.395*** (0.072)	0.359*** (0.076)	0.273*** (0.066)	0.252*** (0.065)	0.352*** (0.081)	0.240*** (0.073)	0.226*** (0.070)
<i>Panel B: Is Lot Ownership Endogenous to 421-a?</i>							
Dev. % 421-a	0.003 (0.025)	0.073*** (0.021)	0.001 (0.024)	-0.007 (0.023)	0.040* (0.021)	-0.006 (0.022)	-0.014 (0.023)
<i>N</i>	10,683	10,683	10,674	10,667	10,675	10,666	10,659
Clusters	178	178	178	178	170	170	170

Notes: This table extends, in Panel A, the specification of Equation 5 to include a measure of developers' 421-a specialization, the leave-out share of buildings they enter into 421-a, weighted by the number of residential units. In Panel B, I regress the 421-a incentive $\Delta\tau_i$ on the measure of developer specialization. Standard errors are clustered at the neighborhood level. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A6: Bootstrap Estimates of Citywide Supply Response and Marginal Fiscal Cost, Offsite versus Onsite Inclusionary Units

	Offsite	Onsite
	(1)	(2)
Supply Response	0.04*** (0.01)	0.61*** (0.07)
Marginal Fiscal Cost	\$2,342,047*** (1,484,622)	\$1,180,016*** (316,382)

Notes: This table presents citywide supply responses and citywide average marginal fiscal costs, for both offsite and onsite inclusionary units. All results come from the multinomial logit specification of Equation 20. In the first row, I report estimates of the citywide supply response of inclusionary units to changes in buildings' 421-a incentive. In the second row, I report marginal fiscal costs per inclusionary unit per year. In the left column, my estimates are for offsite inclusionary units, whereas in the right column, my estimates are for onsite inclusionary units. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A7: Developer Response to 421-a Incentive: Borough and Year Fixed Effects

	(1)	(2)	(3)	(4)
Fixed Effects:	None		Borough	
Controls:	None	None	Lot	Lot & Block
Boroughs:				
Bronx		0.611* (0.362)	2.036*** (0.552)	1.629*** (0.507)
Brooklyn		-0.033 (0.294)	0.994** (0.477)	0.885** (0.451)
Queens		-0.527 (0.451)	0.625 (0.593)	0.523 (0.529)
Staten Island		-0.434 (1.030)	0.538 (0.988)	0.483 (0.959)
Permit Years:				
2004		-0.138 (0.236)	-0.050 (0.249)	-0.065 (0.252)
2005		-0.050 (0.220)	0.016 (0.244)	-0.036 (0.203)
2006		-0.068 (0.230)	-0.077 (0.232)	-0.141 (0.231)
2007		0.016 (0.201)	-0.182 (0.215)	-0.309 (0.191)
2008		-1.320** (0.586)	-1.907*** (0.614)	-2.253*** (0.528)
2009		0.672 (0.593)	0.105 (0.536)	-0.033 (0.522)
2010		0.612 (0.485)	0.005 (0.473)	-0.147 (0.456)
2011		0.693 (0.551)	0.093 (0.585)	-0.101 (0.613)
2012		0.425 (0.515)	-0.276 (0.556)	-0.498 (0.544)
2013		-0.044 (0.471)	-0.638 (0.496)	-0.840* (0.460)
2014		-0.243 (0.530)	-0.868 (0.593)	-0.997* (0.577)
2015		-0.601 (0.466)	-1.328*** (0.459)	-1.644*** (0.427)
<i>N</i>	11,666	11,666	11,645	11,638
Clusters	179	179	179	179

Notes: See notes to Table 3. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A8: Developer Response to 421-a Incentive: Lot and Block Controls

	(1)	(2)	(3)	(4)
Fixed Effects:	None		Borough	
Controls:	None	None	Lot	Lot & Block
Log Land Value			0.277*** (0.107)	0.290** (0.116)
Log Lot Area			0.285** (0.135)	0.347*** (0.131)
Residential Maximum FAR			0.069 (0.053)	0.051 (0.057)
In Special District			-0.524 (0.494)	-0.400 (0.496)
In R10-DA IHP			0.511** (0.260)	0.447* (0.253)
In Liberty Zone			0.959 (0.612)	1.263** (0.538)
Lot Type:				
Entire Block			-0.931* (0.549)	-1.099* (0.589)
Corner			-0.662** (0.291)	-0.846*** (0.270)
Through			-1.314*** (0.454)	-1.546*** (0.492)
Inside			-0.805*** (0.220)	-1.017*** (0.210)
% Less than HS				0.493 (1.166)
% HS Graduate				2.220** (1.012)
% Some College				1.118 (0.963)
% College Graduate				1.858 (1.220)
% Car Commuter				-2.899* (1.527)
% Bus Commuter				-0.346 (1.735)
% Subway Commuter				-0.810 (1.270)
% Walk Commuter				-1.107 (1.399)
Log Median HH Income				-0.104 (0.294)
Median Age				-0.042** (0.020)
% Rental Units				1.270*** (0.488)
% Vacant Units				0.558 (0.555)
% Hispanic				0.279 (0.417)
% Non-Hispanic White				-0.898* (0.537)
% Non-Hispanic Black				-0.419 (0.565)
% Asian				0.431 (1.106)
N	11,666	11,666	11,645	11,638
Clusters	179	179	179	179

Notes: See notes to Table 3. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A9: Incidence Analysis of 421-a: Actual Versus Incidence-Minimizing Counterfactual

	(1) Actual	(2) Counterfactual
Developer Incidence Share	0.456*** (0.117)	0.356*** (0.089)
<i>N</i>	12,146	12,146
Clusters	179	179

Notes: This table presents estimates of the share of the total citywide fiscal cost of 421-a captured by developers as profits. The left column contains the estimated developer incidence share under the actual 421-a program, whereas the right column contains the developer incidence share under a counterfactual program that minimizes developer incidence subject to a constraint that it achieves the same take-up rate. Standard errors are computed by a cluster-bootstrap at the level of Neighborhood Tabulation Areas. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A10: Developer Participation Responses to Variation in 421-a Incentive: Simulated Instruments GMM Results

	(1)	(2)	(3)	(4)
	None	Borough Fixed Effects		
	No Controls	No Controls	Lot Controls	Lot & Block Controls
421-a Incentive	5.419*** (0.384)	6.506*** (0.986)	5.525*** (1.009)	6.588*** (1.060)
<i>N</i>	11,460	11,460	11,448	11,445
Clusters	178	178	178	178
Std. Dev. of e_i	0.335*** (0.024)	0.279*** (0.042)	0.328*** (0.060)	0.296*** (0.053)

Notes: This table reports estimates from Equation 6 of the participation effects of the present-value tax-rate differential $\Delta\tau_i$ from providing onsite inclusionary housing under the 421-a exemption. All specifications include fixed effects for the borough and year of initial permit issuance. For comparison, see Columns 1–4 of Table 3, which estimates Equation 5, which is otherwise identical to Equation 6 but does not instrument for $\Delta\tau_i$ using the simulated instruments approach. Standard errors are clustered at the neighborhood level. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A11: Bootstrap Estimates of Citywide Supply Response and Marginal Fiscal Cost, Alternative Specifications

	RE Logit	Probit	LPM
	(1)	(2)	
Supply Response	0.57*** (0.06)	0.53*** (0.05)	0.59*** (0.06)
Marginal Fiscal Cost	\$1,751,636*** (318,971)	\$2,480,480*** (419,883)	\$1,627,548*** (339,004)

Notes: This table presents citywide supply responses and citywide marginal fiscal costs. In the first row, I report estimates of the citywide supply response of onsite inclusionary units to changes in buildings' 421-a incentive. In the second row, I report marginal fiscal costs per inclusionary unit per year. In Column 1, 2, and 3, my estimates are from respectively a random-effects logit specification, and a linear probability model (LPM) specification. The random effects are Neighborhood Tabulation Area intercepts. All specifications assume building features are exogenously determined and include fixed effects for borough and year as well as lot and block controls. Standard errors are computed by a cluster-bootstrap at the level of Neighborhood Tabulation Areas. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A12: WTP and Cost by Quartile of MVPF Distribution

	Quartile of Neighborhood MVPF Distribution			
	1	2	3	4
WTP	\$45,269 (\$1,222)	\$96,934 (\$956)	\$72,655 (\$660)	\$58,288 (\$487)
Net Cost	\$217,088 (\$4,842)	\$204,675 (\$2,248)	\$79,641 (\$861)	\$19,139 (\$466)
Gross Cost	\$205,113 (\$4,958)	\$199,533 (\$2,308)	\$87,497 (\$874)	\$42,818 (\$484)

Notes: This table reports, by quartile of the distribution of neighborhoods ranked by the estimated marginal value of public funds (MVPF) for 421-a, the willingness to pay (WTP) and net and gross cost (net and cost of fiscal externalities). All values are computed on a per-household, not per-unit, basis (see Appendix D). Neighborhoods are ranked and grouped once on their actual MVPF estimate, and standard errors are computed by a cluster-bootstrap at the level of Neighborhood Tabulation Areas. Therefore, the standard errors incorporate three sources of sampling variance: reported standard errors from the Opportunity Atlas mobility estimates (Chetty et al., 2018), standard errors in my estimates of marginal cost per unit, and misclassification of neighborhoods into quartiles due to uncertainty in the estimated MVPFs. This analysis is therefore responsive to concerns in Mogstad et al. (2020) about the implications of sampling variance for the analysis of Moving to Opportunity and similar policies. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A13: Robustness to Lot Area Threshold

	(1)	(2)	(3)
	Lot Area Threshold		
	20%	10%	30%
No Controls	6.506*** (0.986)	6.681*** (1.133)	6.525*** (0.965)
Lot Controls	5.525*** (1.009)	5.864*** (1.197)	5.455*** (0.989)
Block Controls	6.662*** (1.102)	6.927*** (1.307)	6.649*** (1.062)
Lot & Block Controls	6.136*** (1.100)	6.577*** (1.330)	6.091*** (1.081)

Notes: This table examines the robustness of the GMM logit coefficient estimates for the effect of the 421-a tax incentive. The specification estimated is always Equation 6, changing the construction of the simulated instrument. Column 1 reports baseline estimates, in which we use the simulated instrument where comparable lots must be within 20 percent of the lot area as the actual lot. Column 2 tightens this comparison to lots within 10 percent of the actual lot area. Column 3 loosens the comparison to 30 percent. Each row presents a specification within a given set of controls. All specifications include fixed effects for the borough and year of initial permit issuance. Standard errors are clustered at the neighborhood level. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

B Model Appendix

This appendix extends the main model presented in Section 4 in two respects. First, I introduce an intensive margin: Developers choose the number of units n_i to build. Second, I allow developers to be monopolistically competitive rather than, conditional on λ , taking market rents as given.

B.1 Preliminaries

A housing developer that owns a vacant parcel of land i faces two choices: (1) how many units of housing to build and (2) whether to participate in the 421-a program. Participating developers must reserve units for low-income tenants who pay below-market rents, but in return, they pay a lower property tax rate. Let $\Delta\tau_i \geq 0$ denote the present value of the tax savings as a share of building value. Developers are monopolistically competitive, reflecting imperfect substitutability between parcels. Conditional on its participation choice, a developer builds as many units maximize its profit from the parcel, which is the difference of the present value of rental income and the cost of construction. Taxes are in proportion to profit.

Let n_i be the number of units built on i . If the developer participates in 421-a, it must reserve a fraction λ of these units for inclusionary tenants. Inclusionary tenants pay a fixed rent r to the developer. In market-rate units, the rent $m_i(n_i, \lambda)$ is a function of parcel characteristics, the reserved share of units λ , and the total number of units n_i . For example, if $\partial m_i / \partial \lambda < 0$, then reserving units for inclusionary tenants reduces the willingness to pay of market-rate tenants. The average rent of units in a parcel is a weighted average of the fixed and market rents:

$$p_i(n_i, \lambda) = \lambda r + (1 - \lambda)m_i(n_i(\lambda), \lambda),$$

where $n_i(\lambda)$ is the developer's choice of the number of units conditional on an inclusionary share λ . It will be convenient to define the log-difference in average rent between the participation and nonparticipation options:

$$\Delta \log p_i \approx \frac{p_i(n_i(0), 0) - p_i(n_i(\lambda), \lambda)}{p_i(n_i(0), 0)} = \lambda \mu_i + (1 - \lambda) \delta_i(\lambda), \quad (13)$$

where I define

$$\mu_i = \frac{m_i(n_i, 0) - r}{m_i(n_i, 0)} \quad \text{and} \quad \delta_i(\lambda) = \frac{m_i(n_i(0), 0) - m_i(n_i(\lambda), \lambda)}{m_i(n_i(0), 0)}.$$

The term μ_i reflects the inclusionary discount: It is the difference in rent between an inclusionary unit and a market-rate unit in the nonparticipation counterfactual. The term $\delta_i(\lambda)$ reflects the disamenity: It is the discount on rent for market-rate tenants that the developer must offer to market-rate tenants compensate for the presence of a share λ of inclusionary tenants.¹⁷

B.2 The Developer Participation Decision

The developer participates in 421-a if its profit under participation exceeds its profit under nonparticipation:

$$D_i = 1[\Delta\pi_i \geq 0].$$

Let the developer's pre-tax profit function for parcel i be $\pi_i(p_i, n_i(p_i))$, where $n_i(p_i)$ is the demand for units in parcel i . This difference in profits is

$$\Delta\pi_i = (1 + \Delta\tau_i)\pi_i(p'_i, n_i(p'_i)) - \pi_i(p_i, n_i(p_i)), \quad (14)$$

reflecting the trade-off between a lower tax rate and lower pre-tax profit, as $\Delta\tau \geq 0$ but $\pi_i(p_i, n_i(p_i)) \geq \pi_i(p'_i, n_i(p'_i))$. This specification contains two important assumptions. First, insofar as λ only enters $\pi_i(\cdot)$ via p_i and n_i , the developer's marginal costs cannot depend upon whether units are market-rate or inclusionary. Second, I ignore any fixed costs of participation in 421-a, although this assumption is trivial to relax.

To obtain $\pi_i(p'_i, n_i(p'_i))$, I take a second-order Taylor expansion of the profit function around the equilibrium profit under nonparticipation:

$$\pi(p', n(p')) = \pi(p, n(p)) + \left[\frac{\partial\pi}{\partial p} + \frac{\partial\pi}{\partial n} \frac{\partial n}{\partial p} \right] \Delta p + \frac{1}{2} \left[\frac{\partial^2\pi}{\partial p^2} + \frac{\partial^2\pi}{\partial n \partial p} \frac{\partial n}{\partial p} + \frac{\partial\pi}{\partial n} \frac{\partial^2 n}{\partial p^2} \right] (\Delta p)^2. \quad (15)$$

By the Envelope Theorem, developers set the number of units to equate the marginal revenue

¹⁷However, this intuition ignores an effect of λ on m_i which flows through reoptimization of n_i .

gain from the next unit with the marginal revenue lost from lower rents on inframarginal units. Therefore, at $\lambda = 0$ and the optimal rent under nonparticipation p^* , it follows that

$$\frac{\partial \pi(p^*, n(p^*))}{\partial p} + \frac{\partial \pi(p^*, n(p^*))}{\partial n} \frac{\partial n(p^*)}{\partial p} = 0,$$

and thus the first-order term in Equation 15 equals zero. By Hotelling's lemma, and given constant-elasticity demand for units, Equation 15 simplifies to

$$\pi(p', n(p')) = \pi(p, n(p)) - \varepsilon_D (p - mc) n(p) (\Delta \log p)^2, \quad (16)$$

letting ε_D denote the price elasticity of demand. I substitute Equation 16 into Equation 14, and then I divide both sides by $\pi_i(p_i, n_i(p_i))$, yielding

$$\frac{\Delta \pi_i}{\pi_i(p_i, n_i(p_i))} = [1 + \Delta \tau_i] \left[1 - \varepsilon_D \frac{(p_i - mc_i) n_i(p_i)}{\pi_i(p_i, n_i(p_i))} (\Delta \log p_i)^2 \right] - 1. \quad (17)$$

Given a marginal-cost function with constant elasticity ε_S ,¹⁸ and using logarithmic approximations, Equation 17 simplifies to

$$\Delta \log \pi_i = \Delta \tau_i - \frac{\varepsilon_D + \varepsilon_S \varepsilon_D}{1 + \varepsilon_S \varepsilon_D} (\Delta \log p_i)^2.$$

As $\Delta \pi_i \geq 0$ if and only if $\Delta \log \pi_i \geq 0$, the 421-a participation condition is

$$\Delta \tau_i \geq \frac{\varepsilon_D + \varepsilon_S \varepsilon_D}{1 + \varepsilon_S \varepsilon_D} [\lambda \mu_i + (1 - \lambda) \delta_i(\lambda)]^2, \quad (18)$$

using the expression for $\Delta \log p_i$ in Equation 13, where μ_i and $\delta_i(\lambda)$ are the inclusionary discount and disamenity on market-rate units.

In deciding whether to participate in 421-a, a developer compares the tax savings (the left-hand

¹⁸Consider a total cost function $c(q) = Aq^{1+\varepsilon_S}$. The relationship between the average and marginal cost functions is $ac(q) = \frac{mc(q)}{1+\varepsilon_S}$. Therefore, by the Lerner markup rule,

$$\frac{p - mc}{p - ac} = \frac{p - mc}{p - mc/(1 + \varepsilon_S)} = \frac{(1 + \varepsilon_S)/\varepsilon_D}{1/\varepsilon_D + \varepsilon_S} = \frac{1 + \varepsilon_S}{1 + \varepsilon_S \varepsilon_D}.$$

side of the inequality in Equation 3) and the impact of the lower average rent on its pre-tax profit (the right-hand side of the inequality in Equation 3). The profit loss due to the lower average rent is greater when the demand for units is more elastic (higher ϵ_D) or the reserved share of units λ , the inclusionary discount μ_i , or the disamenity $\delta_i(\lambda)$ is larger. As the price elasticity of demand $\epsilon_D > 1$, a higher price elasticity of supply ϵ_S reduces the profit impact of the rent reduction.

It is worth comparing Equation 18 to the paper's Equation 3. In Equation 3 the rent loss enters as a first-order term, whereas the rent loss is second-order in Equation 18. Nevertheless, in both models, a control vector x_i of detailed lot and block characteristics can potentially explain variation in forgone rental income. It is therefore not of much consequence for the empirical analysis whether we conceive of developers as having an intensive margin and market power.

C Data Appendix

C.1 BIN-to-BBL Crosswalk

The relevant date at which a building's 421-a eligibility is determined is recorded on its initial new-building (NB) construction permit, held by the New York City Department of Buildings. However, these permits do not contain accurate BBL identifiers, so considerable effort is required to match building permits to tax records. Here I outline the steps I took to ensure that my BIN-to-BBL mapping is as complete as possible.

I first use the address file of the Property Address Directory (PAD) of the New York City Department of City Planning, as explained in Section 5. This file allows me to match about 41,000 permits to tax records but provides an incomplete crosswalk of BINs. These match failures could be for two broad reasons: abandoned construction or incompleteness in the crosswalk. If construction is abandoned before completion, then a permit never yields a matching tax lot.

I then supplement the PAD. I identify all BINs that are not matched to BBLs in the PAD, and using these BINs, I obtain the borough, city block, and string address from their permits. Using regular expressions, I clean the string addresses to match their format in PLUTO. I use these fields to then match this otherwise-unmatched subset of permits to PLUTO. This step resolves matching issues for approximately 2,500 permits. The work, however, is not finished: About 1,400 permits still do not have matching BBLs.¹⁹

I reviewed and matched each of these permits to tax lots by hand. For about 600 of these permits, I found tax lots that were, in my judgment, a convincing match. To hand-match permits, I used "ZoLa," a project of the New York City Department of City Planning that digitizes the city's land-use maps,²⁰ as well as Google Street View (GSV). For a given unmatched permit, I searched for its string address in ZoLa/GSV. In general, I found that most cases of unmatched permits reflected a re-addressing of the building. For example, if a building on a corner or through lot

¹⁹I use another Department of City Planning crosswalk, the BBL file of the Property Address Directory, to collapse owner-occupied units up to the building level.

²⁰Available at zola.planning.nyc.gov.

moved its main entrance from one side of the building to another, then the address of the building as stated on the permit did not match the address of the building in tax records. However, it is straightforward to spot this case of match failure on ZoLa/GSV. A similar match failure occurs when the street numbers a developer expects to use at the beginning of construction (and therefore states on the permit) change by the end of construction. This case is particularly common for townhouses in the Outer Boroughs, where developers appear to sometimes change their plans during construction as to how many units to develop. GSV was particularly useful as a confirmation device: In much of New York City, GSV photography begins as early as 2007. This means it is possible to confirm abandoned construction—notably, during the 2008 financial crisis—and also to confirm that a building was built at the time a permit suggests. For most unmatched BINs, I can thus be confident that these do not reflect completed buildings elsewhere.

C.2 Property Tax Liability Calculator

Appendix Table A1 lays out the steps to calculate a building’s tax liability from its market value. In constructing and explaining of my calculator, I follow two sources closely: publicly-available guides from the Department of Finance and the 2020 report of the New York City Advisory Commission on Property Tax Reform, a non-partisan official body that has been delegated the responsibility to study and recommend changes to the property tax code.

Step 1. The Department of Finance estimates market value. Assessment methods vary according to the number of units in a building but are performed using mass-appraisal software, as in other U.S. cities. Throughout, I assume that developers treat these market-value estimates as unbiased, except in one case I discuss below. Per the Commission, these methods are consistent with best practices in property tax assessment.

New York City assessment practices deviate from these standards for condominiums. In particular, condominiums are valued “as if” they were rentals: The Department of Finance predicts the value of a condominium using its characteristics but having estimated its valuation model on only rental buildings. This approach produces a downward bias in estimated market values of con-

dominiums, most acutely at the top of the quality distribution. I assume that developers are aware of this bias, and as I explain in Appendix C, I account for it by estimating condominium market values from a spatial regression on condominium transaction prices, and then adjusting the tax rate for this undervaluation.

Step 2. If the building has three or fewer residential units, the estimated market value is multiplied by 0.06 to obtain the assessed value. If the building has four or more units, the estimated market value is multiplied by 0.45. This system of “fractional assessment,” introduced in a 1981 tax reform, formalized a pre-existing but discretionary practice of underassessment.²¹

Step 3. The 1981 reform also established caps on annual increases in assessed value. These caps vary with the number of units in a building.²² These caps create distortions in the “transitional” assessed value of buildings that vary with building age and across neighborhoods. As detailed in Appendix C.6, I use panel data on assessments from 2010 to 2018 to estimate the time path of undervaluation that a developer could have forecasted for a building. Over time, most buildings’ transitional assessed values deteriorate markedly relative to their actual assessed values.

Step 4. I apply any exemptions to taxable value that appear on the property’s assessment roll. By implication, in my baseline results, I take as predetermined developers’ claiming behavior for all exemptions other than 421-a. Insofar as developers would pursue other tax exemptions if they did not accept 421-a, this approach will overstate the value of the tax incentive for those who accept it. This measurement issue has two competing effects on the empirical analysis. Overstating fiscal

²¹For a political history of the New York City property tax code, see the series of articles by Ethan Geringer-Sameth, “An Old, Unfair System: New York City’s Property Tax Conundrum,” *Gotham Gazette*, 1–5 August 2019.

²²For buildings with three or fewer units, assessed value cannot increase by more than 6 percent per year, nor more than 20 percent over a five-year interval. For buildings with four to ten units, assessed value cannot increase by more than 8 percent per year, nor more than 30 percent over a five-year interval. For buildings with 11 or more units, changes in assessed value not capped but instead phased in uniformly in percentage terms over five-year intervals.

costs among participants may cause me to understate developer supply responses and overstate marginal costs. On the other hand, it also makes the tax savings more informative of participation choice, which would induce the opposite bias. Section 7 further discusses this measurement issue.

Step 5. New York City sets different property tax rates for buildings with three units or fewer versus those with four or more units. Tax rates usually change in each fiscal year. Appendix Figure A4 plots these tax rates since 1981. I assume that, in making 421-a participation decisions, developers' expected future statutory rates equal the average statutory rate in the year of filing their initial permit. A building's tax liability is the product of its taxable value and the appropriate statutory tax rate.

C.3 Additional 421-a Policy Details

This appendix subsection reviews further provisions of the 421-a exemption that I incorporate into estimation but which are ultimately of secondary importance to its design and impacts. The relevant statute is in Article 4, Title 2 of the New York State Real Property Tax Law. See also Chapter 6, Title 25 of the Rules of the City of New York.

Geographic Exclusion Area Over Time. I identify whether lots are in the original and expanded GEAs using lists from the Department of Housing Preservation and Development (HPD). In particular, the HPD lists identify lots in the expanded GEA but not the original GEA. As the original GEA was only in Manhattan, I can infer a lot was in the original GEA if it is in Manhattan but does not appear on the HPD list. See Appendix Figure A5 for a map of the GEA.

I further identify lots in the two small GEA regions. For the Greenpoint–Williamsburg waterfront expansion, I refer to the original zoning amendment text.²³ For the area south of 110th Street but not in the original GEA, I identified the relevant 2010 Census blocks by hand.

Mixed-Use Buildings. A typical mixed-use building has a commercial unit on the ground floor but residential units in the floors above ground. Since the 421-a exemption applies to the entire

²³The amendment's reference number is N 050110(A) ZRK and is available at <https://www1.nyc.gov/assets/planning/download/pdf/about/cpc/050110a.pdf> (pp. 109–110).

assessed value of a building, up to a cap in some years, but the inclusionary requirements apply only to residential units, there are additional provisions in 421-a that apply to mixed-use buildings. These provisions limit the incentive for adding commercial units, which would otherwise reduce the inclusionary units required but not the value of the exemption. In particular, if more than 12 percent of floor area is for commercial use, then the value of the tax exemption is reduced by the commercial share of floor area in excess of 12 percent. I implement this adjustment with data on floor area by use from PLUTO. For more information, see Section 1(d) of Section 421-a of the New York State Real Property Tax Law.

Liberty Zone Program. In the Job Creation and Worker Assistance Act of 2002 (JCWAA2002), an economic-stimulus package passed in response to the terrorist attacks of September 11, 2001 and the subsequent recession, the U.S. Congress included a place-based policy intended to support business and residential investment in Lower Manhattan. Of the \$5 billion cumulative fiscal cost of JCWAA2002's aid to Manhattan, \$1.6 billion was allocated for tax-exempt private-activity bonds for multifamily residential investment within a designated "Liberty Zone." These "Liberty Bonds" were issued by the New York City Housing Development Corporation and New York State Housing Finance Agency (HFA) from 2002 to 2007.²⁴ For further background on the policy, see Comptroller of the City of New York (2002) and U.S. Office of Management and Budget (2003).

The relevance of Liberty Bonds for 421-a is that, if a multifamily residential development which received any Liberty Bond financing, its eligibility for the 421-a exemption changed. In particular, due to the receipt of "substantial government assistance," such developments qualified for a 20-year 421-a exemption if they reserved five percent, rather than 20 percent, of units for lower-income tenants. For further discussion, see New York City Independent Budget Office (2003). To account for this policy difference, I introduce an indicator that equals one if a development was

²⁴A building is in the "Liberty Zone," as specified in U.S. Internal Revenue Service Publication 3991, if it is "located on or south of Canal Street, East Broadway (east of its intersection with Canal Street), or Grand Street (east of its intersection with East Broadway), in the Borough of Manhattan."

permitted between 2002 and 2007 in the Liberty Zone and equals zero otherwise. This indicator equals one for 28 buildings in my dataset. I control for it in my analyses, except in bootstrapped estimates.

Inclusionary Housing Program. 421-a is one of several policy instruments for inclusionary housing in New York City. Beyond 421-a, the Inclusionary Housing Program (IHP) offers developers density bonuses in select neighborhoods if they create onsite or offsite inclusionary units. Since 2016, after my sample ends, the Mandatory Inclusionary Housing (MIH) and Zoning for Quality and Affordability (ZQA) programs have further complicated incentives and requirements for inclusionary housing. I control for IHP eligibility in my analyses.

In 1987, New York City created the R10 Inclusionary Housing Program (IHP), which allows for “bonus” allowable density for residential buildings in R10 zoning districts that provide inclusionary housing. R10 zoning refers to areas with a maximum allowable floor area ratio of 10, is the highest-density zoning in New York City, and applies to buildings on major Manhattan avenues. Eligibility for the R10 IHP also applies to commercially-zoned areas if the zoning district’s “residential equivalent” is R10. In particular, for each square foot of floor area used in inclusionary housing, the development receives between 1.25 and 3.5 square feet of bonus floor area. Starting in 2005, the IHP was gradually expanded to “designated areas” that had been rezoned to allow for greater density. For further information, see Article II, Chapter 3, Section 23-154 of the New York City Zoning Resolution.

The IHP interacts with 421-a because, if onsite, inclusionary units for IHP also may also be applied towards the 20-percent inclusionary share under 421-a. Therefore, buildings in R10 districts or in IHP designated areas that qualify for 421-a also receive a density bonus. I use shapefiles from the Department of City Planning (publicly available on its “Bytes of the Big Apple” data archive) that map IHP designated areas and provide their date of designation. As New York City releases shapefile data in its State Plane Coordinate system, I convert the shapefiles to the geodetic coordinates (that is, latitude and longitude) using the U.S. National Geodetic Survey’s Coordinate Conversion and Transformation Tool (NCAT). I use the PLUTO data to identify buildings in R10

districts or commercial districts equivalent to R10 (C1-9, C2-8, C4-6, C4-7, C5-1, C5-2, C5-4, C5-5, C6-2, C6-4, C6-5, C6-6, C6-7, C6-8, C6-9). To account for this policy difference, I introduce an indicator that equals one if a development was IHP-eligible and equals zero otherwise. This indicator equals one for 350 buildings in my dataset.

Ignored 421-a Rules. Rents in market-rate units in 421-a buildings are also nominally rent-stabilized. During the period I study, however, these stabilization provisions were not enforced. 421-a also applies for up to three years of construction. However, during construction, a property's assessed value is generally only land, which is non-exempt under 421-a. I ignore this component of 421-a in my empirical analyses. 421-a contains provisions that reduce the value of the exemption when a development replaces an existing residential building. However, these provisions are rarely applied, and so I ignore them.

C.4 Other Housing Programs

This paper uses data on four other low-income housing programs beyond 421-a: the federal Low-Income Housing Tax Credit (LIHTC), tenant-based Section 8 (that is, the voucher program), project-based Section 8, and public housing operated by the New York City Housing Authority (NYCHA). This appendix subsection details provides details on the data sources, as well as on my data-cleaning procedures.

LIHTC. I use publicly-available microdata from the National LIHTC Database of the U.S. Department of Housing and Urban Development (HUD). The May 2019 release contains detailed building characteristics and financial information on all LIHTC buildings from 1997 to 2017.²⁵

The variables I use are the project's unique HUD identifier (HUD_ID), building name (PROJECT), and street address (PROJ_ADD, PROJ_CTY, PROJ_ST, PROJ_ZIP), the annual dollar amount of the tax credit allocated for each building (ALLOCAMT), and the dollar amount of other federal-government funds (HOME_AMT, TCAP_AMT, CDBG_AMT, HPVI_AMT, TCEP_AMT).

I restrict the sample to New York City using county FIPS codes. As coverage of the 2010

²⁵ Available at <https://www.huduser.gov/portal/datasets/lihtc.html>.

county FIPS codes variable is incomplete, I replace missing values for 2010 with 2000. No projects are missing county FIPS codes in both 2010 and 2000. Using the recorded project address and the Census public-use geocoding utility,²⁶ I also geocoded all LIHTC projects down to their 2010 Census block, which I use to measure area characteristics. I restrict the sample of LIHTC buildings to those funded between 2003 and 2015.

Tenant-Based Section 8. Building-level microdata on the tenant-based Section 8 program are not publicly available. I therefore use the HUD Picture of Subsidized Households (PoSH) database.²⁷ I use the 2015 release. For tenant-based Section 8, the lowest level at which public-use tabulations is by tract. The PoSH data are compiled from Form HUD-50058 filings, in which local housing authorities record the demographic characteristics and rent payments of each subsidized household.

The PoSH data include the number of Section 8 voucher units by tract (`assisted_unit_count`). Contributions by the government are fully determined by the information on Form HUD-50058. The PoSH data include, from this form, the government's calculation of fiscal costs per unit per month (`averagehudexpenditurepermonth`), which I annualize. I also match area characteristics by 2010 Census tract.

Project-Based Section 8. I use publicly-available data from the Office of Multifamily Housing Programs, which sits within the Federal Housing Administration and HUD, on the portfolio of active project-based multifamily housing contracts as of March 27, 2020.²⁸

I restrict the data to New York City using county FIPS codes (`county_code`) and to buildings on project-based Section 8 using the indicator variable `is_sec8_ind`. The variables I use are the building's identifier as assigned in HUD's internal Housing Enterprise Real Estate Management System (HEREMS) database (`project_id`), the building's address (`address_line1_text`), the total number of units (`total_unit_count`), and the total number of assisted units (`total_`

²⁶ Available at <https://geocoding.geo.census.gov/geocoder>.

²⁷ Available at <https://www.huduser.gov/portal/datasets/assthsg.html>.

²⁸ Available at https://www.hud.gov/program_offices/housing/mfh/presrv/mfhpreservation.

assisted_unit_count). I use the Census public-use geocoding utility to map addresses to 2010 Census blocks, which I use to measure area characteristics.

I do not use project-based Section 8 to benchmark 421-a for two reasons. First, the program contains a variety of sub-programs that serve different purposes and have different funding sources, making such comparisons less meaningful and making it difficult to assemble the necessary cost data. Second, in New York City, much of project-based Section 8 is either part of the Rental Assistance Demonstration (and thus converted NYCHA, with no useful estimates of per-unit cost) or supportive housing (that is, a program that integrates housing with social services, for which per-unit costs are likely non-representative of 421-a units).

Public Housing. I use the publicly-available 2020 NYCHA Development Data Book, which contains detailed information on resident demographic and building characteristics.²⁹ As New York City has not seen significant new construction of public housing in decades, no reliable data on marginal costs exist. I therefore do not benchmark 421-a units against public-housing units.³⁰

The variables I use are the building's Tenant Data System number (tds) and total number of units (totalnumberofapartments). Using the TDS identifier, I match buildings to their state plane coordinates as released by the New York City government.³¹ I convert these to the geodetic

²⁹Available at <https://data.cityofnewyork.us/Housing-Development/NYCHA-Development-Data-Book/evjd-dqz>.

³⁰As HUD explains in its documentation for the PoSH data: "These figures on current spending do not reflect what it would cost to expand any particular program. For example Public Housing costs appear low, since HUD paid off the construction costs several years ago. New Public Housing would cost much more than the average current spending, since construction costs would have to be paid." See U.S. Department of Housing and Urban Development, "A Picture of Subsidized Households: General Description of the Data and Bibliography," 1998, <https://www.huduser.gov/portal/datasets/assthsg/statedata98/descript.html>.

³¹Available at <https://data.cityofnewyork.us/Housing-Development/Map-of-NYCHA-Developments/i9rv-hdr5>.

coordinates (that is, latitude and longitude) using the U.S. National Geodetic Survey's Coordinate Conversion and Transformation Tool (NCAT).³²

C.5 Case Study for Simulated Instrument

This subsection illustrates some of the policy variation in the simulated instrument as defined in Section 7.2. Let $\Delta\tau(b,x)$ be the potential 421-a tax incentive a developer would face if they built a building with characteristics b on a lot with characteristics x . Indexing lots by i , the realized tax incentive is $\Delta\tau_i = \Delta\tau(b_i, x_i)$. Had a developer built another building in their feasible set ($b \in \mathcal{F}_i$), they would have faced a counterfactual tax incentive $\Delta\tau(b, x_i)$. The instrument averages over $\Delta\tau(b, x_i)$ for all $b \in \mathcal{F}_i$ to purge the endogenous choice of b from $\Delta\tau_i$.

To illustrate the policy variation in isolation, as distinct from variation due to the endogenous building choice, consider a given building b and the variation in $\Delta\tau(b,x)$ created by changing lot characteristics x . In other words, how large of a tax incentive would the same building have received if it were built on a lot with different tax-relevant characteristics? This policy variation is substantial: The within-building standard deviation of $\Delta\tau(b,x)$ is 6.3 percent of property value, as compared to an unconditional standard deviation of 7.8 percent of property value. If property taxes were uniform in lot characteristics, holding the building fixed, the within-building standard deviation of $\Delta\tau(b,x)$ would be zero. Thus policy variation in $\Delta\tau(\cdot, x)$ induces 65 percent of the variation in $\Delta\tau(b, x)$ while endogenous variation in building characteristics explains 35 percent.

As a case study, I focus on one building, 168 East 112nd Street in the East Harlem neighborhood of Manhattan, which received its initial construction permit in 2011. The developer provided onsite inclusionary units and received a tax incentive that I estimate is worth, in present value, 81.6 percent of market value. In constructing my instrument, I find the building b at 168 East 112nd Street is in the feasible sets \mathcal{F}_i of 21 other lots. In what follows, I describe three, all in East Harlem, of the 21 lots. Appendix Figure A28 maps the four locations.

- Had this building instead been built at 324 East 112nd Street in 2004 (in lieu of the actual

³²Available at <https://www.ngs.noaa.gov/NCAT/>.

building permitted there), it would have faced a 421-a tax incentive worth 0 percent of building value. This is because, in 2004, 324 East 112nd Street was in an NPP district and thus received 421-a as of right without providing inclusionary housing.

- Had this building instead been built at 15 East 105th Street in 2004, it would have faced a 421-a tax incentive of 47.5 percent of building value. The incentive is smaller because this lot is outside the GEA in 2004 but would become part of the GEA in its expansion.
- Had this building instead been built at 238 East 106th Street in 2014, it would have faced a 421-a tax incentive of 77.9 percent of building value. Like 168 East 112nd Street in 2011, this lot was covered by the GEA expansion area by 2014 and thus would have received similarly generous incentives.

To see the endogenous variation in building choice that the simulated instrument purges, one could perform the opposite exercise: Holding fixed the lot characteristics x_i , consider different buildings $b \in \mathcal{F}_i$ and the variation in tax incentives $\Delta\tau(b, x_i)$ among feasible buildings.

C.6 Modeling the Impact of Assessed Value Growth Caps

An input to the calculation of the 421-a tax incentive is the degree of underassessment at the building level. As Section 3 explains, this underassessment has two components: (1) estimation errors of market value and (2) growth caps on taxable assessed value that cause it to diverge from actual assessed value. This subsection focuses on the latter source of underassessment.

I use a simple age–period–cohort model of the taxable share of actual assessed value to forecast and “backcast” this share for all years and buildings in the dataset. I do so using the 2010–2018 Real Property Assessment Databases. The model therefore serves two purposes. First, it solves the missing-data problem for buildings built before 2010. Second, it provides structure to my forecasts, as I explain below.

Let s_{it} be the ratio of taxable to actual assessed value for building i in year t and Age_{it} the years since the completion of construction. Suppose also that c indicates the tax class of the building: $c = 1$ if i has exactly three units, $c = 2$ if it has between four and ten units, and $c = 3$ if it has 11

or more units. I model the log odds ratio as a linear function of a building-specific fixed effect and linear slope in age plus a flexible class-specific age effect:

$$\log\left(\frac{s_{ict}}{1-s_{ict}}\right) = \alpha_i + \beta_i \text{Age}_{it} + \sum_{a,c} [\gamma_{ac} \cdot 1(\text{Age}_{it} = a, \text{Class}_i = c)] + u_{ict}.$$

I estimate this model using only the buildings in my dataset. This specification omits period effects. To construct my forecasts, I assume that the flexible class-specific age effects are linear beyond the oldest age I observe. In particular, my linear extrapolation is

$$\gamma_{ac} = \gamma_{15,c} + \left(\frac{\gamma_{15,c} - \gamma_{10,c}}{5}\right)(s - 15)$$

which fits a slope from the age effects at age 10 and age 15. In Figure A16, I plot the average taxable share of assessed value by class and age. Assessment growth caps depress taxable shares substantially, and the degree of underassessment rises in building age.

C.7 Modeling the Underassessment of Owner-Occupied Buildings

In Section 3, I explain that the the New York City Department of Finance systematically under-assesses owner-occupied units (condominiums and cooperatives) relative to their true market values. To estimate the degree of underassessment at the building level, I use data from the Department of Finance on all property transactions in New York City from 2003 to 2019.³³

I restrict the dataset to residential units on non-vacant lots which trade at a price above \$25,000. Some units appear to be sold at very low prices or prices of exactly zero, which likely reflects property transfers between legal entities rather than arm’s-length sales. I also hand-correct data-entry errors related to the number of units sold.³⁴

The first step of my process is to estimate market values in 2019 so that they align with my

³³Available at <https://www1.nyc.gov/site/finance/taxes/property-annualized-sales-update.page>.

³⁴In particular, I noticed that sales of entire condominium buildings were sometimes incorrectly entered as “one” unit sold, making these transactions look like sales of single condominium units within a building. I found 28 data-entry errors of this kind.

assessed values. I use a repeat-sales method to estimate the citywide time path of appreciation:

$$\log p_{it} = \alpha_i + \alpha_t + e_{it},$$

where p_{it} is the price, α_t is a year fixed effect and α_i is a fixed effect for each unit. The sample is limited to likely-arm's-length sales of single residential units. I use the estimated time fixed effects to correct prices to the 2019 base year:

$$\log \tilde{p}_i = \log p_{it} + \widehat{\alpha}_{2019} - \hat{\alpha}_t.$$

After using repeat-sales to correct prices across years, I retain only initial sales, so that prices capitalize the full value of the 421-a exemption. I then compute underassessment ratios:

$$U_j = s_j^R (V_j/N_j) \Big/ \frac{\gamma_j}{n_j} \sum_{i \in j} \tilde{p}_i,$$

where s_j^R is the residential share of floor space in building j as recorded in the assessment data, V_j is the total assessed value in 2019, N_j is the number of residential units in the building, and n_j is the number of initial sales observed in j . The coefficient γ_j is for fractional assessment: $\gamma_j = 0.06$ for buildings with three or fewer units, and $\gamma_j = 0.45$ for buildings with four or more units. We censor U_j at 1.5 to address potential remaining data entry errors.

Appendix Figure A27 plots a histogram of underassessment ratios U_j . The average building's assessed value is 38 percent of its estimated market value, and the standard deviation of the assessment ratio is 29 percent. The degree of underassessment is consistent with estimates by the New York City Advisory Commission on Property Tax Reform, and the significant cross-sectional inequality in assessment conditional on market value is consistent with the nationwide analysis of Avenancio-León and Howard (2019).

C.8 LIHTC Cost Estimates

- U.S. Government Accountability Office (2018) reports per-unit total development costs (TDC) of \$282,000, surveying projects in New York City from 2011 to 2015. Under a 70-percent PDV (9-percent) LIHTC, this amounts to a cost of \$197,400.

- Abt Associates (2018) estimates a regression model of LIHTC TDC on data from 2011 to 2016. I use their reported coefficients to construct a typical LIHTC development in New York City—in particular, a Mid Atlantic principal city with high construction wages. I again assume that 70 percent of TDC is financed by the LIHTC. I obtain an estimated cost of \$204,680.
- Adjusting for inflation to 2015, Deng (2005) estimates a cost of \$232,340.
- I use HUD’s per-unit TDC limit in 2015 for development in “inner” (not metro-area) New York City. In particular, I restrict to elevator buildings and take a weighted average of per-unit TDCs by number of bedrooms, weighting by the actual share in New York City per the HUD LIHTC database. I again assume that 70 percent of TDC is financed by the LIHTC. I obtain an estimated cost of \$234,580.

C.9 MVPF Details

I build my MVPF calculator program from the code base of Hendren and Sprung-Keyser (2020), which provides a set of resources to simulate long-run effects of programs similar to 421-a on after-tax incomes and tax payments. In particular, I closely followed their MVPF calculator programs for the Moving to Opportunity experiment (Chetty et al., 2016) and the Chicago Housing Authority voucher experiment (Jacob et al., 2015). As Hendren and Sprung-Keyser (2020) note and as I remind in my Section 9, there are many “judgment calls” in going from a program treatment effect to an MVPF. Here I describe several places where I customized the assumptions from Hendren and Sprung-Keyser (2020) to fit my context.

1. I replace values for several parameters to match New York City character for which Hendren and Sprung-Keyser (2020) use a U.S. national average. First, I replaced the intergenerational income rank–rank function. To do so, I used the estimates of individual income mobility at the 1st, 25th, 50th, 75th, and 100th percentiles from Chetty et al. (2018), pooled among native-born adults across sex and race (variables beginning `kir_native_pooled_pooled_`),

for the five New York City counties. I weight by number of observations to construct city-wide averages. I use these points to adjust the national rank–rank function, which is included in the Hendren and Sprung-Keyser (2020) data files for all percentiles. In particular, I estimate the following regression specification:

$$r_{q,\text{NYC}} = \beta_0 + \beta_1 r_{q,\text{USA}} + \beta_2 r_{q,\text{USA}}^2 + u_q,$$

where $r_{q,j}$ is the expected individual income rank of a child born to a household at the q th percentile of the national income distribution in region j (i.e., either New York City or the U.S. overall). A quadratic specification in national expected rank provided an extremely close fit for the five percentiles that I can directly match between New York City and the U.S. national average. I then interpolated to all other percentiles using the predicted values $\widehat{r_{q,\text{NYC}}}$ from the regression. For a visualization, see Appendix Figure A22. In general, low-income children in New York City have higher expected individual income ranks than typical low-income children.

2. I replace the age–earnings profile using the 2015 5-year American Community Survey public-use microdata to match those of New York City workers. Appendix Figure A23 presents the age profile. This adjustment helps to capture the higher expected lifecycle earnings of New York City children.
3. I replace the state marginal tax rate (MTR) on labor earnings. Hendren and Sprung-Keyser (2020) draw upon a CBO report which estimates it to be 2.6 percent. Consulting the data files for NBER TAXSIM, and taking the simple average of values from 2003 to 2015, I find a New York State MTR of 7.253 percent, consistent with high state income taxes in New York. Furthermore, New York City assesses a local income tax. As it is nearly linear in income, I use a rate of 3.8 percent, which is approximately the statutory rate applying to earnings above \$21,600 as of 2020. Consequently, I assume a state-and-local MTR on labor earnings of 11.053 percent. Unless otherwise stated, I present “consolidated-government” MVPFs, which combine fiscal externalities across the federal, state, and local governments.

4. I draw on available data to approximate the characteristics of inclusionary-tenant households. Following Hendren and Sprung-Keyser (2020), and as I explain below, I use this information to forecast earning impacts on children and parents from 421-a.

- I assume that the average household earns 45 percent of Area Median Income (AMI) in the year of their application (\$37,355 in 2016). To arrive at this estimate, I considered, as a maximum, 421-a’s statutory income cap at 60 percent of AMI. As a minimum, I considered an affordability requirement which says that rent cannot exceed 35 percent of household income, along with an inspection of 421-a rents listed on NYC Housing Connect. My estimate is consistent with contemporary reporting on inclusionary tenants.³⁵
- I approximate inclusionary-tenant household composition using data from the 2003–2015 U.S. Current Population Surveys. Among households in New York City that received housing subsidies (per the Supplemental Poverty Measure module question SPMCAPHOUS), the average number of children (age 18 or younger) per family was 1.44. Of children, 60.8 percent were age 13 or younger. Among younger children, the average age was 6.04 years; among older children, the average age was 15.44 years. Among households with young children, the average age of the head of household was 40.5 years; among those with young children, the average age of the head of household was 44.6 years.

³⁵For NYC Housing Connect, see <https://a806-housingconnect.nyc.gov/nyclottery/lottery.html#home>. For evidence on income ranges, see Marc Santora, “Across the Hall, Diversity of Incomes,” *The New York Times*, 2 September 2011. Data on all New York City social-housing lotteries from 2014 to 2019 suggest that the average minimum and maximum household income for a unit was \$44,092 to \$55,210, implying a midpoint of \$49,650. See Rachel Holliday Smith, Ann Choi and Will Welch, “Affordable Housing Lottery Odds Worst for Those Who Can Afford the Least,” *The City*, 28 June 2020.

5. I predict earnings impacts on parents and children, leveraging my New York City data described above. For children, I closely follow the approach of Bergman et al. (2020). This approach forecasts the average impact of a move on a child’s future individual income rank of children as an adult using data from the “Opportunity Atlas” of Chetty et al. (2018). Given origin and destination Census tracts and the child’s parent household income rank, I compute the difference in future individual income ranks of children from those tracts from that parent household income rank. I then multiply that difference in ranks by 0.62, consistent with evidence from Chetty and Hendren (2018) on the share of their estimated neighborhood effects which are causal effects versus selection. I then convert this difference in ranks into a percentage change in earnings. With the assumption that this percentage change is fixed over the child’s adult lifetime, I use my age–earnings profile, an annual discount rate of 5 percent, and the estimated ages of children to compute a present discounted value impact on earnings. These earnings changes further imply changes in after-tax earnings and tax payments following my assumptions on marginal tax rates. For parents, I assume earnings effects are the same in percentage terms as in the Moving to Opportunity experiment (Chetty et al., 2016), which I scale into dollar terms using my estimate of inclusionary-tenant household income.
6. I use a calibrated model of neighborhood choice to approximate the distribution of source neighborhoods from which households are diverted by inclusionary units. Unlike in Bergman et al. (2020), I do not have an experiment by which I can estimate counterfactual residential choice, so this step allows me to compute the change in child earnings as outlined above. In particular, I assume that, if an inclusionary unit is added in Neighborhood Tabulation Area n , inclusionary-tenant household i would have otherwise lived in Census tract j with a probability that is given by

$$\Pr(R_i = j | n) = s_j d_{j,n}^\beta / \sum_{j'} d_{j',n}^\beta,$$

where s_j is the share of New York City households earning between \$30,000 and \$60,000

in 2015 real dollars who live in tract j and $d_{j,n}$ is the distance in miles between the centroids of tract j and Neighborhood Tabulation Area n . I obtain the tract distribution of these households from the 2011–2015 American Community Surveys, and I calibrate β so that the average move distance $E[d_{j,n}]$ equals 10 kilometers.

7. To estimate inclusionary tenants' WTP for their units, I use the distribution of counterfactual housing expenditure from the 2011–2015 ACS microdata along with assumptions that I argue provide plausible lower and upper bounds on WTP. For the upper bound, I assume households value the unit at the developer breakeven. The lower bound uses the expenditure data. For households earning \$30,000 to \$60,000 in real 2015 dollars in New York City, I compute the housing expenditure distribution and define the WTP in neighborhood n as

$$\text{WTP}_n = E[\min\{y_i, \text{Breakeven}_n\}]$$

where y_i is housing expenditure of a household i .

8. The two MVPF modules from Hendren and Sprung-Keyser (2020) upon which I build my own module estimate the WTP and net cost on a per-household basis. For this reason, I must estimate the average number of children per family in 421-a units and the average age of those children. These data are not publicly available for 421-a, and so as a substitute, I estimate these values using U.S. Current Population Survey data from 2003 to 2014 on housing-aid recipients in New York City. These households are primarily households that receive Section 8 vouchers or are in public housing.

D Additional Results

D.1 Event Study Analysis of the GEA and the NPP Reforms

Section 6 shows that reforms to the Geographic Exclusion Area (GEA) and Neighborhood Preservation Program (NPP) changed the 421-a incentives for some buildings relative to others, resulting in relative changes in take-up of 421-a. Here I repeat that analysis in event-study regressions.

The first set of regressions, analogous to the left panels in Figure 1, employs the following event-study specification:

$$\Delta\tau_i = \alpha_b + \alpha_t + \beta_t \text{Region}_i + X_i\gamma + u_i,$$

where i indexes buildings, $\Delta\tau_i$ is i 's incentive to provide onsite inclusionary housing, α_b is a borough fixed effect, α_t is a year fixed effect, Region_i is an indicator for whether building i is in the GEA or NPP (depending upon the specification), and X_i is a matrix of lot- and block-level controls. The time-varying coefficients β_t therefore capture the change in the 421-a tax incentive for buildings in the treated region relative to observably-similar buildings in the untreated region.

The second set of regressions, analogous to the left panels in Figure 1, employs a similar specification:

$$D_i = \alpha_b + \alpha_t + \delta_t \text{Region}_i + X_i\rho + u_i,$$

where D_i is an indicator for whether building i provides onsite inclusionary housing. The time-varying coefficients δ_t therefore capture the change in the 421-a take-up rate for buildings in the treated region relative to observably-similar buildings in the untreated region. In the NPP analysis, I restrict the control group to buildings in a Census tract that overlaps with an NPP area, to be consistent with Figure 1.

Appendix Figure A9 presents the event-study estimates $\hat{\beta}_t$ and $\hat{\delta}_t$. Consistent with Figure 1, the average 421-a tax incentive in the GEA and NPP regions rises relative to similar buildings in untreated regions after the reforms occur, with little evidence of pre-trends. The 421-a participation rate of buildings in the GEA and NPP regions also rises relative to that of similar buildings in the untreated regions.

D.2 Supply Curve

I define the supply of inclusionary housing in an arbitrary neighborhood n as the share of buildings in n that, under a common 421-a incentive $\Delta\tau$, would choose to take up 421-a:

$$S_n(\Delta\tau) = \int \Pr(D_i = 1 | x_i, \Delta\tau) dF_n(x_i).$$

I obtain $\widehat{\Pr}(D_i = 1 | x_i, \Delta\tau)$ from the estimated logit parameters. Appendix Figure A24 displays the estimated $\widehat{S}(\Delta\tau)$. At $\Delta\tau \approx 0$, almost no buildings take up 421-a, and participation rises most steeply around $\Delta\tau \approx 0.6$, the modal breakeven in Appendix Figure A13. I conduct simultaneous inference over the supply curve following the bootstrap procedure of Chernozhukov et al. (2013). The proper statistical interpretation of these confidence bands is that, sampling from the population, 95 percent of the resulting bands would enclose the entire true supply curve: that is, they apply jointly, not pointwise, over the function. I weight $S(\Delta\tau)$ by the number of residential units in a building and round down fractional inclusionary units.

To construct the confidence bands, I first take B bootstrap samples, clustering on Neighborhood Tabulation Area.³⁶ For each bootstrap sample $b = 1, \dots, B$, I re-estimate the logit specification in Equation 5. I then use my estimates $(\widehat{\sigma}^{(b)}, \widehat{\beta}^{(b)})$ to predict 421-a participation probabilities, and thereby $\widehat{S}^{(b)}(\Delta\tau)$, on a grid of $\Delta\tau$. In implementation, I use grid steps of size 0.05 from 0 to 1.5. At each $\Delta\tau$ value on the grid, I compute the bootstrap mean $\widetilde{S}(\Delta\tau) = B^{-1} \sum_b \widehat{S}^{(b)}(\Delta\tau)$ and the variance $\widehat{V}_S(\Delta\tau) = B^{-1} \sum_b \left[\widehat{S}^{(b)}(\Delta\tau) - \widetilde{S}(\Delta\tau) \right]^2$. To obtain confidence bands that apply simultaneously over the whole curve, rather than pointwise, I define the relevant critical values over the maximum z -score on each bootstrapped curve:

$$c(1-a) = (1-a) \text{ quantile of } \left\{ \max_{\Delta\tau} \left[\widehat{V}_S(\Delta\tau) \right]^{-1/2} \cdot \left| \widehat{S}^{(b)}(\Delta\tau) - \widetilde{S}(\Delta\tau) \right|, \quad b = 1, \dots, B \right\}.$$

³⁶My presentation of the bootstrap procedure of Chernozhukov et al. (2013) draws heavily on lecture notes that have graciously been made publicly available by Victor Chernozhukov and Iván Fernández-Val on MIT OpenCourseWare: https://ocw.mit.edu/courses/economics/14-382-econometrics-spring-2017/lecture-notes/MIT14_382S17_lec5.pdf

The confidence band is then a Cartesian product of sets at each grid value of $\Delta\tau$ ($\Delta\tau_k$ for $k = 1, \dots, K$):

$$\text{CB}_{1-a} = \times_{k=1}^K \left[\widehat{S}(\Delta\tau_k) \pm c(1-a) \left[\widehat{V}_S(\Delta\tau_k) \right]^{1/2} \right].$$

A caution about this analysis is that it quite heavily relies upon the single-index assumption on $x_i\beta$ to extrapolate from marginal units to nonmarginal units. That is, at extreme values of $\Delta\tau$, the counterfactual marginal building is, under the observed $\Delta\tau_i$, a building that is an “almost-always-taker” or “almost-never-taker” of 421-a. There is little information in the data about the breakeven thresholds of such buildings, and therefore there is more uncertainty about their behavior than the confidence bands communicate. Another limitation is that nonmarginal changes in 421-a may have general-equilibrium effects that my results would not capture. An example, one to which I return below, is a large change in 421-a may affect the equilibrium assignment of lots to developers.

D.3 Changing the Set-Aside Share

How would the inclusionary share of all units change in response to a change in the set-aside share λ ? I define I_n , the inclusionary share of all units in an arbitrary neighborhood n , as

$$I_n(\lambda; \{\Delta\tau_i\}) = \lambda \int \Pr(D_i = 1 | x_i, \lambda, \Delta\tau_i) dF_n(x_i),$$

again weighting by the number of residential units and rounding down fractional inclusionary units. At very low λ , the rounding is consequential: A policy of $\lambda = \frac{1}{200}$, for example, will produce almost no inclusionary units because very few buildings have more than 200 residential units. As above, I follow Chernozhukov et al. (2013) to conduct simultaneous inference.

There is no variation in λ under 421-a. I therefore adopt a further parametric assumption on forgone rent that allows me to extrapolate from variation in $\Delta\tau_i$. In particular, I assume that, for building i , the disamenity on market-rate units is $\delta_i(\lambda) = \min\{\frac{\lambda}{1-\lambda} \delta_0(x_i), \mu(x_i)\}$. Up to the point at which the market rent hits the regulated rent ($\delta_i(\lambda) = \mu_i$), this functional form implies that the disamenity is increasing in λ and nests the special case of no disamenity ($\delta_0 = 0$). By this

assumption, I can rewrite Equation 3 as

$$\frac{\Delta\tau_i}{\lambda} \geq (\mu + \delta_0)(x_i),$$

where the composite function $(\mu + \delta_0)(\cdot) = \mu(\cdot) + \delta_0(\cdot)$. The benefit of this assumption is it provides an equivalence between $\Delta\tau_i$ and λ : For any ε , $\Delta\tau$, and λ , developers are indifferent between $(\Delta\tau, \lambda)$ and $((1 + \varepsilon)\Delta\tau, (1 + \varepsilon)\lambda)$. An important caveat is that the plausibility of this assumption is hard to assess, particularly for extreme values of λ , where my results are most dependent upon this assumption for extrapolation. For example, at high λ , there may be a “tipping” equilibrium that my data would not reveal. I therefore stop the extrapolation at $\lambda = 0.5$.

Panel B of Figure A24 reports estimates of $I_n(\lambda; \{\Delta\tau_i\})$ along with the 95-percent simultaneous confidence band. At a very low set-aside share for 421-a buildings, a marginal increase in the requirement raises the inclusionary share of all units citywide, as more units are gained from 421-a buildings than are lost due to exit from 421-a. However, at higher set-aside requirements, the effect of exit fully undoes the mechanical increase in inclusionary units from inframarginal buildings. I also find that citywide inclusionary share of units does not go to zero as λ becomes large. A small share of buildings have near-zero estimated breakevens, and so they may still take-up 421-a even at high set-aside requirements. However, as discussed above, my environment is not well-suited to support strong claims about the production of inclusionary units at high λ .

It is important to note that, in inclusionary housing, the Laffer curve lacks the welfare implications it has in the context of optimal taxation (Werning, 2007). As there are fiscal costs to increases in 421-a participation, and as the characteristics of marginal units change as λ varies, being to the right of the Laffer peak does not imply the existence of a Pareto-improving reform. Nevertheless, λ is a key parameter in inclusionary housing policies, and these results can inform relevant policy discussions.

D.4 Developer Sorting

In a competitive land market, the equilibrium assignment of developers to lots is endogenous to tax policy. This point affects the interpretation of my results, as estimated breakevens may change

if lots were assigned to different developers. It also has testable implications. Suppose developers vary in a fixed cost ψ_j of taking up 421-a on any given building. Anecdotally, some developers invest in administrative capacity so that they can cooperate effectively with government agencies, implying a low ψ_j for such “specialized” developers. This motivates a specification that extends Equation 2:

$$\Delta\pi_{ij} = \Delta\tau_i + x_i\beta - \psi_j + \Delta e_i. \quad (19)$$

Such a model makes two predictions. First, conditional on x_i , specialized developers should sort into lots with high $\Delta\tau_i$, as they have a comparative advantage in building on such lots. Second, conditional on $\Delta\tau_i$ and x_i , a lot is more likely to participate in 421-a if the developer is specialized. Appendix Table A5 finds evidence for the latter prediction, but not the former. As a proxy for ψ_j , I use the leave-out share of the developer’s buildings that have onsite inclusionary housing, weighted by their residential unit counts. Comparing a developer that has never built a building with inclusionary housing to one that exclusively does such development, the latter behaves as if it has a 421-a cost advantage worth about 25 p.p. of building value on the former.

D.5 Offsite versus Onsite Inclusionary Housing

Here I consider models that include an offsite–onsite choice for inclusionary housing. Letting $j = 1$ denote offsite and $j = 2$ denote onsite, I estimate a multinomial logit model with alternative-specific coefficients. Offsite and onsite 421-a take-up probabilities take the following form:

$$\Pr(D_i = j | x_i, \tau_{ij}) = \frac{\exp[(\Delta\tau_{ij} + x_i\beta_j)/\sigma]}{1 + \exp[(\Delta\tau_{i,1} + x_i\beta_1)/\sigma] + \exp[(\Delta\tau_{i,2} + x_i\beta_2)/\sigma]}. \quad (20)$$

In Appendix Figure A11, I calculate developer breakevens for offsite and onsite inclusionary housing. I find lower breakevens for onsite units than for offsite units, especially in in higher-rent neighborhoods. This pattern is consistent with breakevens determined by the correct notion of opportunity cost, which follows from the neighborhoods in which the inclusionary units are ultimately located. In Appendix Table A6, I report citywide supply responses and marginal fiscal costs for offsite and onsite units. I find a very low supply response for offsite units, leading the

marginal fiscal cost of onsite units to be below that of offsite units, despite higher breakevens.

D.6 Extensions and Robustness Checks for MVPF Analysis

Inference. To what extent is neighborhood variation in MVPFs mere sampling variance? Following Mogstad et al. (2020), I estimate the simultaneous confidence set over the MVPF percentile ranks of neighborhoods. I incorporate here cost-side uncertainty from my analysis as well as benefit-side uncertainty from Chetty et al. (2018). Appendix Figure A26 depicts the upper and lower bounds of this confidence set. Neighborhood MVPF ranks are relatively precise by comparison to the Mogstad et al. (2020) reanalysis of Bergman et al. (2020). Why do my results differ? Neighborhood variation in costs is large relative to variation in benefits and is more precisely estimated. In Appendix Table A12, I split neighborhoods by quartiles of the MVPF distribution and compute the average WTP and net cost by quartile. Whereas WTP is roughly constant across quartiles, the net cost per unit is about ten times higher in the bottom-quartile neighborhoods than in the top-quartile neighborhoods. These results highlight the importance of cost variation, not only variation in mobility, in designing policies to increase opportunity.

Crowding-Out. Whether 421-a crowds-out unsubsidized low-income housing across the street or across town is important to the welfare analysis. Crowding-out at a very fine spatial scale would make the policy impotent. However, research on subsidized housing generally does not find such proximate crowding-out. Baum-Snow and Marion (2009) find that LIHTC units are 20-percent crowded out within a one-kilometer distance, and Eriksen and Rosenthal (2010) find full crowding-out, but at a distance of about 16 kilometers.

The key parameter regulating the scale of crowding-out in my MVPF analysis is the distance cost of moves in the calibrated model of neighborhood choice. In my baseline analysis, I calibrate the model so that the average move distance is 10 kilometers. Bergman et al. (2020) find an average move distance of about 17 kilometers, which would imply less proximate crowding-out. I also test the robustness of my analysis to a smaller move distance. In particular, I calibrate the model to match the crowding-out estimates in Baum-Snow and Marion (2009) and Eriksen and

Rosenthal (2010). These results are shown in Appendix Figure A21. The MVPFs are somewhat lower, as changes in neighborhood opportunity are smaller between nearby neighborhoods, but the qualitative conclusions are essentially unchanged from above.

Targeting. Across-the-board increases in the generosity of the 421-a exemption would be badly targeted at high-MVPF neighborhoods. Comparing Figures 4 and 5, supply responses are lower in high-MVPF neighborhoods than in low-MVPF neighborhoods. Typical buildings in high-MVPF neighborhoods are single-family and small multifamily developments, which pay low effective tax rates and thus have low potential tax savings under 421-a. Indeed, these neighborhoods are usually zoned only for low-density development, making it difficult for developers to build the housing that is likeliest to accept 421-a. Targeted increases in incentives for these neighborhoods and types of buildings could increase the citywide MVPF of 421-a, but rezonings are also likely required to yield significant production of inclusionary housing.

D.7 Correlates of Neighborhood MVPF

This subsection examines the association of neighborhood characteristics with MVPFs. In particular, I correlate percentile ranks of neighborhood MVPF with percentile ranks of the demographic characteristics listed in Section 5. Appendix Figure A29 presents the results. In summary:

- The relationship between neighborhood MVPF and measures of socioeconomic status generally follows an inverted-U shape. That is, neighborhoods with very low average incomes, rents, and levels of schooling typically have low MVPFs, as do neighborhoods with very high levels of these variables.
- Neighborhood MVPFs are increasing in the commuting mode share of cars and median resident age and decreasing in population density. These patterns are consistent with zoning regulations in many high-MVPF neighborhoods that are ill-suited to inclusionary housing and to low-cost housing more generally, as noted in Section 9.
- Racial and ethnic composition is related to neighborhood MVPF. MVPFs are decreasing in

the black and Hispanic share of residents, increasing in the white and Asian shares. This pattern is consistent with welfare gains from integrating neighborhoods.

These considerations suggest that, if between-neighborhood patterns in New York City generalize to other cities, then inclusionary housing policies may be most beneficial in middle-income neighborhoods that have been historically resistant to multifamily housing and as a policy lever for residential desegregation. However, the external validity of these results to other cities is unknown.

References for Appendices

Abt Associates, “Variation in Development Costs for LIHTC Projects,” 2018.

Avenancio-León, Carlos and Troup Howard, “The Assessment Gap: Racial Inequalities in Property Taxation,” Working Paper 2019.

Baum-Snow, Nathaniel and Justin Marion, “The Effects of Low Income Housing Tax Credit Developments on Neighborhoods,” *Journal of Public Economics*, 2009, 93 (5-6), 654–666.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F Katz, and Christopher Palmer, “Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice,” Working Paper 26164, National Bureau of Economic Research 2020.

Chernozhukov, Victor, Iván Fernández-Val, and Blaise Melly, “Inference on Counterfactual Distributions,” *Econometrica*, 2013, 81 (6), 2205–2268.

Chetty, Raj and Nathaniel Hendren, “The Impacts of Neighborhoods on Intergenerational Mobility I: Childhood Exposure Effects,” *The Quarterly Journal of Economics*, 2018, 133 (3), 1107–1162.

—, **John N Friedman, Nathaniel Hendren, Maggie R Jones, and Sonya R Porter**, “The Opportunity Atlas: Mapping the Childhood Roots of Social Mobility,” Working Paper 25147, National Bureau of Economic Research 2018.

– , **Nathaniel Hendren, and Lawrence F Katz**, “The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment,” *American Economic Review*, 2016, 106 (4), 855–902.

Comptroller of the City of New York, “One Year Later: The Fiscal Impact of 9/11 on New York City,” 2002.

Deng, Lan, “The Cost-Effectiveness of the Low-Income Housing Tax Credit Relative to Vouchers: Evidence from Six Metropolitan areas,” *Housing Policy Debate*, 2005, 16 (3-4), 469–511.

Eriksen, Michael D and Stuart S Rosenthal, “Crowd Out Effects of Place-Based Subsidized Rental Housing: New Evidence from the LIHTC Program,” *Journal of Public Economics*, 2010, 94 (11-12), 953–966.

Hendren, Nathaniel and Ben Sprung-Keyser, “A Unified Welfare Analysis of Government Policies,” *The Quarterly Journal of Economics*, 2020, 135 (3), 1209–1318.

Jacob, Brian A, Max Kapustin, and Jens Ludwig, “The Impact of Housing Assistance on Child Outcomes: Evidence from a Randomized Housing Lottery,” *The Quarterly Journal of Economics*, 2015, 130 (1), 465–506.

Mogstad, Magne, Joseph P Romano, Azeem Shaikh, and Daniel Wilhelm, “Inference for Ranks with Applications to Mobility across Neighborhoods and Academic Achievement across Countries,” Working Paper 26883, National Bureau of Economic Research 2020.

New York City Independent Budget Office, “Worth the Cost? Evaluating the 421-a Property Tax Exemption,” Report 2003.

Thaden, Emily and Ruoniu Wang, *Inclusionary Housing in the United States: Prevalence, Impact, and Practices*, Lincoln Institute of Land Policy, 2017.

U.S. Government Accountability Office, “Low-Income Housing Tax Credit: Improved Data and Oversight Would Strengthen Cost Assessment and Fraud Risk Management,” Report GAO-18-637 2018.

U.S. Office of Management and Budget, “Analytical Perspectives: Budget of the U.S. Government, Fiscal Year 2004,” 2003.

Werning, Iván, “Pareto Efficient Income Taxation,” Mimeo, MIT 2007.