The Impact of Greenbelts on Housing Markets: Evidence from Toronto*

Alexander Hempel[†]
University of Toronto

January 26, 2024

Click Here for Latest Version

Abstract

Greenbelts are a widespread policy tool used to protect natural spaces from urban sprawl. With rising housing costs in many metropolitan areas, numerous questions have been raised about the impact of greenbelts on housing markets. Yet despite the intense policy debate, there is little empirical evidence to assess how greenbelts affect housing supply and prices across a metropolitan region. In this paper, I set out a new approach to estimate the impact of greenbelt policies on housing market outcomes and use it to evaluate the introduction of the world's largest contiguous greenbelt, which formed a protected zone around Toronto in the early 2000s. Using rich project-level data on housing developments, I first show that the Ontario Greenbelt affected housing development patterns, where restricted, developable census tracts saw less housing built relative to unrestricted tracts. Next, to quantify the effects across the metropolitan area, I build an estimable model of housing supply and demand with heterogeneous supply elasticities at the census tract level. I estimate the model using instrumental variable approaches including a novel heritage designations instrument. Using the estimated model, I simulate the scenario in which no Greenbelt was implemented, finding that the Greenbelt led to an average increase in housing costs of 2.25%; this corresponds to an increase in housing rent of around C$550 a year. In addition, I show that had the Greenbelt been paired with a small relaxation of zoning regulations within the city, these negative consequences from the Greenbelt would disappear, suggesting a viable alternative to developing greenbelts in the face of rising housing prices.

*I would like to thank Robert McMillan, Stephan Heblich and Jonathan Hall for their continuous support throughout the project. I would also like to thank Arthur Blouin, Michael Baker and Kory Kroft for their helpful feedback and Hans Koster, Gabriel Ahlfeldt, Alvin Murphy, Patrick Bayer, Ian Herzog, Torsten Jaccard and Kate Pennington for productive discussions.

[†]University of Toronto, Department of Economics, Email: alex.hempel@mail.utoronto.ca.

1 Introduction

As urbanization proceeds apace in numerous countries across the world, the phenomenon of urban sprawl is becoming an increasingly prominent and contested issue. While the expansion of typically low-density residential housing into the surrounding countryside can help to relieve housing supply constraints, it also generates losses of valuable farmland and natural ecosystems as well as increased pollution from greater car usage (Kahn, 2000). Such concerns have led to the creation of greenbelts and other urban growth boundaries, which restrict housing construction on undeveloped land. Examples of greenbelt policies abound, in diverse cities that include Amsterdam, London, Portland, Seoul and Toronto.

Greenbelts are designed to stop sprawl, but in doing so, they may impose substantial costs. By blocking housing development, they limit the supply of housing and drive up prices – the 'cost' of protecting the greenbelt. Beyond house price impacts, they may also affect the broader economy through losses in productivity and labour mobility (Hsieh & Moretti, 2019; Behrens et al., 2014). Such competing considerations make greenbelts an increasingly contentious policy option, with environmental concerns pitted against housing affordability, which has been brought into especially sharp focus given rising housing prices in many major cities. Despite the intense policy debate, the magnitude of the costs imposed by greenbelts are poorly understood.

In this paper, I develop a new framework for evaluating the housing market impacts of greenbelt policies and their potential interaction with other land use restrictions. This framework consists of a city where the demand for housing is increasing over time and a disaggregated housing supply and demand model with heterogeneous supply elasticities at the census tract level. When a greenbelt is introduced, some locations become restricted and the people who would have lived in those locations are forced to substitute to alternative destinations. This displaced demand can lead to building in other areas of the city, such as on unrestricted, developable land near the urban fringe or on under-developed land in the urban core, which can offset the housing supply impact of a greenbelt.

How large this displaced demand effect is will depend on how easy it is to build housing in the rest of the region. If housing supply is elastic elsewhere in the region, the supply and price effects from a greenbelt will be small because it is easy to develop other locations. But if housing supply is inelastic, such as if the greenbelt binds significantly or if land use regulations within the city are strict (Gyourko & Krimmel, 2021; Glaeser et al., 2006), the supply and price effects from the greenbelt could be substantial. This conceptualization illustrates how the impact of greenbelt policies depends on economic factors beyond the boundaries of the policy and cannot be thought of in isolation.

Using this framework, I evaluate the impact of the largest contiguous greenbelt in the world, the Ontario Greenbelt, which surrounds Toronto, Canada. The Ontario Greenbelt serves as a useful setting for evaluating greenbelt policies. In addition to its sheer scale, the introduction of the policy in the early 2000s allows for a before-after comparison of housing development using high-quality panel data, which is often challenging for other greenbelts introduced decades prior. Exploiting this policy variation, I present motivating evidence that the Ontario Greenbelt had a material impact on housing development patterns in the region using an event-study framework. I compare housing development trends in developable census tracts with more Greenbelt coverage to those with less coverage, before and after the policy was implemented, to see if it was binding. I find that while housing development in these census tracts evolved in parallel before the policy, census tracts more restricted by the Greenbelt saw considerably less development relative to the less restricted tracts after.

Although these results establish that the Greenbelt was a binding policy intervention, they do not answer the broader question of how the Greenbelt affected housing prices across the metropolitan area. One reason for this is that it is challenging to interpret the magnitude of the event-study results on housing supply due to the presence of spillovers. Since the Greenbelt potentially pushes housing development from Greenbelt areas into unrestricted ones, which serve as the control group, the size of the effect may be biased upwards. Another reason the event-study results do not answer the key research question is because they pertain solely on areas with developable land. This is done because it is challenging to draw conclusions about the impact of a greenbelt on census tracts in the urban core without more structure on the problem. For these reasons, I develop a model of housing supply and demand for the region that I can estimate.

The estimable model is based off the conceptual framework of greenbelts discussed earlier, with housing supply and demand curves that vary at the census tract level. I model the housing supply curves for each census tract as a function of observable tract-level characteristics, such as the share of developable land and the distance to the central business district (CBD), which allows for rich heterogeneity across the region. This setup also allows for housing supply elasticities to depend on observable land use regulations such as the greenbelt, which enters the model as a component of developable land, and 'urban growth centers' (UGCs), which are zones targeted for increased density. The housing demand curves are derived from a discrete, location choice model, where households choose the location and type of house to buy. In this setup, housing demand is a function, not only of prices in a given location, but of all locations across the metropolitan area. It is this feature, which allows the effects of a greenbelt to extend to the region as a whole.

To estimate the model, I take advantage of granular data on housing construction and transactions from 2000-2010 to recover census tract level housing supply elasticities and the price elasticity of demand for housing. Information about housing developments comes from Altus Group's Residential New Homes database, which covers the universe of housing projects in the GTA from 2000 to present, while housing transaction data comes from Teranet's GeoWarehouse. To address the endogeneity challenge inherent in estimating housing supply curves, I use an instrument recently introduced by Baum-Snow & Han (2023) for estimating housing supply elasticities at the census tract level. The instrument is the change in *simulated* residential market access (RMA), which is a proxy for the distance-discounted availability of job opportunities in a a given location. The reason this is simulated is because RMA, which can be computed using data on population and employment, is not a good instrument on its own due to its correlation with features of a census tract such as population, which is connected to housing supply. Instead, employment in a location is simulated using Bartik-style labour demand shocks that are uncorrelated with census-tract level characteristics. When estimating the results, I find that the housing supply elasticities vary substantially across the region, where locations with the most undeveloped land have the most elastic predicted supply elasticities (between 0.8 and 1.2) while locations with no developable land are considerably more inelastic (effectively 0).

For the housing demand estimation, I instrument for price using a new heritage designation instrument. Heritage designations prevent buildings from being demolished for new developments and therefore act as a supply shifter that can be used to identify the demand curve. I collect data on all the heritage designations in the Toronto area and the time they were enacted and compute a heritage designation exposure measure. Conditional on a number of control variables, identification comes from comparing similar census tracts exposed to different degrees of heritage restriction. I find that the annual housing demand price elasticity ranges from between -0.66 and -1.14, which is consistent with other estimates used in the literature.

Using these estimates for housing supply and demand, I can simulate the effect that the introduction of the Ontario Greenbelt had on the housing market. To do this, I adjust the share of developable land in a census tract by removing the Greenbelt portion, which makes the supply curves more elastic and then solve for the new equilibrium price and quantity in the market under this counterfactual scenario. I find that the Greenbelt had a substantial impact on housing prices just five years after the policy was introduced. The Greenbelt led to a reduction in housing construction of 60% in Greenbelt affected areas while boosting housing construction by around 4.5% outside of the Greenbelt relative to the amount of construction had the Greenbelt not been in place. This housing supply shock translated into an increase in housing prices of 2.25% on average across the region by 2010, which corresponds to an increase in housing rent of around C$550 a year.

While these effects are somewhat large, for renters this amount to almost 1% of annual pre-tax income, I show that land use regulations elsewhere in the region amplify the impact of the Greenbelt on housing prices. One finding from the housing supply model is that census tracts located in an 'urban growth center' (UGC) exhibit much larger condominium supply elasticities compared to non-UGC tracts. These UGCs, which are regional subcenters targeted for greater density, do not contain any binding incentives, but instead give targets to municipal planning officials. If these designations were expanded to more census tracts, this could increase the housing supply response within the city. I simulate a counterfactual to capture what would have happened if UGC targets were extended to additional census tracts within 1km of the existing boundaries. I find that this small change to land use regulation could entirely neutralize the housing price impacts of the Ontario Greenbelt. Doing so would lead only a 1% increase in housing costs compared to a 2.25% increase before.

These results have clear implications for how policymakers approach land use reform. Given the choice between removing greenbelts and reforming land use restrictions, a policymaker could achieve similar price effects based on the results of the counterfactual exercises. However, the two policies differ in how they affect other important outcomes. Specifically, greenbelts also provide a number of benefits both directly through open space amenities (Anderson & West, 2006; Koster, 2023) and indirectly through the promotion of density. The benefits of increased density tend to outweigh the costs (Combes et al., 2019) because of agglomeration effects (Ahlfeldt & Pietrostefani, 2019), but also because of other factors such as reduced traffic congestion which leads to less air pollution (Gibson & Carnovale, 2015). As a result, greenbelts would appear to be the preferable policy approach given the two alternatives.

Yet my results also highlight one reason why land use reform within cities may not occur. The costs of higher housing prices fall predominantly on renters and first-time home buyers as opposed to incumbent homeowners, who benefit from higher housing prices. In line with the homevoter hypothesis (Fischel, 2004), homeowners have a financial interest in preserving the value of their home and will oppose efforts to build more housing in their neighbourhood. Empirical research has supported this theory in the United States (Duranton & Puga, 2023; Dehring et al., 2008) and in Toronto specifically, where city councillors with more homeowners in their ward are shown to oppose more housing construction (Fang et al., 2023). If homeowners also value the benefits of greenbelt policies, then the following situation holds in practice: a restrictive greenbelt, restrictive land use policies within cities, and higher prices for renters.

This paper makes three main contributions. It contributes to the literature looking at the impact of greenbelt and other anti-sprawl policies, where it is the first paper to incorporate panel data into a model of an urban area to understand these impacts. Prior work from Walsh (2007) and Koster (2023) examined the effect of greenbelt style policies across a metropolitan area using a static, model-based estimation approach. Koster (2023) studied the impact of the English Greenbelt, which was introduced in the 1950's and covers 13% of England, and argued that the amenity benefits from the policy made the Greenbelt welfare-improving. Due to the more recent introduction of the Ontario Greenbelt, I am able to study how the Ontario Greenbelt actually altered development patterns looking before and after the policy. Other empirical work studying urban growth boundaries that did look before and after a greenbelt policy focused solely on the areas close to the boundary, but not the metropolitan area as a whole (Cunningham, 2007; Deaton & Vyn, 2010). There is also a theoretical literature that looks at the impact of greenbelt policies (Quigley & Swoboda, 2007) and compares urban growth boundaries to alternative policy solutions (Brueckner, 2001, 2007; Anas & Rhee, 2007; Bento et al., 2006).

This paper also contributes to the literature studying land use regulation and development frictions by looking at the interaction of land use policies within cities with those on the urban fringe. Land use restrictions within cities have received significant attention in recent years as researchers have sought explanations for rising housing costs (Anagol et al., 2021; Kulka et al., 2023; Glaeser & Gyourko, 2018; Turner et al., 2014; Saiz, 2010; Glaeser & Ward, 2009; Ihlanfeldt, 2007; Mayer & Somerville, 2000). However, much of the land use literature focuses on policies such as minimum lot sizes and maximum built-area ratios, which regulate within-city development more than development on the urban fringe. This paper is the first to incorporate heterogeneity in housing supply elasticities into a model of a greenbelt to explore how the impact of a greenbelt policy depends on the stringency of existing regulation.

Finally, this paper contributes to a newer literature examining the vertical structure of cities. Historically, most literature has focused on the horizontal expansion of the urban area, but recent work from Ahlfeldt & Barr (2022) has highlighted the importance of looking at the vertical structure of cities to understand how cities will grow in the future (Ahlfeldt & McMillen, 2018; Helsley & Strange, 2008). This paper separates housing supply into single family housing and condominium apartments, where over 250 apartments were constructed in this period that were over 25 storeys tall. This paper offers a potential explanation for the increased skyscraper construction: as cities reach either policy induced boundaries like greenbelts or natural boundaries such as the ocean, skyscraper construction will become more profitable as demand is displaced from the urban fringe into the urban center.

2 Context & Data

2.1 Policy Context

The city of Toronto is a major metropolitan area that has grown significantly in recent decades. Toronto is the largest city in Canada and the sixth largest metropolitan area in North America. From 1991 to 2001, the Toronto area grew by 770,000 residents, from 3.9 million to 4.7 million, or 20% (Statistics Canada, 2016), with much of this growth coming in the form of immigration. As the city expanded rapidly around the turn of the millennium, there were growing concerns about the destruction of environmentally sensitive land, which prompted protests against large developments and calls for restrictions on sprawl.

In the early 2000s and in response to political pressure, the Ontario government created the largest contiguous greenbelt in the world around the Greater Toronto Area (see Figure 1). Greenbelt style protections had already existed in the region starting in 1973 with the Niagara Escarpment, which runs mostly outside of the area studied in this paper. Then, the 500,000 acre Oak Ridges Moraine was protected in 2001. Finally, the *Greenbelt Act (2005)* was introduced which combined the Niagara Escarpment, Oak Ridges Moraine and newly protected land into a single Greenbelt almost 2 million acres in size. The Ontario Greenbelt protects some of the best agricultural land in Canada as well as forests, wetlands and the headwaters that are essential for providing clean water for the region.

The initial introduction of greenbelt policies by the Ontario government in December 2001 came about as a surprising departure from the government's previous stance on environmental issues. The provincial premier at the time, Mike Harris, had previously extolled the virtues of a "common-sense revolution", where the government reduced its involvement in several policy portfolios including the environment. However, a water contamination scandal in 2001 prompted Harris' resignation and a shift towards more environmental policy engagement including the creation of the Oak Ridges Moraine (Winfield, 2012). This shift was not successful electorally as the government lost the provincial election in 2003 to the Ontario Liberal party. This sudden turnaround however, does make the initially policy shock somewhat unanticipated.

After the protection of the Oak Ridges Moraine, the completion of the Greenbelt in 2005 was more anticipated, but there was still uncertainty over the location. The Ontario Liberal party campaigned in 2003 on a promise to further expand the Greenbelt and ultimately won the election. Over the following year and a half, there were several consultations about the plan, but the government did not commit to a specific boundary in order to prevent speculation on land around the boundary. While there was lobbying from several parties including developers, farmers and environmental groups, the ultimate Greenbelt boundary left several

Figure 1: Map of the Ontario Greenbelt

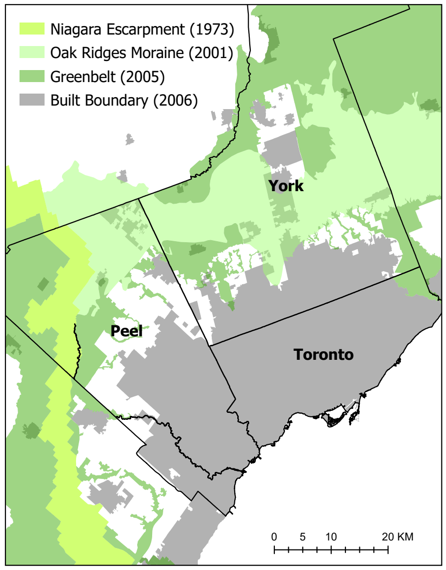

Note: Greenbelt protections in the Greater Toronto Area were introduced in three phases. There were initial protections on the Niagara Escarpment in 1973, which mostly bypasses the region around Toronto. Then, in 2001, the Oak Ridges Moraine was protected, which runs across the northern portion of the city. Finally, the *Greenbelt Act* (2005) was introduced, which brought the prior two protections under the same umbrella and expanded protections to 600,000 acres of land. Black lines represent census divisions.

unexpected winners and losers, with some developers suing the provincial government for including their parcels of land within the Greenbelt (Bradburn, 2022).

The Ontario Greenbelt has been strictly enforced in the years since. The Greenbelt Act (2005) stopped all changes to official plans and development applications in the protected region that were not approved as of December 2004. This meant that some building still occurred on the Greenbelt after 2005, but only through applications that were accepted prior to this cutoff. The Greenbelt boundary has not been changed since 2005 other than to *add* a number of urban river valleys in 2017. In 2022, the premier of Ontario tried to remove 7,200 acres of protection from the Greenbelt, however he reversed the decision in the face of significant political pressure. While there are decennial reviews of the policy, it is meant to be a permanent feature of the planning framework for the region, which differs from some other urban growth boundaries, such as the one in Portland, Oregon, which is updated fairly regularly to population pressures. The permanent nature of the Ontario Greenbelt means that the effects of the greenbelt on housing supply and prices would only be expected to increase over time.

2.2 Data

To evaluate the impact of the Greenbelt on housing development requires combining several sources of data on housing in the Greater Toronto Area. This paper focuses on the three largest regions in the GTA: the City of Toronto and the regions of Peel and York during the years 2000-2010. The primary data source for this analysis is Altus Group's New Residential Homes database, which contains information on all 11,500 plus housing development projects in the GTA from 2000 to the present day. This dataset includes information on the location of the project, the number of units, the type of unit (apartment, row or single), the size of the units the average price of the units and the developer who built the project as well as detailed information on the timing of the project starting with the date a project first sold an assignment to the occupancy date. The information on the first date sold for the project is useful as it gives a precise time when a project would have "entered" the market. I combine the Altus data with information from the 2001 Canadian Census of Population to construct an annual time series of the number of houses of each type available to be purchased in a census tract in a given year.[1] Census tracts contain between 2,500 and 8,000 people and there are 832 census tracts in the GTA.

I add information on housing sales using transaction-level data from Teranet's GeoWarehouse. This is the official land registry for the province and contains all transfers of land, the prices they were transacted at, the type of unit (freehold or condominium), the date of the transaction and a PIN that can be matched to Teranet's parcel data to geolocate each property. In total there were over 1 million housing transactions between 2000 and 2010. The parcel data contain information on lot sizes, which I supplement with publicly available data on housing footprints from Peel and York regions and the City of Toronto. I use this information to create an annual series of housing prices of single family homes and apartments for each census tract. In addition, I create a price index that strips away variation in the characteristics of houses sold across the years to ensure the variation in price is not driven by compositional effects.

Finally, I add information on the land use status of each parcel. Agriculture and Agri-Food Canada's (AAFC) Land Use time series is a series of satellite maps at a semi-decadal frequency dating back to 2000. The spatial resolution of the maps is 30x30 meter pixels classified into 7 broad categories including forest, cropland and settlement, with breakdowns at the settlement level into categories including settlement housing, roads, vegetation and high reflectance areas. Based on these designations, I assign the number of square feet of

[1]This does not account for any housing that is torn down during this period and not rebuilt in some capacity. This could be an issue if say, an apartment building is replaced by another, larger apartment building. However, there appears to be little evidence of this type of activity during this period.

Table 1: Summary Statistics for Census Tracts in 2010 by Housing Type

	Mean	Min	Median	Max
Condominiums				
# Units	1,188	5	795	14,042
Δ # Units 2000-2010	318	0	0	12,242
Sale Price ($)	300,362	65,081	273,037	1,128,263
Δ Sale Price 2000-2010 (%)	67	-49	58	548
Lot Size (sqft)	.	.	.	.
Footprint (sqft)	.	.	.	.
Distance to CBD (km)	17	0	17	74
Census Tract Size (acres)	602	13	203	27,060
Undeveloped Land %	4	0	0	90
Greenbelt %	1	0	0	87
Single Family Homes				
# Units	1,435	120	1,177	18,472
Δ # Units 2000-2010	235	0	0	15,048
Sale Price ($)	486,532	196,761	446,017	1,198,905
Δ Sale Price 2000-2010 (%)	87	6	82	670
Lot Size (sqft)	19,992	2,065	8,939	1,130,959
Footprint (sqft)	2,908	918	2,120	29,353
Distance to CBD (km)	18	1	17	82
Census Tract Size (acres)	985	30	219	40,857
Undeveloped Land %	6	0	0	94
Greenbelt %	2	0	0	92

Note: The following summary statistics are for the year 2010 and for single family homes and condominium apartments separately. There are 721 census tracts with single family homes and transactions and 490 census tracts with condominiums. There is no information on the lot size and footprint for condominium apartments.

each land use type to each parcel in 5 year intervals. I then aggregate up to the census tract level and compute the share of land that is developable (cropland or forest) and the share that is not developable (the remaining categories). I also use the location of each parcel to determine whether it falls inside or outside the Greenbelt and to compute the distance to the Central Business District (CBD) of the City of Toronto (in this case this is the City Hall).

Table 1 presents basic summary statistics of the resulting dataset. For the table, I include any census tract-by-unit type for which I observe greater than zero initial units and sales for more than two years.[2] The resulting dataset includes 721 census tracts with single family homes and 490 census tracts with condominium apartments. I do not have any information on the lot size and footprint of condominium apartments as they fall within a building.[3]

[2] I impute prices for some census tracts with missing transactions for some years using linear interpolation. There are only around 30 census tracts with any imputed values. Results from the paper that drop these tracts or the missing years yield similar overall results.

[3] It should be noted here that these distinctions are made based on ownership structure and not necessarily building structure. This is relevant only for a subset of cases where some townhouses will be classified as condominiums, while others will be considered single family. There are also some areas with only rental housing, which does not factor into this analysis.

The average census tract-by-housing type observation has just over one thousand units, but there is significant dispersion in how much construction occurred over the 2000s. In some census tracts, there were upwards of 12 to 15 thousand units built, while over half of census tracts saw no housing constructed. Sale prices increased substantially over this period in the Toronto region, rising 87% for single family homes and 67% for condominiums on average.[4] Census tracts which include single family homes tend to be twice as large on average as those which include condominiums. This makes sense as census tracts with condominiums are more dense and smaller boundaries are needed to capture the standard number of people. Finally, most census tracts have no undeveloped land and no Greenbelt land, but those that do have undeveloped and Greenbelt land can have substantial amounts. These statistics show that there is also considerable heterogeneity across the region in terms of construction, prices and observable characteristics. This heterogeneity has important implications for evaluating a greenbelt policy because greenbelts are not applied randomly across an entire region. Greenbelts are typically placed in areas that would be expected to see development, while pushing housing demand pressures into places with less developable land. When this heterogeneity is ignored, the housing supply and price impacts of a greenbelt would be underestimated.

3 Motivating Evidence

Although basic economics suggests that a large greenbelt like the Ontario Greenbelt should affect the housing market, it is not certain. Some local advocates argue that because the Greenbelt does not block all potential development in the region (see the white areas in Figure 1), the effect of the Greenbelt on housing prices is minimal and therefore there is no reason to remove the protections. Another reason why the Greenbelt may have minimal economic consequences is if it is located in areas where neither households nor developers want to develop. If this is true, then the Greenbelt would not distort housing development patterns or prices at all.

In this section, I will present motivating evidence that the Ontario Greenbelt did appear to restrict where housing was able to be built. First, I will look at the raw trends in development by housing type and across the region. Then, I will compare trends in housing development patterns in census tracts affected by the greenbelt to those less affected using an event-study design. Although the magnitudes of these effects cannot be interpreted causally due to concerns over spillovers in development within the region, the event-study framework

[4] Extreme transaction values were curbed from the sample. Those in the top and bottom percentiles were cut to avoid extreme outliers driving results.

Figure 2: New Units Brought to Market Over Time by Type in the GTA

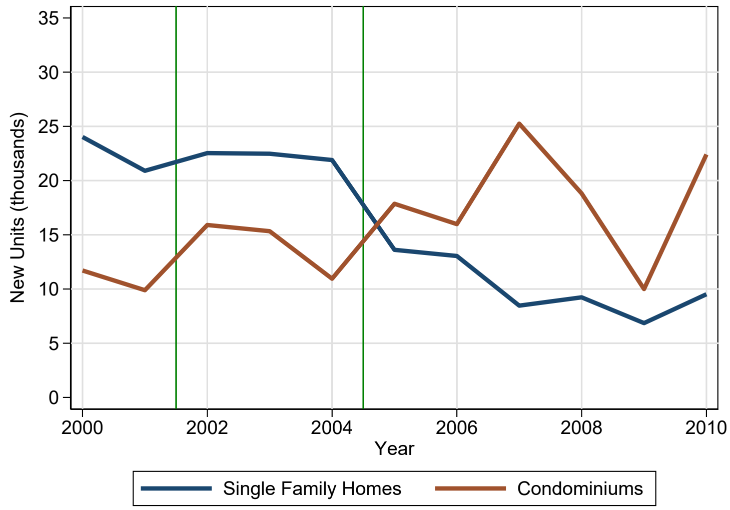

Note: This figure shows the trends in how many units were brought to the market in each year by type. The date used is the date of the first unit sold in a development project. This date usually occurs around the same time that building permits are acquired and that the construction process is getting underway which is a good proxy for developer intentions and behaviour. The dates of the Greenbelt policies are indicated with vertical lines.

will demonstrate that the Greenbelt did appear to cause a sharp break in development trends. This suggests that the Greenbelt is a binding policy intervention.

3.1 Descriptive Facts

Housing development in the GTA underwent significant change during the period of study. In this section, I define housing development as the process of bringing a project to the market and offering units up for sale. Typically, the process of selling units occurs concurrently to the acquisition of building permits and start of the construction process so I use the date that the first unit in a development project is sold to represent the beginning of a development project generally. One advantage of using the date of the first sale is that unlike building permits, which can be acquired and never acted upon, once units are sold there is pressure from buyers to follow through with the project.[5] Another advantage is that this date reflects when the project effectively starts competing with other housing units in its neighbourhood. Throughout this paper, references to housing development or construction timing will be referencing this date.

In Figure 2, I show the trends in new units brought to the market from 2000-2010. For

[5] I only include projects that are completed in the analysis. This means that projects that started selling units, but were not completed are not included in the analysis.

the first half of the decade, single family homes were being developed at a rate of between 20 to 25 thousand units a year. However, starting in 2005 this number began declining to below 10 thousand units a year, a reduction by over half. Conversely, the trend for condominiums increased during this period, rising from around 10 thousand to as much as 25 thousand in 2007. By 2005, condominium construction became the dominant form of construction in the region. These trends suggest that there was a discernible change in housing development patterns during this time.

Figure 3 plots development patterns across the region from 2000-2010. I show that the majority of single family home construction occurred on the urban fringe in a handful of census tracts that saw 5 to 15 thousand new units constructed. Within the urban footprint there was very little construction of single family homes during this period. The majority of condominium construction occurred in a very concentrated set of locations in downtown Toronto and some of the regional subcenters, where there were 5 to 10 thousand units built. There were only a few apartments built in the suburban and rural areas.

The Ontario Greenbelt largely affected single family housing development on the urban fringe and therefore the trends and patterns observed are consistent with a story where the Greenbelt slowed urban sprawl. Because there is an uptick in condominium construction, the downward pattern for single family homes cannot solely be explained by aggregate economic effects such as the Great Recession. It does remain possible however that there were secular declines in suburban housing demand or supply shocks to that sector that coincided with the policy and as such, these are merely suggestive figures. To address this in greater detail, I will examine these trends in an event-study framework in the next section.

3.2 Event-Study Regressions

In order to determine whether the declines in single family housing construction are driven by the Greenbelt or some other factor, I employ an event-study framework where I compare census tracts that were more or less exposed to the Greenbelt policy. In this way I can see whether locations that became highly restricted by the policy saw declines in homebuilding relative to those that were less restricted. Because the Greenbelt protects undeveloped land and not developed land, the focus in this section is on single family homes and not on condominiums, which tend to be built mostly within existing urban areas.

Figure 3: Total Units Brought to Market by Census Tract and Type, 2000-2010

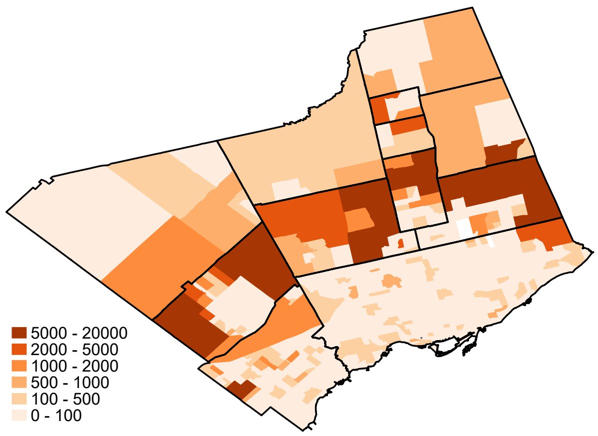

(a) Single Family Homes

(b) Condominiums

Note: This figure presents the total number of units brought to the market between 2000-2010 by census tract. The black lines represent census subdivisions (or municipalities). For single family homes, the majority of construction occurs at the urban fringe with very little in the urban core. For condominiums, development occurs predominantly in downtown Toronto in very concentrated areas.

3.2.1 Empirical Specification

I estimate the following regression equation to study the effect of the Greenbelt on housing supply.

$$\ln H_{jt} = \sum_{g=-G}^{-2} \alpha^g D_{jt}^g + \sum_{k=0}^{K} \alpha^k D_{jt}^k + \nu_j + \eta_t + \varepsilon_{jt} \tag{1}$$

The variable of interest is the log of housing stock in a census tract j at time t for the years 2000-2020.[6] D_{jt} is a treatment indicator for whether a census tract is exposed to the Greenbelt policy at a lead time of g or lag time of k. I control for census tract (ν_j) and year (η_t) fixed effects, which makes this a two-way fixed effects specification. The parameters of interest are the α's, which as I will discuss shortly, reflect a relative, but not necessarily causal effect of the Greenbelt on housing development patterns due to the presence of spillovers.

Treatment in this specification is defined as having more than half of the census tract covered by the Ontario Greenbelt. Because the share of land in a census tract covered by the Greenbelt is continuous, a threshold needs to be chosen to represent the point at which one considers a census tract to be treated. The reason this threshold is not set at 100% is that most census tracts are not entirely covered by the Greenbelt. In addition, partial treatment by the Greenbelt may still affect housing supply either by increasing competition

[6]The Altus New Residential Homes database extends to the present day, which means I can run this analysis until 2020. This differs from the broader dataset because I only have transactions data up to 2010.

for land parcels and driving up the cost of land, making desirable parcels for construction undevelopable or by reducing the number of parcels available for land assembly (Brooks & Lutz, 2016). The reason this threshold is not set at anything over 0% is because some locations may only have a small share of land covered by the Greenbelt and therefore may not "behave" like a census tract with restricted land supply. I use 50% as the threshold and show that the results are robust to some alternative thresholds.

Treatment is staggered in this empirical setting as some census tracts receive treatment due to the *Oak Ridges Moraine Protection Act* in December 2001, while others are only treated after the introduction of the *Greenbelt Act* in February 2005. A tract that is partially protected in 2001, but less than 50%, that crosses the 50% threshold in 2005 is considered to be treated in 2001. Given the recent literature on the problems with staggered difference-in-differences with heterogeneous and dynamic treatment effects (Callaway & Sant'Anna, 2021; Sun & Abraham, 2021; Goodman-Bacon, 2021), which I suspect may be present in this setting, I also estimate the effects of the Greenbelt using the method of Callaway & Sant'Anna (2021). This approach involves estimating average treatment effects for each group-time treatment cohort and compares them exclusively to not yet treated units, which avoids the issue of comparing to already treated units. When I estimate the model using this approach I find the results to be very similar to the original event-study findings.

The control group is the set of census tracts where at least 25% of the land is deemed developable in the initial period of 2000 and where they are not already considered covered by the Greenbelt. The reason the remaining census tracts do not serve as a good control group is because these are places that have ostensibly been "treated" already in the sense that they have restrictions that prevent them from building - namely, a lack of developable land. These tracts see hardly any development in either the pre-period or the post-period. Appendix Figures 2 and 3 show that census tracts where undeveloped land is less than 25% of the tract have virtually no Greenbelt exposure and see almost no housing development on average. I also show that results are robust to different cutoffs.

Appendix Figure 1 shows which census tracts fall into the treatment, control and undevelopable groups. In the map, the City of Toronto and the main suburb, Mississauga (which is west of Toronto), are mostly developed already in 2000, leaving most of the analysis to be elsewhere around the urban fringe. In general, the treated tracts are farther away from the urban fringe compared to the control group although there is some variation to this along the corridor of suburbs to the north of the city. While the treatment and control groups may not be perfectly balanced in terms of characteristics, I will show that parallel trends holds.

One challenge in interpreting the estimates of this regression is the presence of spillover effects from the treated group into the control group. When the Greenbelt is implemented, some development may shift from the treated areas to the control areas - in fact, this is exactly what one would predict to happen. As a result, the magnitude of the estimates cannot be interpreted as the causal effect of the Greenbelt. This is because the amount of housing built outside the Greenbelt would be larger than the "true" counterfactual leading to an upward bias in the estimate. However, the parameters of interest can still be interpreted as a relative effect between the treated and control groups because the presence of spillovers is dependent on the Greenbelt distorting the housing development behaviour in the region in the first place. This limitation however, highlights one of the primary reasons why a model is necessary to disentangle these spillovers from the true effect of the policy.

There are two additional threats to identification in this setting. First, there may be concerns that the policy was anticipated leading developers to push projects forward in time and leading to an overestimate of the size of the effect. While this may be true for the second phase of the policy, this was unlikely for the first phase, which came about as a surprise. I also show similar trends for each treatment group separately and see that the effects carry on many years into the sample, which suggests anticipation is not driving the results. Second, there may be concerns that the policy boundaries were endogenous to developer interests. If this was the case, then there would be selection concerns that only places unlikely to be developed received the Greenbelt. However, I present a few pieces of evidence to suggest that any selection concerns were orthogonal to development patterns. The first is that several developers were impacted by the policy. I take a sample of 100 large parcels of land that fell within the Greenbelt in 2005 and find that 30 of them were owned by numbered companies or companies with a variation of "Realty" or "Developments" in their name, which is indicative of developer interest in the location. A second is that developers sued the government after the Greenbelt policy for the lost value of their land (Winfield, 2012). Finally, the government had campaigned on the policy to bolster its environmental credentials and if anything, included more land than was expected initially (Winfield, 2012). Together, this suggests that it is unlikely the ultimate Greenbelt boundaries were influenced by developer lobbying in a manner that would bias the results.

3.2.2 Main Results & Robustness

Figure 4 presents the results of the event-study regression. Prior to the implementation of the policy the treatment and control group follow parallel trends in terms of housing development. However, after the policy was introduced there was a steady downward trend in the amount of housing in the Greenbelt tracts relative to the non-Greenbelt tracts. Ten years after the treatment date, the difference between the treated tracts and control tracts reached over 20% and fifteen years after, the difference surpassed 30%. The effects are

Figure 4: Housing Development in Greenbelt Tracts Versus Non-Greenbelt

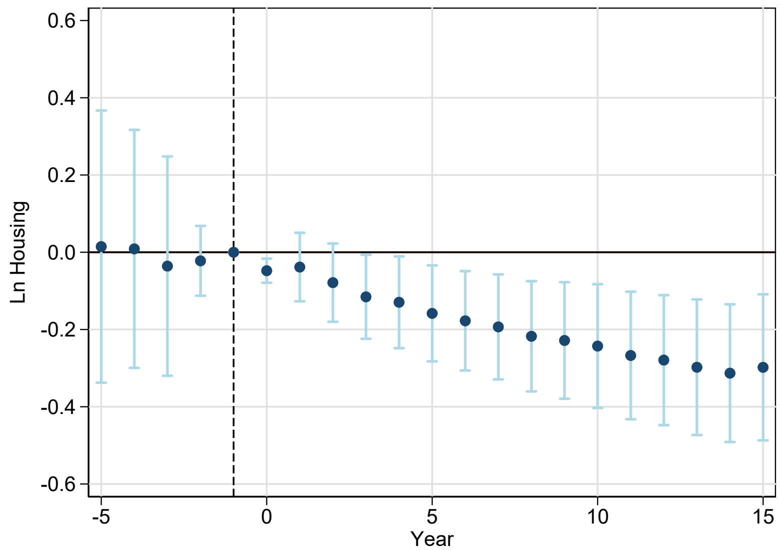

Note: This figure depicts the estimates from a linear event-study model with the log of housing in a census tract as the dependent variable and the timing variable the distance to treatment in either 2001 or 2005. Treatment is defined as having more than 50% of the census tract in the Greenbelt. The coefficients can be interpreted as the percentage difference in housing between the treated and control tracts. Standard errors are clustered at the census tract level.

statistically significant starting four years after the policy, suggesting that it took time for the effects to emerge, which is consistent with the longer timelines of housing development projects. Overall, these results show that untreated census tracts followed a similar pattern relative to treated tracts prior to the Greenbelt being introduced, but after the policy, they saw much more rapid development compared to those in the Greenbelt.

I validate these results using several alternative approaches. First, the dynamic nature of the effects raises concerns that the results may be biased due to a contaminated treatment group. To address this, I estimate the model again using the Callaway & Sant'Anna (2021) estimator and plot the treatment effects for each treatment cohort separately in Appendix Figure 4. The figure shows that for both treatment groups, the magnitude of the effect is fairly stable, although the standard errors are large, particularly for the 2001 treatment group. In addition, the pre-trends for the 2005 treatment group remain relatively stable. These results suggest that heterogeneous and dynamic treatment effects are not driving the results of the event-study.

Second, I investigate if using a continuous treatment variable will also show that the Greenbelt had an effect. To do this, I use the same setup as for the discrete case, but the treatment indicator becomes $D_{jt} = GB\% \times \mathbb{1}(\text{year} >= \text{treat year})$, where the treatment will vary according to treatment intensity. Under this specification, the parameter of interest is identified not only from being treated or not treated, but according to treatment intensity, where more highly treated tracts are compared to less treated tracts.

I estimate this regression using several different thresholds for the undeveloped land share that determine the sample of census tracts. In Appendix Figure 1, I show that regardless of the sample used, more intense Greenbelt treatment leads to less housing supplied after treatment. These results suggest that a 10% increase in Greenbelt coverage at the census tract level is associated with a 1.6-4% decrease in housing supply over the post period. The effects become stronger as census tracts with more developable land are compared and are only significant at the 30% undeveloped land share threshold. One explanation for this is that the Greenbelt does skew towards more undeveloped tracts and the control group becomes more comparable at those levels. Another explanation is that in places with more elastic housing supply, the level of development in the control group is expected to be higher. These results suggest that the effect of the Greenbelt on housing supply patterns holds up whether one looks at treatment as discrete or continuous.

I also present results where I vary the treatment and control thresholds for the main event-study results. I present these results in Appendix Figures 5 and 6. These results show that choosing different cutoffs for the Greenbelt threshold and the share of developable land needed to be in the control group does not change the overall trends between the treatment and control groups. Another comparison I present is where I separate the control group into partially treated Greenbelt tracts and those not treated. That is I compare trends in census tracts that are less than 10% covered by the Greenbelt to those more than 50% covered and those between 10-50% covered by the Greenbelt. I present these results in Appendix Figures 7 and 8. I find that the effects of the Greenbelt are attenuated when comparing just to the census tracts that receive little to no treatment while the census tracts that receive 10-50% treatment actually see an increase in housing supply after the policy relative to the less than 10% treated group. These results are consistent with the story of spillovers, where the tracts that are between 10-50% treated by the Greenbelt are geographically closer to those that are more than 50% treated, and are therefore more likely substitutes for housing development in the aftermath of the policy. This is evidence that the Greenbelt may induce spillovers to nearby areas, but less so to regions farther away.

These results demonstrate that the decline in single family home construction during this period is likely linked to the restrictions from the Greenbelt. Although not entirely causal, the sharp break in trend suggests that this was not driven by secular changes in the economy, but by the policy. However, these results are limited in scope - they can only tell us what happened in and around the boundary of the Greenbelt. Of greater interest is what

Figure 5: A Conceptual Model of Greenbelt Implementation

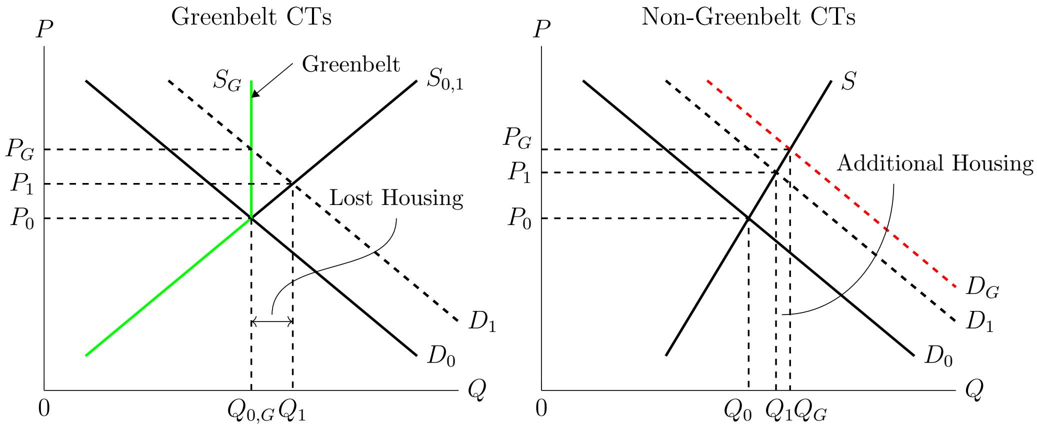

Note: This figure plots a conceptual model for the implementation of a greenbelt policy and the impacts on census tracts restricted by the greenbelt to those unrestricted by the policy. The model assumes a positive demand shock, such as from immigration, that incentivizes a city to grow.

happened to the housing market across the region. This will be explored using a model presented and estimated in the following sections.

4 Model of the Housing Market

A model of the housing market is required to answer the counterfactual of what would have happened had the Greenbelt never been implemented. This involves answering both how much more housing would have been built in areas where the Greenbelt was implemented as well as how much *less* housing would have been built in the rest of the city in response to the policy. Given that the policy crosses any standard political boundaries, it is necessary to conceptualize how housing markets behave at a sub-municipal level. This is a challenge because housing supply responses vary dramatically across a region, especially if we assume the existing stock of housing is fixed in the short run. As seen in Section 3.1, the pattern of development is highly uneven across space with single family homes built around the urban fringe and condominiums built in small pockets in regional subcenters. In this section, I will lay out a model of the housing market that will provide a useful conceptual basis for thinking about the impact of greenbelt policies as well as a clear pathway to estimation that incorporates substantial heterogeneity.

4.1 Conceptual Framework

This section formalizes a conceptual model of greenbelt policies that will motivate the rest of the paper. Greenbelt policies affect a substantial amount of land around cities which should serve as a shock not just to the areas directly affected by the policy, but the rest of the metropolitan area as well. The rest of the metropolitan area will be affected primarily through a reallocation of housing development, where developers who would have built in restricted Greenbelt areas now build in unrestricted areas. I formalize this argument and highlight the key nuances in Figures 5 and 6.

For the conceptual framework, one can imagine a metropolitan area that is growing over time either due to internal migration, immigration or increased demand for space. A growing city is a prerequisite for a greenbelt to have any effect because otherwise the policy would not restrict any housing demand. Within the city, there will be census tracts, which can be thought of as neighbourhoods, which have heterogeneous supply curves and are subject to housing demand shocks. The slope of the supply curves may depend on characteristics of the census tract such as the amount of developable land or the number of land use restrictions on housing construction. Finally, suppose that a government introduces a greenbelt policy that bans housing development within its boundaries. I will not assume that the greenbelt is binding to the entire urban fringe, which allows for the possibility that other suburban areas may continue to expand.

Figure 5 plots supply and demand curves for a given census tract within the Greenbelt and a given census tract outside the greenbelt. I assume that a housing demand shock (for example, from immigration) shifts housing demand from D_0 to D_1 over the course of a year. In the absence of a greenbelt, one would move along the supply curve to a higher quantity and a higher price, P_1, in both census tracts. However, if a greenbelt is introduced and is completely binding, it renders the supply curve highly inelastic in the greenbelt census tracts, where no more housing can be built regardless of the price. This would increase the price in the greenbelt tracts up to an even higher price, P_G, and would also lead to a lower level of housing provided. The decrease in housing development would mean that some people would no longer have housing and would be forced to substitute towards housing either in non-greenbelt census tracts within the metropolitan area or cities outside the region entirely. The substitution to other areas within the metropolitan area is represented by D_G in the non-greenbelt panel. The more restrictive a greenbelt is, meaning the more census tracts it covers, the larger this shift in demand will be.

The extent to which this shift in demand translates into construction within the region and higher prices will depend on the elasticity of housing supply in the other areas of the city. In Figure 6, I plot different scenarios of how housing demand substitution can translate into

Figure 6: A Conceptual Model of Greenbelt Implementation

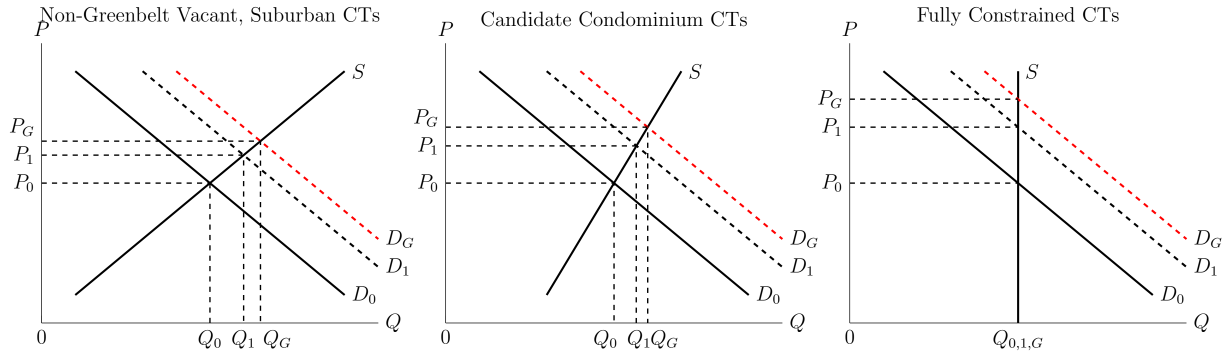

Note: This figure plots a conceptual model for the implementation of a greenbelt policy and the impacts on census tracts of different types within the city. Census tracts with more elastic supply will see larger supply responses and lower price increases compared to those with more inelastic supply curves.

construction and prices throughout the rest of the metropolitan area. If housing supply is more elastic, as it might be around the urban fringe where there is plenty of land to develop, there may be more construction with lower corresponding prices. In this scenario, greenbelts simply relocate housing to other suburban areas, which may lead to smaller impacts on prices and quantities, but would mean the greenbelt does little to curb urban sprawl. A second scenario is where housing supply is somewhat responsive through condominium construction within the urban footprint. In this case, although there is no developable land to build outwards, there could be construction upwards. Depending on how elastic this supply is, the displaced demand from the greenbelt could be met by the supply of denser forms of housing, which means that price effects could be small without leading to more urban sprawl. Finally, there is a scenario where housing supply is unresponsive to demand shocks due to a lack of available land, zoning regulations and local opposition to development. In this scenario, the displaced demand from the greenbelt leads entirely into higher housing prices and people are pushed out of the region.

This conceptual framework highlights how housing supply responses throughout a metropolitan area play a key role in determining the economic consequences of greenbelt policies. When housing supply within cities is sufficiently elastic, a city can more easily accommodate a greenbelt policy without causing prices to increase significantly. However, if supply is too inelastic, then price effects are going to be large and lead to emigration from the region. In the next section I will impose some parametric assumptions on these supply and demand curves in order to estimate housing supply and demand elasticities with variation at the census tract level. Using these estimated curves, I will then be able to simulate the effect of the policy using the logic of this conceptual framework.

4.2 Model of Housing Supply

The literature on housing supply has traditionally focused on differences in housing supply elasticities *across* cities (Saiz, 2010). It is not until recently that greater attention has been paid to differences in housing supply elasticities *within* cities (Baum-Snow & Han, 2023). Understanding within city differences in housing supply responses is crucial for understanding the impact of place-based policies that shift housing demand around within a city. In the context of a Greenbelt, if the housing supply elasticity is assumed to be constant throughout the entire region, then the true effects of the Greenbelt may not be accurately captured. This is because greenbelts are typically located in areas with the most elastic housing supply and push demand into areas with less elastic housing supply, which would lead to larger estimates of the policy impact. In this section, I will derive a housing supply curve for each census tract stemming from a developer's optimization problem.

4.2.1 Developer's Problem

Suppose in each census tract j there is a developer of type $i \in \{S, A\}$ for single family homes and for condominium apartments. Each developer of type i in location j has a unique cost function, $C_{ij}(H_{ijt})$, that depends on the amount of housing in a census tract in year t, H_{ijt}. These developers operate in a competitive market and take housing prices in their census tract as given.

The cost function for a developer is assumed to be of constant elasticity form:

$$C(H_{ijt}) = \rho_{ijt} \left(\frac{H_{ijt}}{\alpha_{\mathbf{ijt}}} \right)^{\frac{1}{\rho_{ijt}}}$$

where ρ_{ijt} and α_{ijt} are census tract by type level parameters. If $\rho_{ijt} < 1$, this is a convex function, where the cost is increasing in the amount of housing in a census tract. The marginal cost for a price taking developer is then:

$$\frac{\partial C(H_{ijt})}{\partial H_{ijt}} = \left(\frac{H_{ijt}}{\alpha_{\mathbf{ijt}}} \right)^{\frac{1}{\rho_{ijt}}-1} \frac{1}{\alpha_{\mathbf{ijt}}}$$

Redefining $\rho = \frac{\varphi}{1+\varphi}$ and with profit maximizing developers, who set $P_{ijt} = MC_{ijt}$, the housing supply curve can be written as:

$$H^S_{ijt}(P_{ijt}) = \alpha_{\mathbf{ijt}}{}^{1+\varphi_{ijt}} \left(P_{ijt}\right)^{\varphi_{ijt}}$$

This can also be written in log-linear form, which is useful for the estimation of the model

$$\ln H_{ijt} = (1 + \varphi_{ijt}) \ln \alpha_{\mathbf{ijt}} + \varphi_{ijt} \ln P_{ijt} \tag{2}$$

where φ_{ijt} is the constant elasticity of housing supply. Assuming that φ_{ijt} is positive, this implies that ρ_{ijt} is less than 1 and that the cost function is convex. Finally, I will assume that φ_{ijt} is a function of observable tract-level characteristics, $\varphi_{ijt} = \gamma_0 + \gamma_1 x_{ijt}$. This is what will generate the heterogeneity in housing supply responses across space.

4.3 Model of Housing Demand

Housing demand is an important consideration in regions surrounded by greenbelt because greenbelts affect the choice set of households and lead to substitution. Without accounting for substitution in the model, one would not observe a broader greenbelt effect, as the reduction in development within the greenbelt would simply lead to a decline in housing supply in that area and higher prices with no spillover to the rest of the city. It is only when accounting for substitution across locations that the impact on broader regional outcomes is observed because the substitution to other locations represents a demand shock that induces developers to build more housing. The requirement for a model to incorporate substitution patterns adds some complexity to the type of housing demand model. In particular, it requires estimating a housing demand *system*, where the price of housing in one location enters into the decision making of a household looking to buy housing anywhere in the region. In this section, I present a discrete-choice model of location choice where households decide where to live in the region.

4.3.1 Housing Demand Model

A household's utility can be written as a function of the characteristics of a house of type i in location j at time t.

$$U_{ijt} = \alpha P_{ijt} + x_{ijt}\beta + \xi_{ijt} + \epsilon_{ijt}$$

where ϵ_{ijt} is distributed Type I Extreme Value (Gumbel), α is the price coefficient, x_{ijt} is a set of time-varying observable characteristics and ξ_{ijt} are the unobserved "product" characteristics, if we think of a house type-location pair as a product. The market share, s_{ijt} of houses of type i in location j at time t is

$$s_{ijt} = \frac{\exp(\delta_{ijt})}{1 + \sum_k \exp(\delta_{ikt})} \tag{3}$$

where $\delta_{ijt} = \alpha P_{ijt} + x_{ijt}\beta + \xi_{ijt}$. We can take logs and then rearrange this equation to obtain the following expression

$$\ln s_{ijt} - \ln s_0 = \alpha P_{ijt} + x_{ijt}\beta + \xi_{ijt} \tag{4}$$

Because the housing shares, s_{ijt}, are observed in the data as the share of all housing of each type across the region in each year, this equation can be estimated using regression methods. The outside option, s_0 is defined in this context as the share of housing in adjacent regions Durham and Halton that are outside of the sample used in the analysis. This choice of the outside option has an important implication for the interpretation of the counterfactual as a decrease in housing supply within the Peel, York and Toronto markets will translate into an increase in housing in more distant markets - typically in the form of more urban sprawl.

This model represents an integrated housing demand system because the location-type shares s_{ijt} depend on the prices of all other locations and unit types. In the case of the Greenbelt, if the price within the Greenbelt rises, this will decrease the share demanded within the Greenbelt and raise demand elsewhere. One limitation of this model is that households will substitute to locations in proportion to their market shares and not as a function of similarity.

4.4 Equilibrium & Dynamics

Equilibrium in this model consists of setting supply in each census tract equal to demand using the supply and demand functions. I rewrite Equation 3 to be in terms of log housing as a function of the total market size, M_t

$$Q^D_{ijt}(P) = \ln M_t + \alpha P_{ijt} + \xi_{ij} + \xi_t - \ln\left(1 + \sum_{ij} \exp\left(\alpha P_{ijt} + \xi_{ij} + \xi_t\right)\right) \tag{5}$$

I write the equation of log housing supply as

$$Q^S_{ijt}(P) = \alpha_{ij} + \gamma_0 \ln P_{ijt} + \gamma_1 x_{ijt} \ln P_{ijt} \tag{6}$$

In equilibrium, quantity supplied will equal quantity demanded, $Q^S_{ijt}(P) = Q^D_{ijt}(P)$. Bayer et al. (2004) show in a similar housing market setting with a discrete choice demand specification, that there is a unique vector of prices that clears the market for housing.

The main mechanism that drives dynamics in this model is the market size. I assume that market size, M_t, evolves exogenously over the existing period, where households either immigrate or arrive from elsewhere in the country. Canada has a very open immigration policy that always reaches the targeted and capped number of individuals, many of whom land in the Toronto area. Although it is possible these migrants respond to the intricacies of Canadian housing policy, it will be assumed this decision is exogenous. As the market size grows, this will increase the housing demand at every price point and put pressure on the supply of housing. The supply and demand curves will be fixed other than the increases in market size and changes to the ξ parameters in the housing demand function and changes to the share of developable land in the housing supply function. The model can then be solved for each period, yielding an equilibrium price and quantity in each census tract at every point in time.

5 Estimation

5.1 Housing Supply Curves

5.1.1 Empirical Specification

The first-difference of the housing supply curve derived in Equation 2, yields the following equation that can be estimated using linear regression techniques.

$$\Delta \ln H_{ijt}^{S} = \tilde{\alpha_{\mathbf{ij}}} + \varphi_{ijt} \Delta \ln P_{ijt}^{S} + \varepsilon_{ijt} \tag{7}$$

I measure the quantity of housing supplied, H_{ijt}^{S}, using the log of housing quantity of type i in census tract j at time t. The time series of housing quantities is constructed using information from the Canadian Census of Population and development flows from the Altus Group's New Residential Homes database. The date used for housing being added to a census tract is the date of the first sale of a development project as this is when a set of units can be considered to be purchased and competing with other units within the census tract. If there is no housing of that type in a particular location, that location-by-type is omitted from the analysis. The main variable of interest is the log of the housing price of housing type i in census tract j at time t. I measure housing price, P_{ijt}^{S}, using a housing price index, which reflects the prices paid for housing in a location after partialling out variation in housing characteristics from year to year.

I include a number of observable characteristics at the neighbourhood level to capture heterogeneity in housing supply elasticities. First, I include the share of undeveloped land in a census tract. The share of undeveloped land is an important determinant of census tract level housing supply elasticities because land that is undeveloped should be cheaper and easier to develop than land with existing structures on it. The greater elasticity stems from several factors including the fact that acquiring the land is less expensive and that there will be less community opposition to projects that are not in an existing community.[7] I measure the share of developable land in five-year intervals using satellite imagery from AAFC. There is exogenous variation in this variable stemming from the introduction of the Greenbelt, where I subtract the share of developable land that is blocked by the policy. This policy variation will be useful when simulating the impact of the Greenbelt policy.

The second key observable characteristic I include is whether a census tract is in an "urban growth centre" (UGC). UGCs are regional subcenters across the GTA that were designated for increased development in 2005 around the time of the Greenbelt's introduction. In Appendix Figure 9, I plot the location of the UGCs that were introduced. The UGCs correspond to centers of business activity across the region, with the largest one in downtown Toronto and were "planned to accommodate significant population and employment growth" (Places to Grow Act, 2005) in cities across Ontario. UGC designations stipulate targets for municipal planners and officials of 400 residents and jobs per hectare in Toronto by 2031 and 200 residents and jobs per hectare throughout the rest of the GTA by 2031. As a result of these targets, these locations would be expected to have a more elastic supply of condominiums and high density towers compared to nearby locations that did not receive these targets. Finally, I can include several control variables in the form of census tract characteristics in the initial period of 2001. These include the median income, the share of university graduates, the employment rate, the working age population, the initial price level, the distance to the central business district and the initial share of developable land.

5.1.2 Identification

Despite purging the tract-specific, time-invariant term and including a rich set of controls, there remain major concerns that ordinary least squares would be biased. First, there is the classic simultaneity problem in estimating supply or demand curves where only the equilibrium price and quantity are observed. Then, there is the concern that housing demand shocks can be correlated with negative supply shocks, such as the mobilization of local residents against greater development pressure in an area. As noted in Davidoff (2016), highly demanded and productive locations may also be correlated with stricter supply restrictions. If this is the case, OLS estimates of the supply elasticity would be biased towards zero. Positive demand shocks and negative supply shocks would lead to large price increases without the requisite change in quantity.

[7] There are reasons to believe however, that the housing supply elasticity on undeveloped land may not be completely elastic either. For example, land assembly is a non-trivial process (Brooks & Lutz, 2016), dynamic behaviour by land owning agents hoping for better future returns can suppress housing supply (Murphy, 2018; Capozza & Li, 1994) and less developed land may also have higher servicing costs.

To address this endogeneity requires an instrument that will shift housing demand without shifting housing supply factors like construction costs or zoning policies. The instrument I use follows the recent work of Baum-Snow & Han (2023), who estimate housing supply elasticities at the census tract level across the United States. The proposed instrument is the simulated change in residential market access (RMA), which summarizes the driving distance weighted number of employment opportunities available from a given census tract (Tsivanidis, 2023). The simulated RMA instrument is relevant because greater access to employment will increase demand for housing in a given location all else equal.

Taken alone however, RMA does not satisfy the exclusion restriction. As RMA is co-determined with a measure of firms' access to workers (FMA), which depends on the population in a given tract, RMA may be highly correlated with housing supply factors, which would render it endogenous. This can be resolved however, by isolating changes in RMA that arise from exogenous changes to labour demand in the region. Using Bartik-style shocks to labour demand in a given census tract will, in turn, shift RMA in a location in a way that is orthogonal to local population and productivity shocks. This variation in RMA, both across locations and over time, can provide the exogenous housing demand shocks required to identify the housing supply curves at the census tract level.

5.1.3 Instrument Construction

Residential market access (RMA) can be measured as the solution to a system of equations, which set firm location demand equal to worker location demand.

$$\begin{aligned} FMA_j &= \sum_i \frac{e^{-\kappa\varepsilon\tau_{ij}}\pi_i}{RMA_i} \\ RMA_i &= \sum_j \frac{e^{-\kappa\varepsilon\tau_{ij}}L_j}{FMA_j} \end{aligned} \tag{8}$$

Using data on employment, L_j, population in place of residence, π_i, and commuting distances, τ_{ij}, I can estimate the parameter cluster $\kappa\varepsilon$ and solve this system of equations for RMA and FMA in each year. First, I estimate a gravity equation of commutes on commuting distances along the road network to recover the parameter cluster $\kappa\varepsilon$.[8]. Using Pseudo-Poisson Maximum Likelihood and origin-destination fixed effects, I recover the parameter cluster $\kappa\varepsilon = -0.067$. This value is somewhat lower than those found by Tsivanidis (2023) and Baum-Snow & Han (2023) for major congested urban areas, but is well within the range of plausible estimates.

[8]For more details on how I construct this instrument, see Appendix Section B.1

Figure 7: Change in Simulated RMA in Different Years

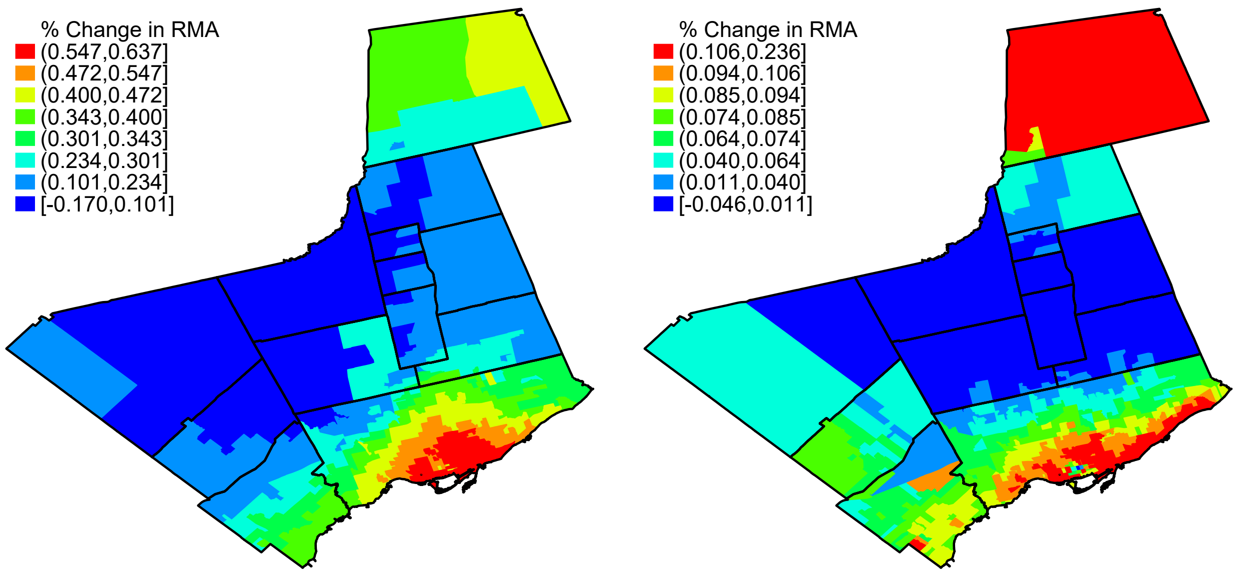

(a) 2004 (b) 2010

Note: This Figure presents the change in simulated $\ln RMA$ in two separate years, 2004 and 2010 in the Greater Toronto Area. Simulated RMA is calculated using a system of equations for RMA and FMA and data on employment by census tract, population and commuting.

Second, I replace actual observed employment and population in a census tract with a simulated Bartik-style counterpart.

$$\tilde{L}_j^t = \sum_k L_{jk}^{01} \left(\frac{\bar{L}_k^t}{\bar{L}_k^{01}} \right) \tag{9}$$

where $\bar{L}_k^t$ is the aggregate employment in industry k at time t across all of Canada excluding the Greater Toronto Area. Working population in a census tract, π_i, is inflated in each year from the 2001 base year value to match the total increase in labour demand from the Bartik-shocks. In order to alleviate further endogeneity concerns, the construction sector and census tracts within 2 km are omitted when counting employment. Plugging in these values for employment and population, solving the model and then differencing across years yields the main instrument, $\Delta \ln RMA$.

5.1.4 First-Stage Results and Unified Elasticity

The simulated RMA instrument constructed in the previous section serves as a housing demand shock that is driven solely by changes in labour demand predicted by national trends in employment. In Figure 7, I plot the change in simulated RMA in different years.

There are two interesting features to note. First, the variation in magnitude differs across years. In 2004, some census tracts saw a decline of 0.17%, while others saw an increase of 0.64%. However, in 2010, the variation is much less, ranging from increases of 0.24% to declines of 0.05%. Second, there is significant variation across space in different years. While the north-western part of the city saw reduced residential market access in 2004, it saw above average growth in RMA in 2010. This variation in housing demand over time and at a fine geographic level will help to identify the supply curves.

One concern may be that the first-stage might be weak and that there is insufficient variation in RMA to be correlated with the endogenous variable - price. In Appendix Figure 2, I present the results of the first stage regression on the set of interaction terms using different specifications. In all cases the main coefficient of the instrument is statistically significant and has the correct sign where higher RMA is correlated with higher prices.

Before estimating the complete model, I will estimate the "unified" supply elasticity that captures the average supply elasticity in the model. I present the results in Figure 2. The first two columns are estimated using no controls or fixed effects, with column 2 including simply a constant term that captures the deterministic trend over time. Column 1 sees a statistically significant estimate for the unified elasticity of 0.17, however with a constant term this becomes 0.046 and insignificant. However, with no control variables, one might be concerned about the exclusion restriction. In particular, that changes in simulated RMA may be correlated with census tract characteristics like income, which may in turn also affect the quantity of housing supplied. To account for potential violations of the exclusion restriction, I use two approaches: the addition of control variables and a census tract by housing type fixed effect.

In Columns 3 and 4, I include a suite of control variables, which are observable census tract characteristics in the initial period. Adding control variables means that identification should be thought of as arising from two census tracts with similar characteristics experiencing different housing demand shocks from simulated RMA. The one experiencing a larger shock should see a larger increase in price and also housing supply. I estimate this regression using both OLS and IV in order to see whether the instrumental variables approach generates the expected change in magnitude. Although the coefficient on price turns out to be small and statistically insignificant in both cases, the IV coefficient is positive as would be expected, while the OLS coefficient is negative. The instrument is fairly strong according to the Kleinbergen-Paap F-stat, as it surpasses the conventional thresholds for weak instruments.

In Columns 5 and 6, I address concerns about the exclusion restriction by including census tract by unit type fixed effects. This approach removes variation across different

Table 2: Results for Unified Supply Curve Regression

	No Controls		With Control Vars		CT FEs	
	IV	IV	OLS	IV	OLS	IV
$\Delta \ln P$	0.170***	0.046	-0.022	0.049	-0.023	0.096
	(0.020)	(0.077)	(0.018)	(0.085)	(0.020)	(0.093)
Constant	X	✓	X	X	X	X
Controls	X	X	✓	✓	X	X
CT FE	X	X	X	X	✓	✓
N	10899	10899	10899	10899	10899	10899
Kleibergen-Paap F	2,394.936	120.269		94.906		79.525

Standard Errors are Clustered at the CT x Unit Type level

Note: This table presents the regression coefficients for the supply curve estimation with no heterogeneity. The first two columns show the results when no controls or fixed effects are used. Columns (3) and (4) show the results when controls are used both in the case of OLS and IV and (5) and (6) do the same, but with census tract by type fixed effects rather than control variables.

census tracts and instead solely exploits variation within a census tract by unit type over time. While this addresses potential endogeneity concerns across space, it reduces a key margin of variation. The results again show two statistically insignificant values, but where the IV estimate is positive and the OLS estimate is negative. However, the instrument is weaker in this specification compared to the one using control variables.

On the surface, these regressions are somewhat confusing as one would expect housing supply elasticities to be positive and significant. However, there are reasons to believe these estimates may be accurate and precise zeros. The unified elasticity reflects the average elasticity across the entire region. In many cases, one would expect an elasticity at the census tract level to be close to zero. This is even more true in my model where I estimate housing supply separately for houses and condominium apartments. In many neighbourhoods with no room to expand it is normal for the elasticity of housing supply of single family homes to be zero.

Comparing these estimates to those in Baum-Snow & Han (2023), I find that mine are somewhat smaller, but that there are a couple of reasons for this. Baum-Snow & Han (2023) find that the total unified supply elasticity is 0.24 for total units compared to something between 0.05 and 0.17 in my case. First, splitting up single family homes and condominiums introduces a number of highly inelastic product categories. Second, Baum-Snow & Han (2023) estimate 10-year elasticities while these are one-year estimates. Given that adjustment is far more likely over ten years than one, it is unsurprising that the longer run estimates are larger. Finally, the sample used in each case may not be comparable. When looking at similarly large, supply-restricted cities in the United States, Baum-Snow & Han (2023) find average elasticities closer to 0.2. In addition, the sample used in this paper for the Greater Toronto Area omits some regions of the metropolitan area that would be included in the US estimates. Because these more distant areas likely have greater supply elasticities, this could be raising the average elasticity relative to the case in this paper. Estimating unified elasticities as the average elasticity also masks critical heterogeneity across space. This is a margin that I explore in the next section.

5.1.5 Supply Curve Main Results

I present the results of the supply curve regression with interaction terms in Table 3. Focusing mainly on columns 3-6, one can see that the uninteracted term remains small and even negative. This reflects the elasticity when all the other interaction terms are equal to zero. The uninteracted term in this case corresponds to a census tract with single family homes and no developable land such as in the middle of an urban area. When thought of this way, a small or negative value makes sense. However, for census tracts that are condominiums in urban growth centers (UGCs) or single family homes with substantial amounts of developable land, the elasticity can get much larger. For example in Column 4 (the preferred specification), a census tract with condominiums in a UGC would see an elasticity 0.605 points higher than one outside a UGC (which would be close to zero). For the share of developable land, a census tract with houses where 50% of the land is developable would see an elasticity 0.63 points higher. These results suggest substantial heterogeneity across locations in housing supply responses. One concern here is that the instruments are somewhat weak, but given that there are many instruments (one for each interaction term) and the instrument is strong in the unified case, this does not appear to be driving the results.

Another thing to note from the results table is that the coefficients using instrumental variables are much larger than the ones using OLS. The returns to a census tract being in a UGC and to a census tract having developable land are far higher in the IV regression than the OLS regression. This suggests that the OLS estimates are biased downward as would be expected if supply shocks are correlated with demand shocks. In Appendix Figure 10 I plot lowess estimates of single family housing elasticities by distance to the city center. It is clear that the IV regressions are far superior in capturing the larger elasticities at the urban fringe compared to the OLS regressions. It is also clear that elasticities only become positive beyond around 20 km from the CBD. The IV results are also fairly stable across different specifications and controls, which suggests that the instrument is isolating variation that is not significantly correlated with the controls.

Using the results from Column 4, I predict elasticities for each census tract and by type and plot them in Figure 8. As predicted, there is considerable heterogeneity across locations

Table 3: Supply Curve Regression

	No Controls		With Control Vars		CT FEs	
	IV	IV	OLS	IV	OLS	IV
$\Delta \ln P$	0.042***	-0.068	-0.082***	-0.099	-0.053***	-0.037
	(0.010)	(0.062)	(0.013)	(0.071)	(0.011)	(0.084)
$\Delta \ln P$ (Condo)	0.126**	0.088	0.046**	0.090	0.022	0.056
	(0.058)	(0.069)	(0.023)	(0.068)	(0.027)	(0.088)
$\Delta \ln P$ (Suburban House)	-0.094	-0.072	0.057	-0.073	0.089	-0.089
	(0.136)	(0.133)	(0.084)	(0.158)	(0.102)	(0.194)
Near UGC = 1 x $\Delta \ln P$	0.599***	0.632***	0.097**	0.605***	0.055	0.579***
	(0.169)	(0.163)	(0.039)	(0.165)	(0.041)	(0.223)
% Dev Land x $\Delta \ln P$ (Condo)	0.812	0.911	0.137	0.877	0.066	0.566
	(0.641)	(0.611)	(0.126)	(0.741)	(0.126)	(0.808)
% Dev Land x $\Delta \ln P$ (Suburban House)	1.281***	1.285***	0.465*	1.261***	0.225	1.144**
	(0.408)	(0.402)	(0.270)	(0.405)	(0.358)	(0.565)
Constant	X	✓	X	X	X	X
Controls	X	X	✓	✓	X	X
CT FE	X	X	X	X	✓	✓
N	10899	10899	10899	10899	10899	10899
Kleibergen-Paap F	7.166	10.752		9.306		6.360

Standard Errors are Clustered at the CT x Unit Type level

Note: This table plots the results of the supply curve regression where the dependent variable is the log of housing in a census tract of type j. A selection of interaction terms is included in this table, which capture heterogeneity in supply elasticities across census tracts.

and housing types. For single family homes housing supply elasticities are close to zero within the existing built up areas of Toronto and the largest suburb of Mississauga (in the west). However, beyond the urban fringe, the elasticities are much larger and surpass 0.8 in a number of locations. For condominiums, housing supply elasticities are largest within urban growth centers and small elsewhere, but the largest elasticities are still less elastic than single family homes. These results are consistent with the story from the conceptual framework. The Greenbelt is located in the most elastic areas of the city and push housing demand pressure to other, less elastic areas which should raise prices.

5.2 Housing Demand Curves

5.2.1 Estimation Approach

With supply estimated, attention now turns to the demand curves. The demand curve from Equation 4 can be estimated using linear regression techniques. The dependent variable is the

Figure 8: Predicted Supply Elasticities by Housing Type in the GTA, 2010

(a) Single Family Homes

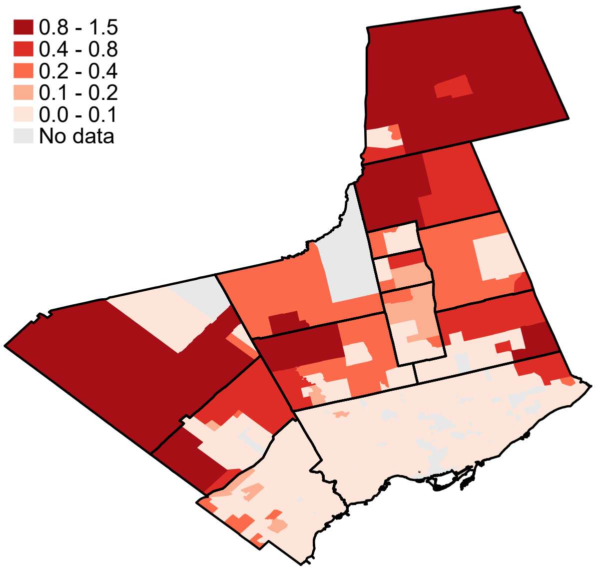

(b) Condominiums

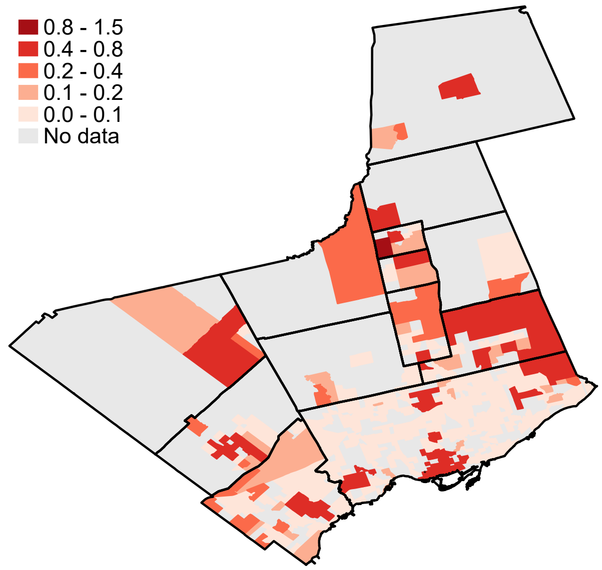

Note: These maps plot the predicted supply elasticities across census tracts in the GTA in 2010. These elasticities capture the percentage change in housing supply resulting from a 1% change in housing prices.

log share of housing of type i in location j, s_{ijt}, minus the log share of the outside option. Housing demand for a given type of housing in a given location depends on a number of characteristics, x_{ij}, including the price, p_{ijt}. One would expect that for comparable housing types and locations, a higher price would lead to less housing demanded. Estimation of the housing demand curve is a well known challenge due to simultaneity bias, where changes in equilibrium prices and quantities are driven in part by demand shocks rather than shifts along the demand curve. Therefore, an instrument that shifts supply while holding demand constant is required. Without one, standard OLS estimates would tend to be biased towards zero or even positive values. In this section, I propose a new instrument for estimating housing demand across a metropolitan area: proximity to heritage designations.

5.2.2 Heritage Designations in Ontario

As cities change and grow over time, there is often concern around losing historically significant buildings. This concern may relate to the architectural beauty of a building, its importance to the community or its connection to a famous person or organization. To address these concerns, many cities implement protections for culturally important buildings and landmarks to protect them from development in the future.

In Ontario, properties can be designated under the Ontario Heritage Act, which means

Figure 9: Heritage Designations Across the Greater Toronto Area, 2010

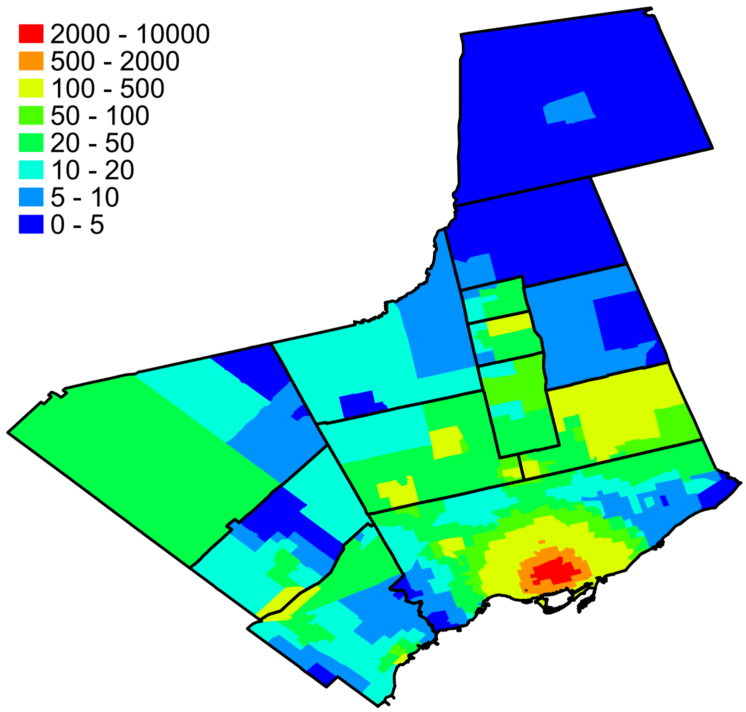

Note: This map plots the distance-discounted number of heritage designations within a 10km radius of a given census tract. A decay factor of $1/km^2$ is used. Only properties designated under the Ontario Heritage Act are included.

they are protected from renovation and demolition.[9] There are three kinds of heritage designation: listed, designated (Part IV) and designated (Part V). Listed properties are simply listed on a register where the city has 60 days to decide whether to designate a property upon receiving a building permit request for the site. Listed properties are often added to the register at the request of local heritage groups and interested residents. Designated (Part IV) properties are buildings that are recognized either for their architectural value, connection to a famous individual or importance to the area. These designations must be formally approved by a city council in order to take effect. Designated (Part V) properties are properties that fall within a Heritage Conservation District (HCD). These districts are areas with a distinct architectural style that are deemed important enough to preserve entire neighbourhoods or blocks of housing from development - there are 20 such districts in the City of Toronto alone.

I collect information on every heritage designated property (Part IV and V) in the Greater Toronto Area as well as the year of designation. In Figure 9, I show the exposure of each

[9]Similar heritage protections can be found in many other countries as well. In the UK, this began in 1953 under the Historic Buildings and Monuments Act.

census tract to nearby heritage designations. For each census tract, I calculate heritage exposure as the total number of heritage designations within 10km discounted by $\frac{1}{km^2}$. There is substantial variation in heritage designations across the region even conditional on municipality size and age. There are also significant changes in the number of designations issued each year as shown in Appendix Figure 11. During the period of interest (2000-2010), there are many heritage designations created.

5.2.3 Instrument Validity

Heritage designation exposure is a valid instrument for housing prices as it is both relevant and plausibly exogenous conditional on a set of control variables. Heritage exposure captures how much land is effectively unavailable for development nearby, which reduces supply and raises prices. As heritage exposure does not vary with unit type, I interact unit type with the instrument to allow for different impacts between the two types of unit.

Variation in this instrument comes from differences in heritage designation both across time and space. Heritage designations do not follow a uniform assignment process, but are more community driven, which can lead to significant differences across locations. It should be noted that the Toronto area has over 20,000 properties either listed or designated in heritage registers, which is a substantial amount for a city founded in only 1793. This also means that the marginal designation is not a large, important municipal landmark, but is instead a small house that someone famous may have lived in for a few years. The unremarkable nature of more recent heritage designations suggests a moderate amount of randomness in the assignment process, at least along the architectural dimension. Identification of the demand curve will then stem from variation in heritage designations across similar types of housing.

I present support for the relevancy of the instrument in the first stage regression in Appendix Table 3. Conditional on a set of numerous control variables, heritage designations have a statistically significant, positive impact on housing prices. This effect is stronger for single family homes than it is for condominiums. One explanation for this is that the supply effect of the designations is stronger for single family homes than it is for condominiums. Or in other words, a condominium requires less land in order to build a given number of units, while building a comparable number of single family homes would require more space and encounter more heritage barriers. Many of the other controls have the expected sign including education and the average age of the housing stock, which are positively correlated with price.

For the heritage instrument to be valid, I include a rich set of control variables to address concerns around the exclusion restriction. One potential violation is if heritage designations induce housing demand in a location. There are a few reasons this is unlikely to be the case. First, heritage designations do not change anything observable about a property - their purpose is in fact to preserve the building exactly as it is - which means designations should not incite changes in demand based on the change in status of a property. Second, while the age and historic nature of a property may be desirable to some buyers (Ahlfeldt & Holman, 2018), I include controls for the average age of the housing stock to address concerns that older buildings are in demand and drive further demand shocks.

A second potential violation of the exclusion restriction is that people that enact more heritage designations also attract more demand. As pointed out in Ahlfeldt et al. (2017), educated neighbourhoods are more likely to enact heritage designations, which could in turn lead to more demand in a location. To address this I include a set of initial sociodemographic characteristics including education, the employment rate and the median income of a census tract. This would mean that any violations would have to be separate from income and education. While it is possible that some have a greater preference for heritage, it is less likely that this characteristic would attract increased demand distinctly from income and education. Finally, heritage designations may lead to greater demand if the stability they offer neighbourhoods is an amenity to individuals. Another framing however, is that this "stability" is simply the negative supply shock that stability offers. Disentangling a preference for stability from simple inelastic supply is challenging and difficult to address.

5.2.4 Housing Demand Curve Results

I present the results of the demand estimation in Table 4. The first two columns present the OLS estimates with no IV, while the remaining columns show the IV results using different radii and lags. One reason to examine the lagged terms is if one thinks that it takes time for heritage designations to affect housing supply. Looking at the OLS estimates without controls, the coefficient is positive and the demand curve is upward sloping. When including controls the coefficient becomes negative, but is almost five times smaller than the IV estimates in the remaining columns. When using the IV, the results are fairly robust across different distances and to the inclusion of lagged terms. In all cases the instrument is larger than the conventional thresholds for instrument strength and the null of valid instruments in the Hansen-J test cannot be rejected.

The coefficients from the demand regression are not easily interpretable on their own, but the resulting housing demand elasticities are more useful. The average annual housing demand elasticities using the 10 km radius range from -0.66 in 2001 to -1.14 in 2010. These numbers are similar in magnitude to those in the somewhat dated literature with housing price elasticities. Albouy et al. (2016) finds an uncompensated price elasticity around -0.66

Table 4: Demand Model Regression Results

	OLS		IV - By Radius			IV - With Lags	
	10km	10km	5km	10km	15km	Lag 1-Yr	Lag 2-Yr
Prices (in $10,000)	0.0101***	-0.0066***	-0.0295***	-0.0278***	-0.0273***	-0.0286***	-0.0288***
	(0.0020)	(0.0010)	(0.0095)	(0.0083)	(0.0080)	(0.0088)	(0.0090)
Controls	X	✓	✓	✓	✓	✓	✓
Unit FE	X	✓	✓	✓	✓	✓	✓
N	12110	12110	12110	12110	12110	10899	9688
Kleibergen-Paap F			23.65	24.05	23.86	21.28	19.37
Hansen-J			.1348	.1771	.1848	.148	.1348

Standard Errors are Clustered at the CT x Unit Type level

The regression table shows the results of a number of regressions of the log share on price and other attributes. The first two columns show the OLS results with and without controls. The other columns show different radii used for the instrument and the use of lags.

and report that most estimates fall between -0.3 and -1. Zabel (2004) summarizes a variety of elasticity estimates ranging from -0.2 to -0.9. Hanushek & Quigley (1980) find elasticities of -0.64 in Pittsburgh and -0.45 in Phoenix. Baum-Snow & Han (2023) suggest using an estimate of -0.8 in their simulation exercise. Given this past research, the estimates I obtain are not unreasonable. With these estimates and the housing supply results in the previous section, I can solve for equilibrium prices and quantities and conduct counterfactual analysis.

6 Counterfactuals

6.1 The Effect of the Greenbelt on the Housing Market

Up to this point, I have established that the Ontario Greenbelt had an impact on housing development patterns and have estimated a model of housing supply and demand for the Greater Toronto Area. Determining the complete impact of the Greenbelt on the Toronto housing market will require simulating what would have occurred had the Greenbelt not been put into place and comparing the results to the baseline case where the Greenbelt was implemented. In the baseline scenario, I use the predicted housing supply elasticities from the supply curve regressions. If predicted elasticities are below zero, I adjust them to be small positive values (0.01) in order for the model to be well-behaved.[10] I fix the intercept to be such that the curve passes through the observed point of housing supply in 2004 before the policy was put in place. I compute the equilibrium prices and quantities over time using

[10] The number of census tracts where this is the case is small and when they are below zero, it is very close to zero.

Figure 10: Effect of the Greenbelt on Housing Construction, 2010

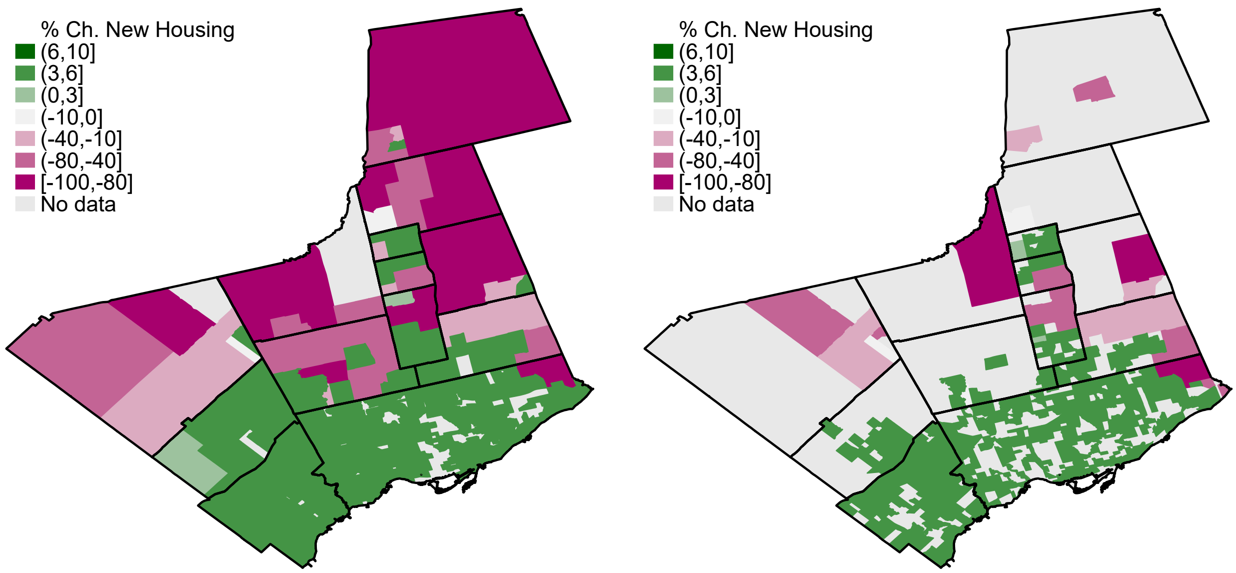

(a) Single Family Homes

(b) Condominiums

Note: This Figure presents the percentage change in housing quantity by census tract and unit type in the region. This percentage is the change in quantity in the actual, baseline Greenbelt scenario compared to the change in quantity that would have occurred had there been no Greenbelt. A value of -50% should be interpreted as a 50% reduction in the amount of construction relative to the counterfactual.

changes in aggregate population in the region.

In the counterfactual scenario where the Greenbelt was never put in place, I recompute the supply elasticity assuming that the restricted Greenbelt land was developable. To do this, I replace x_{ijt} in Greenbelt census tracts with the share of developable land in those tracts prior to the implementation of the Greenbelt. To illustrate this clearly, imagine a census tract where 40% of its land is developable, but half of this land is in the Greenbelt. In the baseline case only 20% of the land will be considered developable, but in the counterfactual scenario this will be adjusted to 40%. Increasing the share of developable land will increase the housing supply elasticity based on the coefficients from Table 3 and rotate the supply curve. I then re-estimate the equilibrium given the new supply curve.

I present the results of the effect of the Greenbelt on housing construction across space in Figure 10. In this figure housing construction is interpreted as the change in estimated equilibrium housing quantities over time in both the actual and simulated scenarios. The percentage change in housing construction as a result of the policy is the change in the amount of construction relative to the case with no Greenbelt. Concretely, a value of -60% for a census tract means that if in the Greenbelt-less scenario 500 houses would have been built, but only 200 were built in the scenario with the Greenbelt, this would be a 60%

reduction. In the figure, one can see the impact of the Greenbelt on housing construction in highly restricted areas where construction was curtailed by over 80%. However, one can also observe the positive spillovers to other parts of the region. On average in Greenbelt areas, construction declined relative to the counterfactual by 50-60%, while construction increased by around 4.5% in non-Greenbelt areas.

To compare these to the event-study results in Section 3.2, I convert these values to stocks rather than flows. I find that in Greenbelt areas, classified the same way as in the event-study, the stock of housing was around 14.5% below the counterfactual without a Greenbelt. This is closely in line with the results found in the event-study at the five-year point. This should be interpreted with some caution however, as this model may not capture the spillover mechanism in sufficient detail.

This negative supply shock to the region translates to an increase in average prices of around 2.25% by 2010. While this shock is a positive for homeowners, around half the residents of the City of Toronto are renters - for whom this would reflect an increase in housing costs. Using the price-to-rent ratio of 20 for the Toronto area, this increase translates into an increase in annual rent of around C$550 a year for single family home renters. According to the Canadian Survey of Labour and Income Dynamics, in 2010, the average pre-tax income of a renter in the Toronto area was C$61,200, which means this translates to an increase in costs worth almost 1% of pre-tax earnings - a consequential amount.

6.2 Alternative Policy Reforms

While these results serve as an interesting result, they lack some context. In order to provide greater context, in this section I will consider two alternative and more extreme policy counterfactuals and compare the impacts. In the first counterfactual, I will consider the scenario where the Greenbelt eliminated all developable farmland land and protected it as part of the Greenbelt. This would make infill development within the urban area the only way that housing could be built. In the second counterfactual, I will consider pairing a relaxation of land-use regulations within the city with the Greenbelt to see whether this could reduce the impact of the Greenbelt policy while continuing to protect environmentally sensitive areas at the urban fringe.

In the second counterfactual, I relax land use regulations through an expansion of the Urban Growth Centers (UGCs) which were discussed in the model of housing supply. I add census tracts within 1 km of an existing growth center boundary to the policy (as seen in Figure 9). This involves changing the observable characteristics for the supply curves in these regions to be considered part of the UGC. Doing this raises the elasticity of housing

Figure 11: Price Effects of the Greenbelt Under Different Counterfactuals

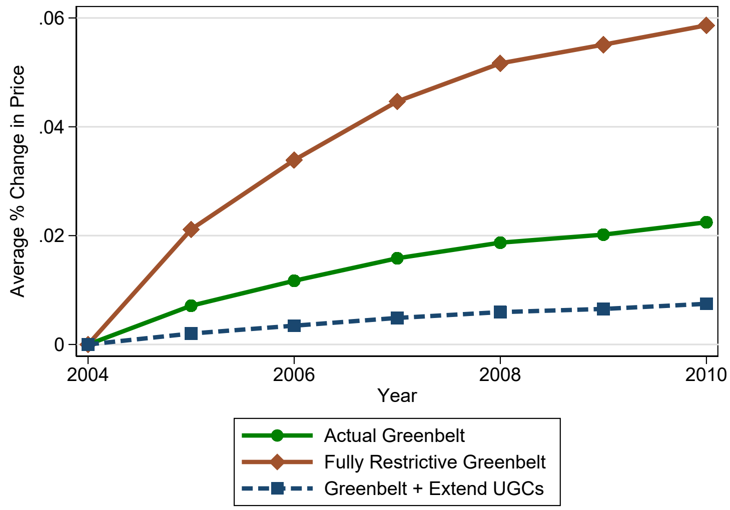

Note: This figure plots the price effects under three separate counterfactuals: the main counterfactual of the actual Greenbelt, a completely binding Greenbelt and a case where zoning regulations are relaxed along with the Greenbelt.

supply in those tracts for condominiums substantially.

After solving for the new equilibrium prices and quantities under these two alternative counterfactuals, I compare the price effects to those from the main counterfactual in Figure 11. I find that in the scenario where the Greenbelt completely restricted housing development at the urban fringe, average prices would have increased almost 6% in five-years relative to the scenario without the Greenbelt. This is almost three times the effect of the actual Greenbelt, which suggests that a more stringent Greenbelt would have had stronger impacts. Conversely, a relaxation of some zoning regulations within the built-up areas would have resulted in a much smaller price impact. Prices would have increased only 1% relative to the no-Greenbelt counterfactual had the Greenbelt also been paired with a small expansion of UGCs. These results suggest that the reduction of housing supply barriers within cities serves as a viable pathway to both environmental protection and lower housing costs.

7 Conclusion

Greenbelts are employed by cities to manage urban sprawl, but their effects on the housing market are not well understood. This paper shows that greenbelt policies can have substantial impacts on housing markets as exhibited by the introduction of the Ontario Greenbelt around the Greater Toronto Area in the early 2000s. One reason for these substantial effects is that greenbelts tend to be placed in locations with elastic housing supply responses while displacing demand to areas with inelastic housing supply responses. This leads to less housing being produced and higher costs. However, this paper also shows that land use reform that allows for greater density within cities can be an effective tool for addressing housing affordability without removing protections on environmentally-sensitive land.

This is an important result because of the myriad of benefits a greenbelt can offer, which includes open space amenities, reduced pollution from car travel, productive local agriculture and the protection of biodiversity. However, we do not have a strong understanding of the magnitude of these benefits because so many of them are diffuse in nature. In particular, while research such as Koster (2023) measures the benefits of living close to a greenbelt through a higher willingness to pay for housing, this represents only one of many possible benefits from greenbelt policies. Other benefits such as the enjoyment a family from within the city has when travelling into nature are more difficult to measure because these public spaces are usually free to consume and therefore there is no data collected. Better estimation of the benefits of greenbelts could result in stronger conclusions about the welfare effects of the policy.

Despite land use reform representing a viable pathway to achieve the many benefits from greenbelts, there are complex political economy factors which prevent these reforms from taking place. Self-interested homeowners, who are narrowly focused on the value of their property, is a classic explanation for the lack of land use reform (Fischel, 2004). However, greenbelt policies introduce an additional layer, where many strong proponents of greenbelt policies are in fact people who live within cities, while those living within greenbelt boundaries are less supportive. Those living within greenbelts may not be supportive because of the lost real option value of their property, where the value of developing the land disappears (Deaton & Vyn, 2010; Cunningham, 2007). Those within cities may support greenbelts because increased land scarcity simply further increases their home values, they value the benefits of the greenbelt directly or because cities are more liberal and more liberal cities tend to have more restrictive land use policies (Kahn, 2011). A stronger understanding of the political motivations of those in favour and against greenbelt policies could provide important lessons for how to reach a stronger policy consensus in the future.

References

Ahlfeldt, G. M. & Barr, J. (2022). The economics of skyscrapers: A synthesis. *Journal of Urban Economics*, 129, 103419.

Ahlfeldt, G. M. & Holman, N. (2018). Distinctively Different: A New Approach to Valuing Architectural Amenities. *Economic Journal*, 128(608), 1–33.

Ahlfeldt, G. M. & McMillen, D. P. (2018). Tall buildings and land values: Height and construction cost elasticities in Chicago, 1870–2010. *Review of Economics and Statistics*, 100(5), 861–875.

Ahlfeldt, G. M., Moeller, K., Waights, S., & Wendland, N. (2017). Game of Zones: The Political Economy of Conservation Areas. *Economic Journal*, 127(605), F421–F445.

Ahlfeldt, G. M. & Pietrostefani, E. (2019). The economic effects of density: A synthesis. *Journal of Urban Economics*, 111, 93–107.

Albouy, D., Ehrlich, G., Liu, Y., Davidoff, T., Deaton, A., Handbury, J., Olsen, E., Peng, C.-W., Song, C., Warren, L., & Wu, J. (2016). Housing demand, cost-of-living lnequality, and the affordability crisis.

Anagol, S., Ferreira, F., & Rexer, J. (2021). *Estimating the economic value of zoning reform*. Technical report, National Bureau of Economic Research, Cambridge, MA.

Anas, A. & Rhee, H. J. (2007). When are urban growth boundaries not second-best policies to congestion tolls? *Journal of Urban Economics*, 61(2), 263–286.

Anderson, S. T. & West, S. E. (2006). Open space, residential property values, and spatial context. *Regional Science and Urban Economics*, 36(6), 773–789.

Baum-Snow, N. & Han, L. (2023). The microgeography of housing supply. *Journal of Political Economy*.

Bayer, P., McMillan, R., & Rueben, K. (2004). An equilibrium model of sorting in an urban housing market.

Behrens, K., Duranton, G., & Robert-Nicoud, F. (2014). Productive cities: Sorting, selection, and agglomeration. *Journal of Political Economy*, 122(3), 507–553.

Bento, A. M., Franco, S. F., & Kaffine, D. (2006). The efficiency and distributional impacts of alternative anti-sprawl policies. *Journal of Urban Economics*, 59(1), 121–141.

Bradburn, J. (2022). How Ontario got its Greenbelt — and who tried to stop it.

Brooks, L. & Lutz, B. (2016). From today's city to tomorrow's city: An empirical investigation of urban land assembly. *American Economic Journal: Economic Policy*, 8(3), 69–105.

Brueckner, J. K. (2001). Urban sprawl: Lessons from urban economics. *Brookings-Wharton Papers on Urban Affairs*, 2001(1), 65–97.

Brueckner, J. K. (2007). Urban growth boundaries: An effective second-best remedy for unpriced traffic congestion? *Journal of Housing Economics*, 16(3-4), 263–273.

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

Capozza, D. & Li, Y. (1994). The intensity and timing of investment: The case of land. *American Economic Review*, 84(4), 889–904.

Combes, P. P., Duranton, G., & Gobillon, L. (2019). The costs of agglomeration: House and land prices in French cities. *Review of Economic Studies*, 86(4), 1556–1589.

Cunningham, C. R. (2007). Growth controls, real options, and land development. *Review of Economics and Statistics*, 89(2), 343–358.

Davidoff, T. (2016). Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors. *Critical Finance Review*, 5(2).

Deaton, B. J. & Vyn, R. J. (2010). The effect of strict agricultural zoning on agricultural land values: The case of Ontario's Greenbelt. *American Journal of Agricultural Economics*, 92(4), 941–955.

Dehring, C. A., Depken, C. A., & Ward, M. R. (2008). A direct test of the homevoter hypothesis. *Journal of Urban Economics*, 64(1), 155–170.

Duranton, G. & Puga, D. (2023). Urban growth and its aggregate implications. *Econometrica (Forthcoming)*.

Fang, L., Stewart, N., & Tyndall, J. (2023). Homeowner politics and housing supply. *Journal of Urban Economics*, 138, 103608.

Fischel, W. A. (2004). *The homevoter hypothesis: How home values influence local government taxation, school finance, and land-use policies*. Harvard University Press.

Gibson, M. & Carnovale, M. (2015). The effects of road pricing on driver behavior and air pollution. *Journal of Urban Economics*, 89, 62–73.

Glaeser, E. & Gyourko, J. (2018). The economic implications of housing supply. *Journal of Economic Perspectives*, 32(1), 3–30.

Glaeser, E. L., Gyourko, J., & Saks, R. E. (2006). Urban growth and housing supply. *Journal of Economic Geography*, 6(1), 71–89.

Glaeser, E. L. & Ward, B. A. (2009). The causes and consequences of land use regulation: Evidence from Greater Boston. *Journal of Urban Economics*, 65(3), 265–278.

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

Greenbelt Act (2005). SO 2005, c 1.

Gyourko, J. & Krimmel, J. (2021). The impact of local residential land use restrictions on land values across and within single family housing markets. *Journal of Urban Economics*, 126, 103374.

Hanushek, E. A. & Quigley, J. M. (1980). What is the price elasticity of Housing Demand? *Review of Economics and Statistics*, 62(3), 449.

Helsley, R. W. & Strange, W. C. (2008). A game-theoretic analysis of skyscrapers. *Journal of Urban Economics*, 64(1), 49–64.

Hsieh, C. T. & Moretti, E. (2019). Housing constraints and spatial misallocation. *American Economic Journal: Macroeconomics*, 11(2), 1–39.

Ihlanfeldt, K. R. (2007). The effect of land use regulation on housing and land prices. *Journal of Urban Economics*, 61(3), 420–435.

Kahn, M. E. (2000). The environmental impact of suburbanization. *Journal of Policy Analysis and Management*, 19(4), 569–586.

Kahn, M. E. (2011). Do liberal cities limit new housing development? Evidence from California. *Journal of Urban Economics*, 69(2), 223–228.

Koster, H. R. A. (2023). The welfare effects of greenbelt policy: Evidence from England. *The Economic Journal*.

Kulka, A., Sood, A., & Chiumenti, N. (2023). Under the (neighbor)hood: Understanding interactions among zoning regulations. *SSRN Electronic Journal*.

Mayer, C. J. & Somerville, C. T. (2000). Land use regulation and new construction. *Regional Science and Urban Economics*, 30(6), 639–662.

Murphy, A. (2018). A dynamic model of housing supply. *American Economic Journal: Economic Policy*, 10(4), 243–267.

Places to Grow Act (2005). SO 2005, c 13.

Quigley, J. M. & Swoboda, A. M. (2007). The urban impacts of the Endangered Species Act: A general equilibrium analysis. *Journal of Urban Economics*, 61(2), 299–318.

Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics*, 125(3), 1253–1296.

Statistics Canada (2016). *Special tabulation of data from the 1971, 1991, 2001 and 2011 Census of Population*. Technical report.

Sun, L. & Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2), 175–199.

Tsivanidis, N. (2023). Evaluating the Impact of Urban Transit Infrastructure: Evidence from Bogotá's TransMilenio.

Turner, M. A., Haughwout, A., & Klaauw, W. v. d. (2014). Land use regulation and welfare. *Econometrica*, 82(4), 1341–1403.

Walsh, R. (2007). Endogenous open space amenities in a locational equilibrium. *Journal of Urban Economics*, 61(2), 319–344.

Winfield, M. (2012). *Blue-green province: The environment and the political economy of Ontario*. Vancouver: UBC Press.

Zabel, J. E. (2004). The demand for housing services. *Journal of Housing Economics*, 13(1), 16–35.

Figure 1: Census Tracts Classified into Treatment and Control Groups

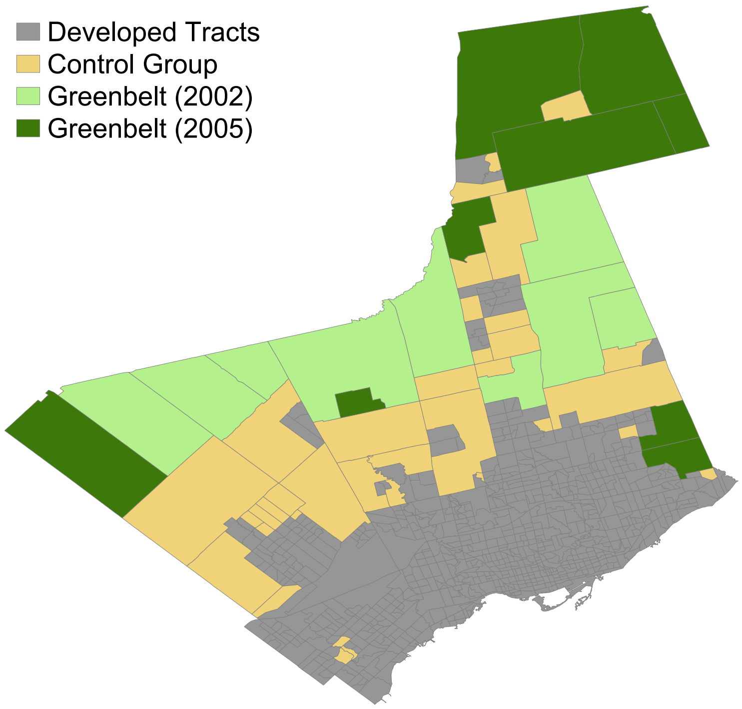

Note: This map plots the treatment status of census tracts across the GTA according to whether it is developed, in the control group or in a particular treatment group. Treatment status is defined as having over 50% of a census tract within the Greenbelt.

A Figures

Figure 2: Average Greenbelt Exposure by Share of Undeveloped Land %

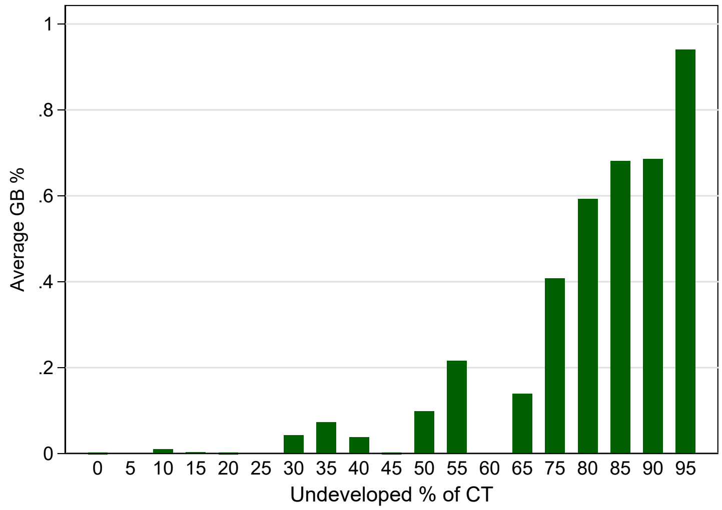

Note: This figure plots the average greenbelt exposure at the census tract level by the share of undeveloped land. Less developed census tracts are more likely to be exposed to the Greenbelt.

Figure 3: Average Development 2000-2010 by Share of Undeveloped Land %

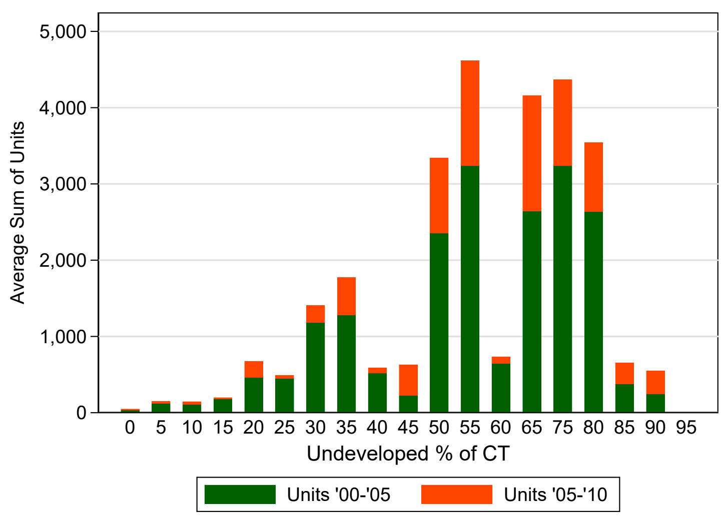

Note: This figure plots the average amount of housing development in the periods before and after the full Greenbelt was implemented. Census tracts that have less than 20% of land undeveloped see virtually no housing construction on average relative to those with more undeveloped land.

Figure 4: Callaway & Sant'Anna (2021) Estimator

(a) Treated in 2001 (b) Treated in 2005

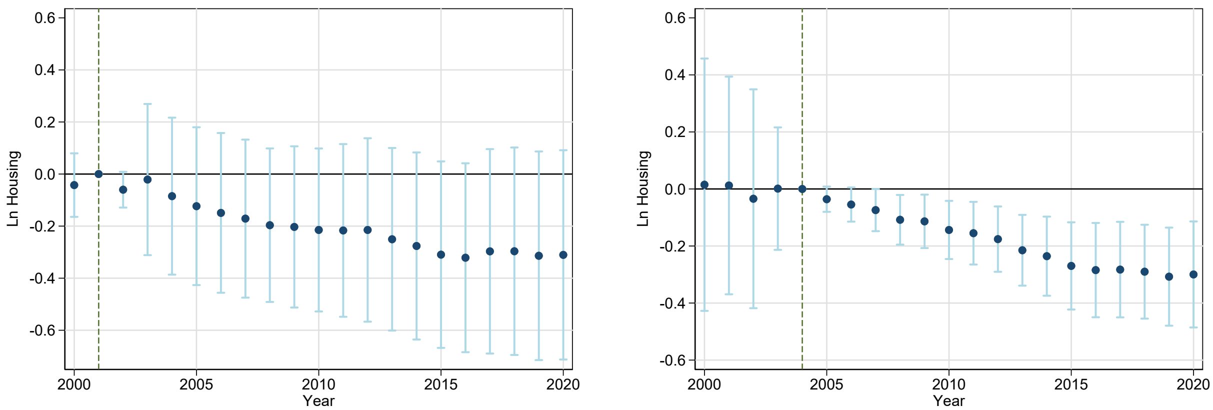

Note: This figure plots the estimates using the Callaway & Sant'Anna (2021) estimator for staggered difference-in-differences. The left panel shows the effects for the group treated in 2001, while the right panel shows the effects for the group treated in 2005. Standard errors are clustered at the census tract level.

Table 1: Continuous Treatment Regression Results by % of Undeveloped Land Threshold

	Undeveloped Land Share				
	20%	25%	30%	35%	40%
Continuous Treatment	-0.16404	-0.17740	-0.27921**	-0.34771**	-0.39829***
	(0.11591)	(0.11655)	(0.12362)	(0.13241)	(0.13899)
N	1617	1365	1197	1071	987
R^2	0.938	0.940	0.938	0.934	0.931

Standard errors in parentheses

Standard Errors Clustered at the Census Tract Level

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: This table presents the regression coefficient from a continuous treatment regression of the log of housing in a census tract on the share of Greenbelt land. Each column refers to the cutoff of undeveloped land for census tracts included in the sample. Standard errors are clustered at the census tract level.

Figure 5: Event Study Results by Greenbelt Threshold

(a) GB Cutoff: 45% (b) GB Cutoff: 50% (main) (c) GB Cutoff: 55%

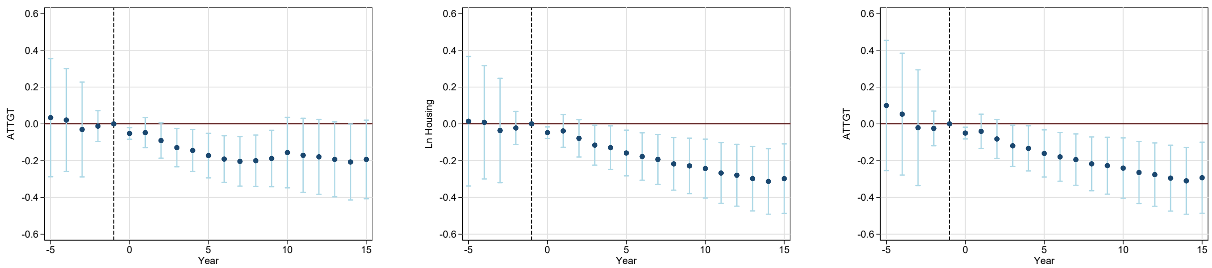

Note: This figure plots the results of the event study regression specification using different thresholds for being treated by the Greenbelt policy. The dependent variable is the log of housing while a census tract is considered treated if there is more than the specified threshold covered by the Greenbelt. The main result is presented in the center for context.

Figure 6: Event Study Results by Developable Land % Threshold

(a) Dev Cutoff: 20% (b) Dev Cutoff: 25% (main) (c) Dev Cutoff: 30%

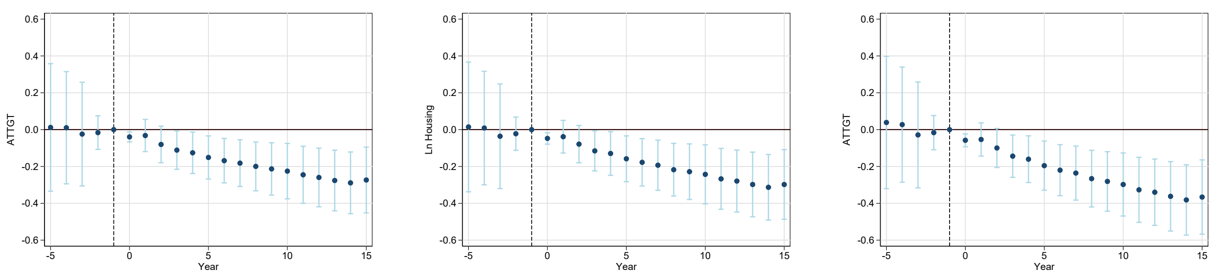

Note: This figure plots the results of the event study regression specification using different thresholds for the amount of developable land available to be included in the control group. The dependent variable is the log of housing while a census tract is considered treated if it is more than 50% covered by the Greenbelt. The main result is presented in the center for context.

Figure 7: Event Study Results Omitting 10-50% Greenbelt Tracts From Control

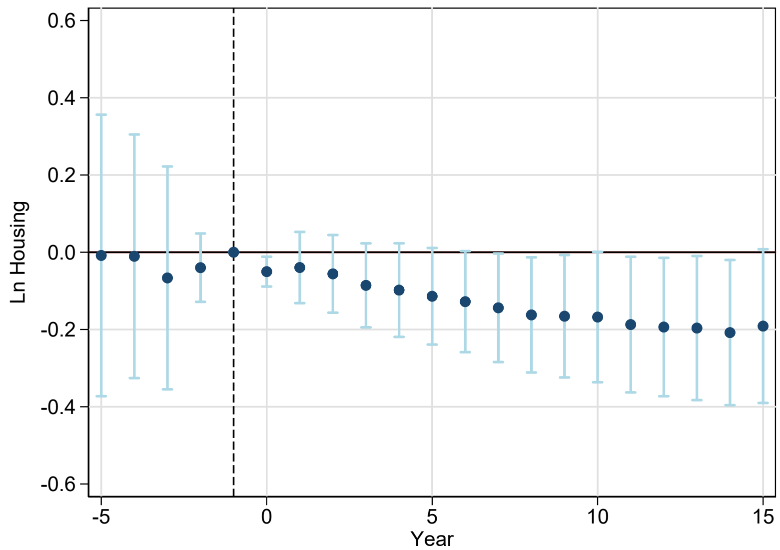

Note: This figure plots the results of the event study regression specification without census tracts treated 10-50% by the Greenbelt included in the control. The dependent variable is the log of housing and a census tract is considered treated if it is more than 50% covered by the Greenbelt. Standard errors are clustered at the census tract level.

Figure 8: Event Study Results Comparing 10-50% Covered to Those <10% Covered

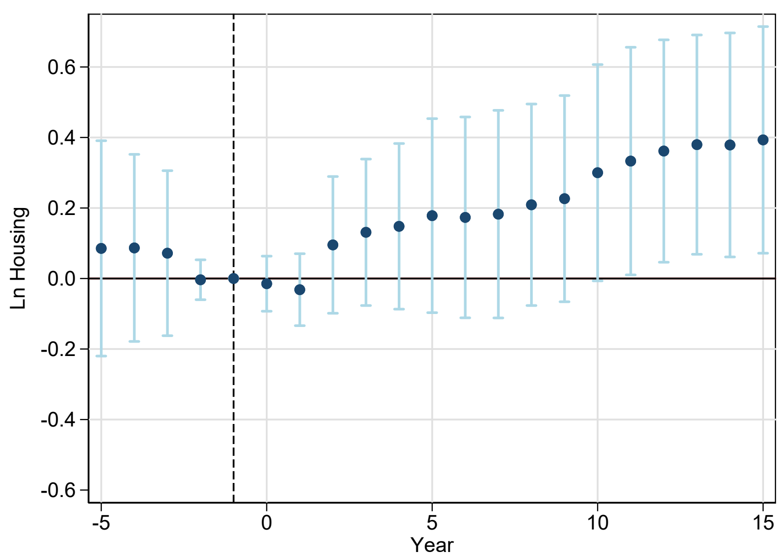

Note: This figure plots the results of the event study regression specification between census tracts treated 10-50% by the Greenbelt and those less than 10% covered. The dependent variable is the log of housing and a census tract is considered treated here if it is between 10-50% covered by the Greenbelt. Standard errors are clustered at the census tract level.

Figure 9: Urban Growth Centers in the Greater Toronto Area

(a) Actual UGCs (b) Proposed UGC Expansion

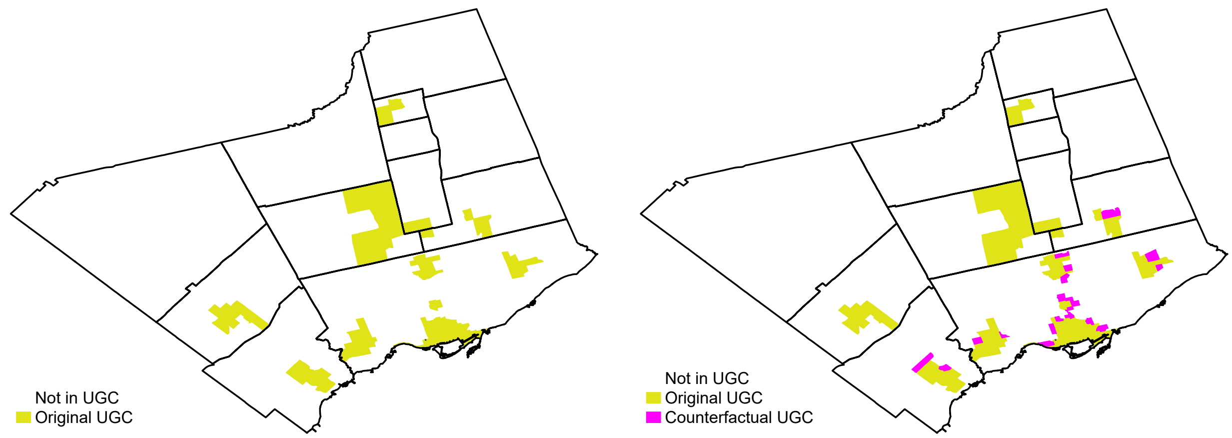

Note: This figure plots the location of "urban growth centers" in the GTA region. In yellow are the actual UGCs that were introduced in 2005 with density targets for the designated areas. In pink are a proposed set of 19 UGCs that fall within 1 km of the existing UGC boundaries to be used in the counterfactual analysis.

Table 2: First Stage Results of Supply Curve Regression

	No Controls		With Control Vars	CT FEs
	(1)	(2)	(3)	(4)
$\Delta \ln RMA$	4.472***	1.408***	1.287***	1.102***
	(0.084)	(0.089)	(0.093)	(0.095)
$\Delta \ln RMA$ (Condo)	-1.218***	-1.156***	-1.028***	-0.773***
	(0.205)	(0.199)	(0.194)	(0.207)
$\Delta \ln RMA$ (Suburban House)	-1.086*	0.132	0.242	0.475
	(0.638)	(0.658)	(0.684)	(0.742)
Near UGC = 1 x $\Delta \ln RMA$	1.174**	1.024**	0.676	0.762
	(0.461)	(0.457)	(0.478)	(0.637)
% Dev Land x $\Delta \ln RMA$ (Condo)	-0.070	2.241	2.241	2.620
	(2.212)	(2.042)	(2.127)	(2.428)
% Dev Land x $\Delta \ln RMA$ (Suburban House)	1.278	0.679	0.499	0.600
	(1.580)	(1.495)	(1.513)	(1.639)
Constant	X	✓	X	X
Controls	X	X	✓	X
CT FE	X	X	X	✓
N	10899	10899	10899	10899
R^2	0.119	0.011	0.276	0.042

Standard Errors are Clustered at the CT x Unit Type level

Note: This table presents the regression coefficients from the first stage regression of prices on simulated ln RMA. The first two columns have no controls with one omitting the constant term. I subsequently control for a set of variables including median income, employment rate, university degree share and others. The final column is estimated without controls, but with census tract fixed effects.

Figure 10: Comparing Predicted Elasticities by Distance to CBD and Regression Type

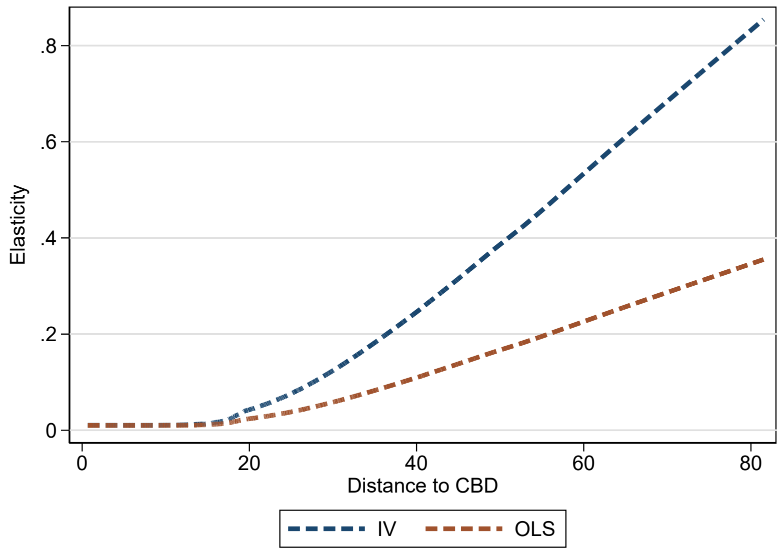

Note: This figure plots lowess regression of elasticity on distance to the central business district (CBD) for both the OLS and IV elasticity estimates. Distance is calculated in kilometers and elasticities are estimated using coefficients in Column 4 of Table 3.

Figure 11: Heritage Designations in the Greater Toronto Area Over Time

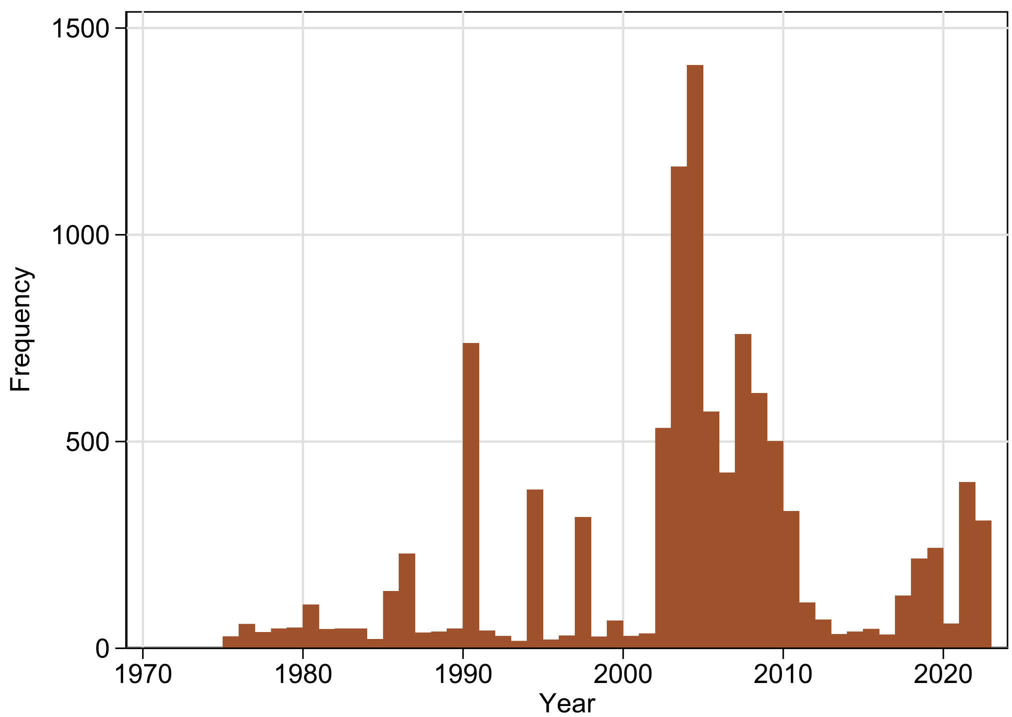

Note: This map plots number of heritage designations issued each year in the Greater Toronto Area.

Table 3: First Stage and Reduced Form Results of Demand Regression

	Price (in $10,000$)	log Share
Unit Type=0 × Heritage Exposure	0.0023***	-0.0001**
	(0.0006)	(0.0001)
Unit Type=1 × Heritage Exposure	0.0068***	-0.0002***
	(0.0011)	(0.0000)
Initial Number of Units	0.0004**	0.0006***
	(0.0002)	(0.0000)
Initial Prices	0.0001***	-0.0000
	(0.0000)	(0.0000)
Initial Share of Unit Type	-2.4314***	1.5992***
	(0.5200)	(0.1163)
Median Income	0.0000**	-0.0000***
	(0.0000)	(0.0000)
Share with University Degree	8.4667***	0.9962***
	(1.9400)	(0.2972)
Employment Rate	-0.0103	-0.0052
	(0.0228)	(0.0034)
Average Age of Housing Stock	0.0630***	0.0027
	(0.0124)	(0.0017)
Average Square Footage	-0.0000	-0.0000
	(0.0000)	(0.0000)
Average Lot Size	0.0000	-0.0000**
	(0.0000)	(0.0000)
Distance to CBD	0.0392	-0.0185***
	(0.0471)	(0.0060)
Distance to CBD × Distance to CBD	-0.0007	0.0002***
	(0.0006)	(0.0001)
Constant	-0.3090	-6.7026***
	(1.7232)	(0.2103)
Unit FE	✓	✓
Observations	12110	12110
R^2	0.753	0.706

Standard Errors are Clustered at the CT x Unit Type level

Note: This table presents the regression coefficients from the first stage regression of prices on heritage exposure and a set of control variables. The second column presents a regression of the dependent variable the log share minus the log of the share of the outside option on the instrument and a set of control variables.

B Additional Details

B.1 Instrument Construction

The residential market access instrument used in the papers closely follows the one used in Baum-Snow & Han (2023). I will provide additional details here that I do not discuss in the main sectino of the paper. In a quantitative spatial model, residential market access reflects the access to employment opportunities as summarized by the following equation, $\sum_k \sum_{j'} [w_{j'k} e^{-\kappa \tau_{ij'}}]^{\varepsilon}$. Here, w reflects the wage earned in commuting location j and this is discounted by the time it takes to get there. The model generates the following equilibrium property:

$$
\begin{aligned}
FMA_j &= \sum_i \frac{e^{-\kappa\varepsilon\tau_{ij}} \pi_i}{RMA_i} \\
RMA_i &= \sum_j \frac{e^{-\kappa\varepsilon\tau_{ij}} L_j}{FMA_j}
\end{aligned}
\tag{10}
$$

Using data on employment, L_j, population in place of residence, π_i, and commuting distances, τ_{ij}, I can estimate the parameter cluster $\kappa\varepsilon$ and solve this system of equations for RMA and FMA in each year.

I estimate the parameter cluster $\kappa\varepsilon$ using a gravity equation and information on commutes and driving distances using the contemporary road network in the Greater Toronto Area. Although this does not reflect the exact commuting pathways of the time, it gives a useful approximation for driving distances. Commutes themselves are taken from the 2001 Canadian Census of Population. The gravity regression uses origin-destination fixed effects and a Pseudo-Poisson Maximum Likelihood approach (due to the many zeros in the commuting matrix) and yields an estimate for the parameter cluster of $\kappa\varepsilon = -0.067$. This is somewhat higher than some estimates in Tsivanidis (2023) and Baum-Snow & Han (2023) for large congested cities, but well within the range of plausible estimates.

Rather than using the exact values for employment and population, Baum-Snow & Han (2023) use simulated values derived from Bartik-shocks. Using employment and population information from the Canadian Census of Population, I can then plug these simulated values into the following equation. Industries are based on 2 digit NAICS codes and there are 15 of these in the sample. Aggregate trends in employment are at the national level minus the Greater Toronto Area. To alleviate further endogeneity concerns, they omit census tracts within 2 km of a given census tract and omit the construction industry from the Bartik-shock

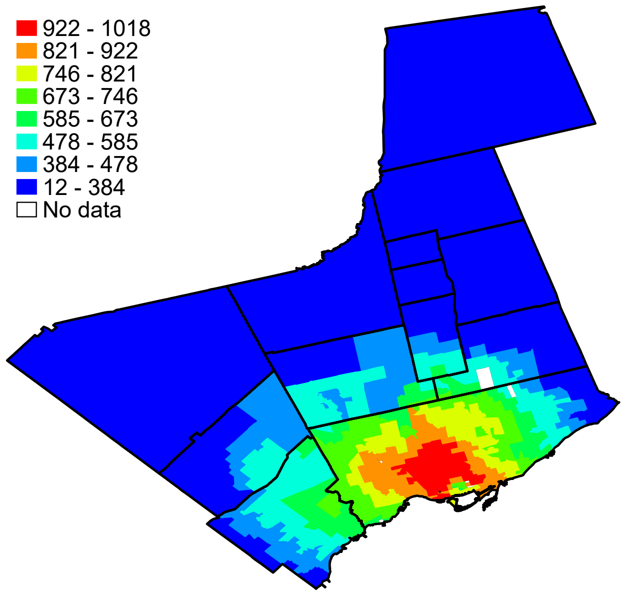

Figure 1: Ln Simulated RMA, 2004

- I do the same in my version.

$$
\begin{aligned}
\widetilde{RMA}_i^t &= \sum_{j \subseteq R(i)} \frac{e^{-(\hat{\kappa})_{r(i)}\tau_{ij}^{01}} 1\left(dis_{ij} > 2km\right) \sum_{k \neq cons} L_{jk}^{01} \left[E^t_{r'(j)k} / E^{01}_{r'(j)k}\right]}{\widetilde{FMA}_j^t} \\
\widetilde{FMA}_j^t &= \sum_{i \subseteq R(j)} \frac{e^{-(\widehat{\varepsilon\kappa})_{r(i)}\tau_{ij}^{01}} 1\left(dis_{ij} > 2km\right) \pi_i^{01} \left[\frac{\sum_{j \subseteq R(i) k \neq cons} L_{jk}^{01} \left(E^t_{r'(j)k} / E^{01}_{r'(j)k}\right)}{\sum_{j \subseteq R(i)} L_j^{01}}\right]}{RMA_i^t}
\end{aligned} \tag{11}
$$

I present the estimated simulated ln RMA is Figure 1. The best access to employment opportunities is to the center of the city, but slightly to the west of center, which reflects that the western suburbs have a higher number of employment opportunities. This can be repeated for every year in the sample and changes can be computed over time.

Differencing over time is important given that areas in the center of the region tend to have better access to employment systematically. In Figure 2, I plot the total change in simulated RMA over time. The range of changes goes from slightly negative in the north-west of the region to almost an 8% increase around the city center. This reflects the growth of finance and technology jobs during this period which were more concentrated in the urban center and the decline in agricultural and manufacturing jobs that took place to the north-west.

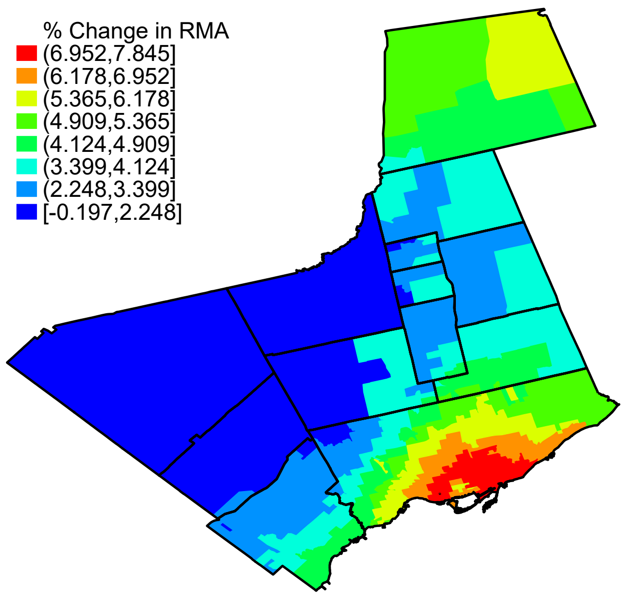

Figure 2: Change Ln Simulated RMA, 2002-2010