Homeowner Politics and Housing Supply*

Limin Fang† Nathan Stewart‡ Justin Tyndall§

October 18, 2023

Abstract

This paper examines whether homeowner opposition to nearby housing development affects local councillors' votes on housing bills. Homeowners benefit financially from restricted housing supply through increased housing prices. City councillors, who approve housing development applications, cater to the needs of homeowners who are often long-term resident voters with a financial stake in neighbourhood amenity levels. Using data from Toronto, Canada from 2009 to 2020, we identify housing bills through a machine learning algorithm. We find that councillors who represent more homeowners oppose more housing bills. In particular, councillors are significantly more likely to oppose large housing developments if the project is within their own ward.

Keywords: Housing supply; Urban development; Land-use regulation; NIMBY-ism

JEL classification: R31, R38, R52

---

*This paper is supported in part by funding from the Social Sciences and Humanities Research Council. We thank the participants of the 2023 North American Meeting of the Urban Economic Association and the 2022 American Real Estate and Urban Economics Association conference, especially Xiaolun Yu and Rhiannon Jerch, for valuable comments. We also thank the participants of the 2021 Canadian Economics Association conference and the University of Hawai'i applied microeconomics seminar series for helpful advice. Thank you to Guenther Lomas of Sigtica for assistance with AI methods. Valuable research assistance was provided by Aurora Marstokk and David Tomalin. Thank you to Michael Fox of MOD Developments for discussing details of Toronto's development process. All errors are our own.

†Sauder School of Business, University of British Columbia, 2053 Main Mall, Vancouver, BC, Canada. *limin.fang@sauder.ubc.ca*.

‡Department of Geography and Planning, University of Toronto, 100 St. George Street, Toronto, ON, Canada. *nathan.stewart@mail.utoronto.ca*.

§University of Hawai'i Economic Research Organization & Department of Economics, University of Hawai'i at Mānoa, 2424 Maile Way, Saunders Hall 540, Honolulu, HI, United States. *jtyndall@hawaii.edu*.

1 Introduction

Housing affordability has become an acute policy concern in both Canada and the United States. Acceleration in home prices has not led to a significant expansion of housing supply. Supply constraints have played a big role (Burn-Murdoch, 2023). These constraints include policies and regulations, such as zoning restrictions, stringent building codes, and long processing times for permitting. Political opposition to new housing from incumbent residents, termed NIMBYism (Not in My Back Yard-ism), has supported restrictive policies and acted to directly impede new housing supply (Glaeser et al., 2005b; Kahn, 2011). In this paper, we empirically demonstrate the role of NIMBYism in housing supply by examining the connection between homeowners, their local political representatives, and the political approval or denial of housing.

Although both homeowners and renters may object to new housing in their neighbourhoods due to potential congestion externalities (Davidoff et al., 2022), incumbent homeowners are particularly motivated to restrict housing supply because they benefit financially from an increase in housing prices. City councillors, who approve housing development applications, are more likely to cater to the needs of homeowners because homeowners are often long-term resident voters, whereas renters are not. Renters may support new housing because additional housing could reduce rent and bring positive urban amenities (Diamond and McQuade, 2019).

In this paper, we ask whether homeowners' opposition to new housing materializes through local politics; in particular, through councillors' votes on housing projects. To answer this question, we build a novel data set containing councillors' voting history on all bills in Toronto, Canada from 2009 to 2020. We gather information from the full text of each bill. Knowing if a particular bill concerns housing requires reading the full bill's text, as well as accompanying documents. We use cutting-edge machine learning algorithms to read through the large volume of text in order to identify housing bills. We connect this data with local demographic information. Using these data, we estimate the relationship between city councillors' voting behaviours and the shares of homeowners in the wards they represent. We find that councillors from wards with greater shares of homeowners are more likely to oppose new housing developments. In particular, councillors are especially resistant to large housing projects within their own wards. Compared with other ward characteristics, homeownership is a strong predictor of councillors voting against housing development. These findings provide direct evidence that NIMBYism is a significant impeding force in housing supply, and that city councillors are influenced by NIMBY sentiments.

Our paper's main contribution is to provide new, direct evidence confirming the role of resident opposition to housing that has often been assumed in the previous literature. Several papers have linked homeownership to housing supply restrictions; examples include Ahlfeldt (2011), Ahlfeldt and Maennig (2015), Cheung and Meltzer (2013), Dehring et al. (2008), Fischel (2001), Fischel (2005), and Hilber and Robert-Nicoud (2013). There is also evidence that anti-growth lobbying movements comprised of both renters and owners have succeeded in blocking housing production (Glaeser et al., 2005b; Kahn, 2011). However, past research does not explicitly study the mechanisms linking homeowner opposition to the eventual blocking of new housing construction. Our modeling approach allows us to trace the mechanisms that start from local opposition to housing and end with city councillors casting votes to block new housing. By measuring councillor voting behaviour and constituent demographics directly, we can estimate the extent to which incumbent residents block new housing through their political representatives.

Understanding the extent to which incumbent homeowners block new housing has important policy implications. In particular, limits on housing supply impede labour mobility, increase housing prices and significantly reduce social welfare (Glaeser et al., 2005a; Glaeser and Ward, 2009; Gyourko and Krimmel, 2021; Gyourko et al., 2021; Gyourko and Molloy, 2015; Hsieh et al., 2015; Hsieh and Moretti, 2019; Ihlanfeldt, 2007; Turner et al., 2014). A better understanding of the most important housing constraints allows policymakers to design more targeted policies that effectively address barriers to housing production.

1.1 Literature Review

Prior literature is built on an assumption that homeowners are able to impose limitations on housing construction by exerting political power. Ortalo-Magné and Prat (2014) provide a formal theoretical model of politically generated housing restrictions, and many papers argue that local politics can constrain housing supply (Feinerman et al., 2004). Our study expands on the existing literature by empirically establishing a mechanism through which homeowners exert political influence.

There is an established correlation in the literature between politics and the permitting of new housing. For example, Kahn (2011) showed that more liberal cities in California permit less housing. In Spain, Solé-Ollé and Viladecans-Marsal (2013) found that left-wing governments tended to release less land for development and Solé-Ollé and Viladecans-Marsal (2012) found that more competitive electoral environments in Spanish cities led to less land being released, particularly in areas with more homeowners.

The ability of homeowners to exercise influence may be a function of the political incentives of the local electoral system. Mast (2020) contrasted cities that elect councillors "at-large" with those that use discrete electoral districts, finding that switching from an at-large to a district-based system reduced new housing supplied by 20%. One explanation is that local representatives have a strong incentive to oppose housing in their own district; an incentive that disappears with the dissolution of electoral wards. Hankinson and Magazinnik (2022) examined the impacts of a state law in California that prompted many municipalities to change from at-large to district-based councillor elections. While cities that made the switch permitted less housing, the reforms also changed the spatial distribution of housing permits, making them less likely to be disproportionately filed in marginalized districts.

Our study is closely related to Lee et al. (2004), who showed that the voting behavior of Representatives in the US House remains static in the face of changing constituent preferences. In our paper, we also estimate how a changing constituency affects politician voting behaviour. Different from them, we find that local representatives do adapt their voting to suit the changing needs of their constituencies.

Our study also contributes to literature on the Homevoter Hypothesis, which finds that homeowner voting is driven by housing price incentives (Fischel, 2002, 2005). Unlike renters, Homevoters benefit financially from improvements in location quality, making them more likely to vote in favour of positive amenities and more likely to resist local disamenities. Ahlfeldt and Maennig (2015) and Dehring et al. (2008) examine public polls on large infrastructure projects; their results conform to the homevoter hypothesis.

Some prior work has argued that homeowner influence in the US has grown over time (Glaeser et al., 2005b). Clarke and Freedman (2019) demonstrated a connection between the rise of Home Owner Associations (HOAs) and strict local land use regulation. Several past studies conclude that NIMBYism is often championed by cadres of vocal homeowners who do not represent larger area populations (Hankinson, 2018; Einstein et al., 2019a,b).

In the context of the literature, our contribution is to establish a direct link between local homeownership and the denial of new housing permits. The endogeneity in households' homeownership choice complicates the causal interpretation of our analysis. People choose owning over renting for many reasons. For example, many people prefer owning for its long-term locational stability, whereas renting introduces a threat of eviction. High-income households are more likely to own, although other factors can prompt high-income households to rent (Aizcorbe et al., 2003; Davidoff, 2006; Henderson and Ioannides, 1983; Sinai and Souleles, 2005, 2013). We discuss these factors in detail in our conceptual framework (Section 4.2), and control for factors such as income in the empirical analysis.

The paper proceeds as follows. Section 2 introduces the empirical setting. Section 3 describes the data used. Section 4 provides the estimation methodology. Section 5 discusses results and Section 6 concludes.

2 Empirical Setting

The City of Toronto is Canada's largest municipality (4th in North America), containing 2.8 million residents. Like many North American cities, Toronto has undergone a period of sustained home price appreciation over recent years (Figure 1). In 2005, the benchmark price[1] of a single family home in the Toronto metropolitan area was $347,000 (CAD); in 2021, it was $1,174,000 (+238%). Over the same period, condominium prices increased by 199%.

Despite rapid price increases, the overall flow of new housing was essentially flat. In 2005, there were 42,000 housing starts in the Toronto metropolitan area (Figure 2). In 2020 there were 39,000, with a significant shift away from single family housing towards condominiums.[2]

The recent price and quantity trends suggest that the supply elasticity of new housing in Toronto is extremely low. The City Council has significant power to dictate the flow of new housing. Projects that require a change to the existing municipal zoning code (often called a "rezoning") require a council vote for approval. In Toronto, nearly all midrise or highrise buildings need to go through rezoning. Smaller projects requiring only a minor variance from zoning, such as single family homes, typically go through

[1]The "Benchmark Home Price" is calculated by the Canadian Real Estate Association (CREA) to represent the transaction price of a "benchmark home," whose attributes are typical of homes traded in the area where it is located (CREA, 2022). The calculation aims to control for the changing attributes of the homes sold over time.

[2]Between 2005 and 2021, annual condominium starts increased by 76% while single family home starts fell by 63%.

Figure 1: Benchmark Home Price in Greater Toronto

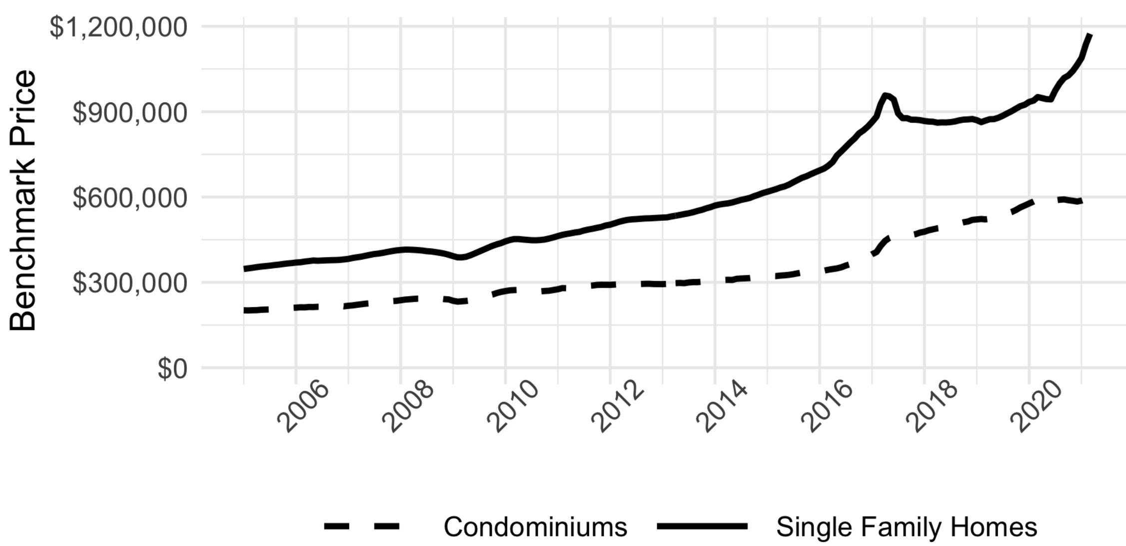

Home values in Toronto have climbed dramatically between 2005 and 2021. Source: CREA MLS Seasonally Adjusted Housing Price Index.

Figure 2: Housing Starts in Greater Toronto

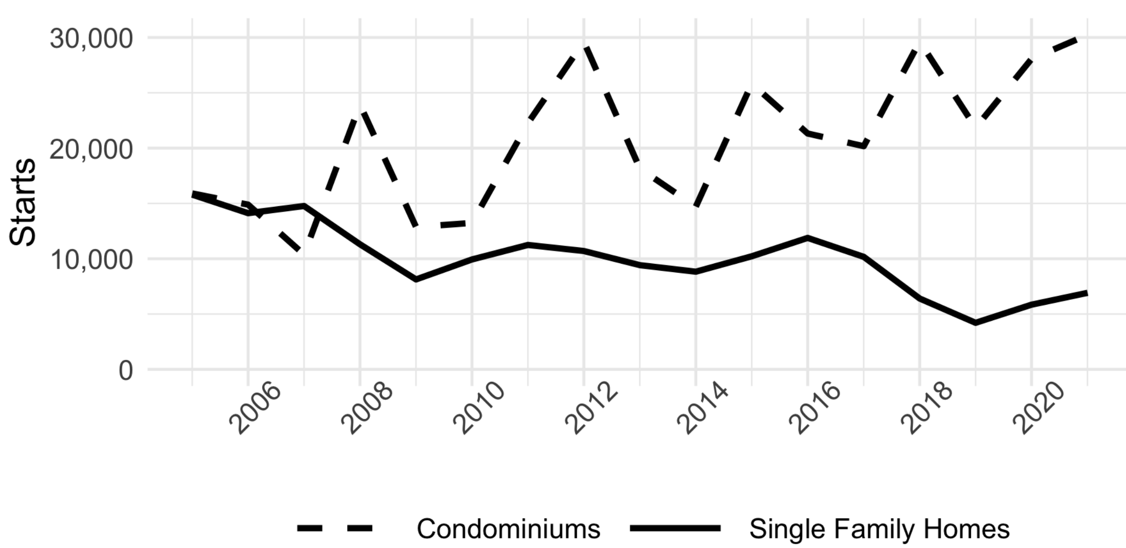

New housing starts have been flat overall and shifted from single family home constructions to condominiums. Source: Canadian Mortgage and Housing Corporation.

the council-appointed Committee of Adjustment.[3] Since there is very little vacant land

[3]Typical of these smaller projects are single-family homes and up to four-storey buildings that seek permission to minimally exceed the existing maximum height or minimum road setback zoning

left in a mature city like Toronto, most proposals to expand housing supply are midrise to highrise structures. The City's legislative structure positions City Councillors to act as the primary decision making mechanism responsible for approving or denying applications for additional housing and densification.

2.1 The Planning Process

When developing housing in Toronto, there is a complex political process that precedes the involvement of City Council. It begins with a developer identifying available plots of land or existing real estate that could serve as the site of future (re)development. Prospective locations are vetted based on several criteria, including the regulatory conditions attached to the site, the location's financial underwriting requirements, and the expected challenges associated with obtaining City Council approval.

Once a site is identified, developers initiate a pre-application consultation with the City. This involves gathering information about development conditions and site requirements from various City departments such as Transportation, Planning, Heritage, and Parks and Recreation. Where required, developers also reach out to representatives of the committees responsible for administering specialized site plans. Site viability is reassessed at this time; if the location is viable, developers will prepare a prospectus for the site. If not, the project is abandoned.

To prepare the prospectus, developers work with city planners and the site's local city councillor to align their proposal with the councillor's preferences. Representatives and developers negotiate specific aspects of the development including building height and density, architectural and landscape characteristics, and the provision of public amenities. Councillor support for projects can vary substantially. In some cases, specific councillors have gained a reputation for outright opposing all proposed developments, whereas others are known to be more pliable. In recent Toronto history, councillors like Jaye Robinson, Mike Colle, and Stephen Holyday have adopted "NIMBYist" philosophies. During their tenure, these councillors regularly invoked concepts like "preserving neighbourhood character" and "opposing billionaire developers" to justify opposition to densification and housing supply initiatives in their wards and across the City. For example, in 2015, Robinson successfully blocked the construction of 22 townhouse units in her electoral ward on the basis of protecting neighbourhood character (Peksa, 2015).[4]

conditions.

[4]Similarly, Colle, Holyday, and Robinson have garnered a reputation for actively voting against

Proposals often undergo several rounds of revision during consultation. For developers, the ideal outcome is the creation of a financially viable and politically supported development. It is common for the council to align their decision with the project's local representative. This is not always the case, however, as meeting records show many instances where proposals are passed despite being opposed by the local councillor.

Completed proposals are submitted to the city planning office for review, which again includes several rounds of application processing such as revisions and public consultation. We provide details of this review and other specific procedures leading to council voting in Appendix A. For applications involving rezoning, the final determination is made through a vote by city council. The council can decide to approve the proposal, to send it back for revisions, to defer the vote, or to deny the project.

If denied, developers can have the decision reviewed by the Ontario Land Tribunal (OLT).[5] Established by the Province of Ontario as the Province's development dispute appellate body, the OLT is charged with determining whether a proposal meets the policy requirements of a given site. If the appeal is upheld the city will be required to accept the proposal without revision. The existence of the OLT curtails council power to capriciously reject proposals. However, the costs and delays associated with an OLT appeal provide strong encouragement for developers to gain council approval.

During the review process, city officials may require developers to include affordable housing or other civic amenities in exchange for allowing more units to be built. This type of negotiation is permitted under Section 37 of Ontario's Planning Act. Section-37 projects often contain more affordable housing units than other developments. In our empirical analysis we examine heterogeneous effects among projects that had a Section-37 agreement.

The complexities of the approval process suggest that councillors wield limited, but meaningful power over housing supply. In our analysis, we only study one stage of the process: the council vote. The voting behaviour of councillors will reflect their desired

densification and housing initiatives while also rallying public and political opposition to broader pro-housing initiatives (Omstead, 2020; Bozikovic, 2022; Vyhnak, 2023).

[5] The OLT went through two rounds of name change and restructuring. It was named the Ontario Municipal Board (OMB) prior to April 3, 2018 and Local Planning Appeal Tribunal between April 3, 2018 and June 1, 2021, during which period, its power was limited and its scope reduced (Willing, 2019). The reduced power of the OLT likely made councillors more willing to say yes to development, as prior to that councillors relied on the OMB to approve development while they catered to the NIMBY demand of residents. OMB changes affected all wards in the same way. Our regression approach will control for time fixed effects (in the form of bill fixed effects) and can therefore account for these changes.

outcomes, but will also serve as a political signal to their constituents. We consider councillor voting behaviour as an indicator of a councillor's support for housing.

2.2 Political Context and Redistricting

Toronto councillors are directly elected from small jurisdictions (called "wards") which usually consist of neighbourhoods with similar characteristics. One councillor possesses one vote, as does the mayor, who is elected at large. From 1998 to 2018 Toronto was divided into 44 wards, but in 2018 the wards were redrawn and reduced to 25. The ward boundaries are shown in Figure 3. After the redrawing of wards, incumbent councillors were forced to compete for the limited council seats. The change was imposed on Toronto by the provincial government as a cost saving measure. The provincial bill to shrink the size of council was not introduced until July of 2018, part-way through an active Toronto election campaign, with the municipal election set for October 22, 2018. Some councillors attempted to block the legislation. The Supreme Court of Canada decided in a 5-4 decision on October 1 that the redistricting measure was valid under the province's powers, allowing the new ward boundaries to be applicable for the election occurring three weeks later.

The redistricting event provides an exogenous shock to a councillor's political incentives as they were confronted with a new set of constituents. Of the 44 original councillors, 21 were re-elected in a new ward. The continuation of councillors across the event allows us to observe the same councillor's voting behaviour under disparate exposure to homeowner and renter constituencies. The redistricting provides additional statistical variation in councillor homeowner exposure, which allows us to better identify the effect of constituency composition on councillor voting behaviour. Additionally, variation generated by the redistricting allows us to isolate the idiosyncratic components of councillor behaviour from effects that are caused by the ward itself, something that has not been possible in prior literature.

The City of Toronto is ideal for examining our question for several reasons: (1) Rising home prices did not translate into a proportionate increase in housing supply, suggesting that regulatory factors are preventing a robust supply response; (2) the ward system and redistricting event provides a simple setting to model councillor incentives, (3) The City maintains complete public records of all bills put before council, extending back to 2009, as well as the complete voting histories of councillors. The data sets are available through the City of Toronto Open Data Portal. Electronic access to these

Figure 3: City of Toronto Ward Boundaries

A. Pre 2018 Ward System **B. 2018 to Present Ward System**

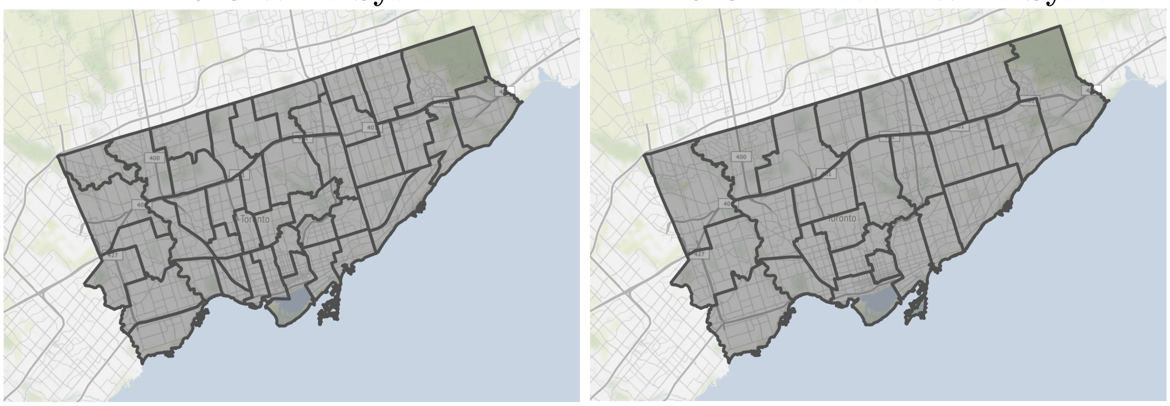

Panel A shows the boundaries of the 44 wards that comprised the Toronto ward system from 1998 to 2018. Panel B shows the ward boundaries after the 2018 redrawing. In the revised system the overall number of wards was reduced from 44 to 25.

public data sets makes it feasible to analyze a large dataset of text through a machine learning (ML) algorithm which we outline in Section 3.

In our analysis we explore whether high-homeownership ward councillors opposed new local housing more strongly than high-renter ward councillors. Figure 4 plots the negative correlation between a ward's homeownership rate and the number of new housing units permitted per square kilometer, using data from 2009-2018 under the original 44 ward system.[6] There is a strong negative correlation between ownership and permitting. New housing is relatively common in high-renter wards and rare in high-ownership wards. Wards with homeownership rates below 50% had an average of 192 units permitted per square km over the period, while wards with homeownership rates above 50% averaged only 51 units permitted per square km. The finding provides some correlative evidence that, in Toronto, new housing is being directed towards low homeownership wards.

3 Data

To examine the influence of homeowners and NIMBYism on new housing construction, we combine four unique data sets covering the City of Toronto from 2009 to 2020: (1) a database of the full text of all city council bills; (2) a full record of city

[6]The data source is the building permit dataset held by the City, also available on the Open Data Portal.

Figure 4: New Units Permitted and Homeownership Rate by Toronto Ward, 2009-2018

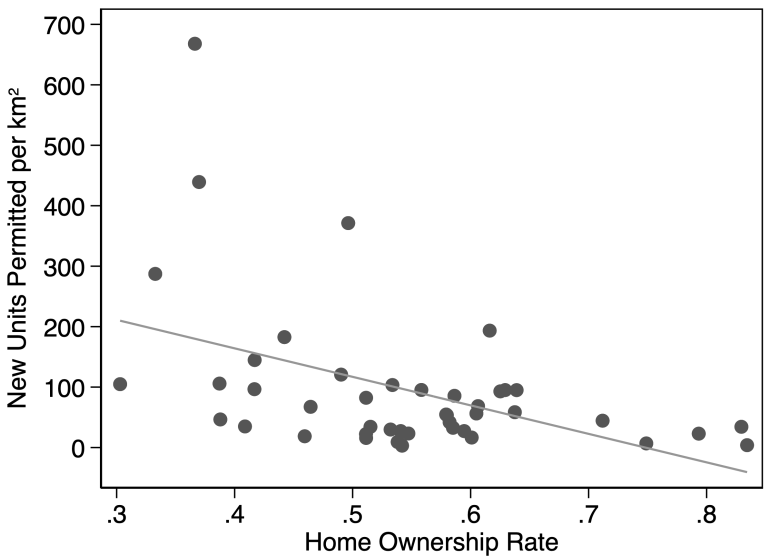

Each point represents one of the 44 wards that existed in Toronto. A fitted line is included. A 10 percentage point increase in the local homeownership rate correlates with 47 fewer units permitted per square kilometer. 21% of the variation in new unit permitting per unit area is explained by the homeownership rate.

councillor voting history; (3) Canadian census data on local homeownership rates and other demographic information at the ward level from the 2016 Canadian Census; and (4) spatial information on ward boundaries through time.[7]

We downloaded these datasets from the City of Toronto's online data repositories, the Open Data Portal. Council bill data covers every bill introduced in the Toronto city council. Each bill is associated with a set of city council meeting minutes as well as the text of the bill. We scraped the full text of all bills and accompanying documents from Toronto's website. Across all bills, there were 8,244 bills and tens of thousands of associated pages.

Manually reading through all bills to identify housing-relevant items would be prohibitively labour intensive. In particular, classifying housing-related bills into those that support higher-density development and those that hinder development requires professional judgment from people who are familiar with urban planning processes in Toronto. This consideration prohibits hiring a large number of research assistants to read the text. For time efficiency and to limit potential human bias, we employed a

[7]The data and code needed to replicate the full results of this paper are available through the Mendeley Data repository. The data and code can be accessed by searching the paper's title on the Mendely Data website.

machine learning (ML) algorithm for textual analysis to classify the bills.[8] We provide an overview of the ML process here and provide further details in Appendix B.

We proceeded as follows: we first broke all documents into about 96,000 "excerpts," with each excerpt containing a paragraph of text. We manually tagged 12,000 excerpts as either "pro-housing", "anti-housing" or "not housing relevant". "Pro-housing" excerpts propose either new housing or zoning changes to allow a higher density of housing. "Anti-housing" excerpts prohibit development by proposing new zoning restrictions. Common examples of anti-housing measures are more stringent restrictions on building heights or proposals to enact local heritage preservation laws. "Not housing relevant" excerpts contain no language related to housing. We use the manually tagged excerpts to train a ML model, which then predicts which category the remainder of the text excerpts fall in. The housing bills we identify cover both new building constructions and upgrades to existing buildings, both of which have been shown to be important components of new housing supply (Schuetz, 2020).

The ML approach deploys a neural network algorithm, using the text of an individual excerpt as an input observation, and outputs percentage probabilities that the excerpt is pro-housing, anti-housing, or irrelevant. We partnered with a private firm who specialized in ML text classification to facilitate the tagging and execute the algorithm.[9] We aggregate this information to the bill level by flagging council bills that contained at least one excerpt that was predicted to have at least a 10% chance of being housing relevant. We identified 2,566 bills that met this criteria. Subsequently, we manually read through these 2,566 bills to check whether they in fact contained any housing relevant measures. We identified 631 bills that contained specific housing measures and included proposed building characteristics, while the remaining bills had no housing relevant language. 5% of the bills identified as housing relevant did not include specific information on building characteristics and were dropped from the data set. We use these 631 bills to populate our final data set. In addition to classifying the housing bills as pro or anti-housing, we manually created variables for each bill

[8] Using a machine learning (ML) algorithm to read large volumes of texts has been used extensively in political science research. Wilkerson and Casas (2017) provides a comprehensive survey. There are two main types of ML algorithms: unsupervised and supervised learning. Unsupervised learning has been used to discover and identify clusters; for example, Quinn et al. (2010) employed this method to detect policy topics and classify Senate speeches by topic. Supervised learning, on the other hand, is mostly deployed as a labour-saving device; for example, Drutman and Hopkins (2013) used this method to identify which emails out of 250,000 Enron emails were political in nature. Our ML algorithm belongs to supervised learning. For a recent ML application to housing supply research see Stacy et al. (2023).

[9] The firm is Sigtica (at sigtica.com)

to capture particular characteristics of the proposed housing. In particular, we create variables that capture the number of housing units in the proposed project, the proposed height of the building in meters and whether the project was covered under Toronto's Section-37 public benefit program. We make use of these variables to test for heterogeneous support across project types. The median height of a proposed housing development in the final data set is 12 meters, the median number of units proposed is four, and 33.1% of proposed projects have a Section-37 agreement.

With the classified housing-related bills in hand, we then matched our bills with the city council's voting records (data set 2). The voting data contains unique identifiers for each bill, allowing them to be linked to bill information. The voting decisions of councillors on housing bills provide the dependent variable of our empirical model, which will be discussed in Section 4.3. For 17% of potential votes, across our sample of 631 bills, the councillor is marked as "Absent" in the data. We drop individual bill-councilor observations where the councillor is marked as absent.

We structure our data set so an observation is at the bill-ward level. Every ward has a representative who can vote on every bill. The final data set is a panel with 23,135 observations, covering 631 bills, 69 unique wards and 78 unique councillors, spanning January 2009 to November 2020.

3.1 Representativeness of Sample

As outlined in Section 2.1 and Appendix A, not all new housing built in Toronto requires action by the City Council. Small projects, such as single-family housing, that can be built with only minor variance or no adjustment to current zoning are often able to proceed without Council involvement. Therefore, our sample is not representative of the entire new housing supply, but is skewed towards larger and more complex projects.

According to city permit data from our study period, among projects that applied for permits, 87% were for projects that would create a single-unit of housing. In our final data set of council bills, 31% are for projects that create a single unit of housing.

While our sample skews towards larger projects, the 631 council bills involve projects in a diversity of neighbourhoods. We map the location of the primary site of each bill in Figure 5. The spatial distribution of bills are correlated with existing population density, but extend to all parts of the city. In Appendix C, we include maps contrasting the spatial distribution of our bills, permitted housing, and the net change in housing stock over our study period. In addition, we also compare the average ward characteristics of permitted housing against average ward characteristics of the housing units covered by our bills. We find average neighbourhood demographics are similar between our bills and the full set of permitted multifamily housing.

Figure 5: Location of Proposed Housing Related Bills

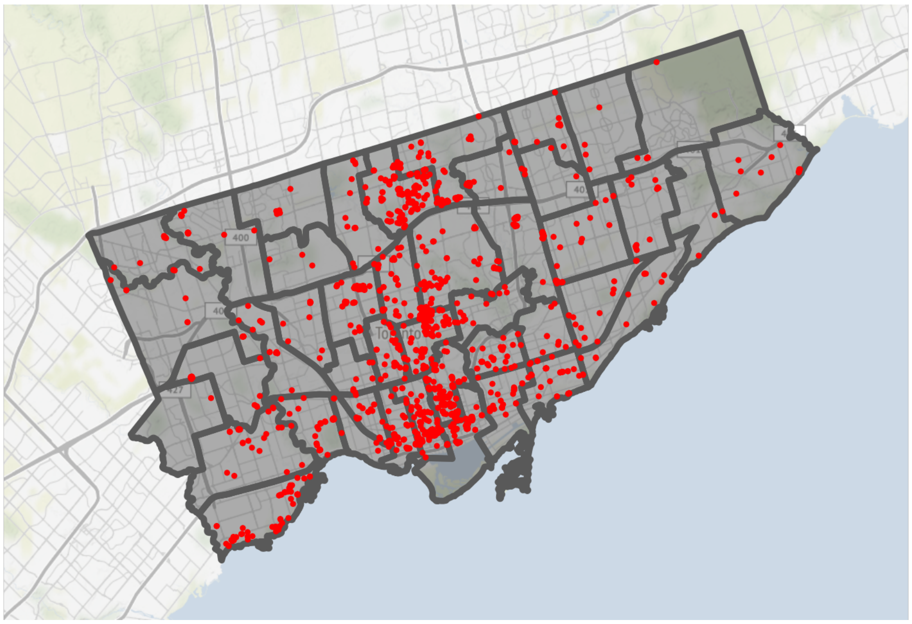

Each red dot indicates the primary location of one proposed bill. The original 44 ward boundaries are shown for reference.

3.2 Construction of Measures and Summary Statistics

We generate a variable that takes a value of one if the councillor voted to support a new housing measure. For pro-housing bills, voting to advance the measure indicates support for housing. For anti-housing bills, voting to oppose the measure indicates support for housing.[10] We observe votes on several different types of motions. We differentiate between motions meant to advance the bill and those meant to prevent the bill from advancing. Motions to "Adopt Item" or "Adopt Item as Amended" are the simplest motions, which represent an up or down vote on a final version of a bill. Similarly, motions to "Introduce Motion without Notice" are meant to advance a new bill within council. Conversely, several types of motions are meant to stop the bill from advancing towards becoming law, including "Receive Item," "Refer Item," "Defer

[10]We treat voting in favour of a pro-housing bill the same as voting against an anti-housing bill. Both of these votes are effectively a vote to support more housing.

Item," "Defer Item Indefinitely," and "Withdraw a Motion." "Receiving" an item means that the Council acknowledges the bill but chooses not to pursue further votes. "Referring" a bill means it is sent back to a lower committee. We treat a vote in favour of these motions as a vote that opposes that bill. We also observe motions to "Waive Referral," which means the council is giving permission for the bill to skip consideration by a committee. We interpret votes in favour of waiving referrals as an indication the councilperson supports the bill. To avoid double-counting bills, we retain only one council vote per bill. We keep the final vote that appears in the data. For example, if council voted to "Introduce a Motion without Notice," then voted down a motion to "Defer Item" and then voted to "Adopt Item," we retain only the vote to "Adopt Item." We only consider full council votes, ignoring committee level activity.

In the final data set, 91.9% of all councilperson votes are in favour of housing.[11] Of 23,135 votes, 1,873 are in opposition to new housing. We attribute this high rate of housing support to the process that precedes a bill coming up for a vote. If a councillor expects a bill to fail, it is unlikely that the bill will be introduced to begin with. From this perspective, the estimated effects in our study can be seen as a lower bound on councillor influence over housing approval.

To estimate the role of homeownership on councillors' voting behaviour, we incorporate Canadian Census data on homeownership rates (data set 3). Ward level data are computed from confidential micro-level Census data by the City of Toronto to reflect ward level averages. We obtain separate data for both the old wards and new wards. Census data is based on responses from the 2016 Canadian Census.[12] The average homeownership rate across observations is 54.1%, with a range of 29-83%. Other demographic characteristic data are also obtained from the ward level census data provided by the City of Toronto. Table 1 summarizes the key variables in our data.

For all bills in the final data set we geocode the address of the proposed new housing development and assign it to a ward (data set 4). Where a bill concerns multiple locations we assign that bill to multiple wards.

[11] Across the full set of votes on 6,146 proposed bills covering all topics, we find that 92% of the 227,281 councillor votes are in favour of supporting the motion. In our final housing specific data set, 97% of votes are in support of the motion. Because some motions are to block new housing, we classify 92% of votes as a "vote in favour of housing." Therefore, all types of bills (but particularly housing bills) seem to garner very high levels of average support.

[12] We elect not to incorporate 2011 Census data due to well documented inconsistencies that resulted from the 2011 survey being voluntary.

Table 1: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.
Vote in favour of housing	0.919	0.273	0	1
Population	64397	15135	45595	129080
Homeowner percentage	0.541	0.121	0.291	0.834
Single person household share	0.303	0.098	0.129	0.538
Immigrant share	0.464	0.121	0.240	0.692
Of European descent share†	0.493	0.198	0.096	0.793
University ed. share	0.354	0.138	0.134	0.609
Labour force participation rate	0.646	0.057	0.542	0.811
Unemployment rate	0.053	0.007	0.041	0.075
Median household income (log)	11.099	0.168	10.788	11.634
Median age	39.538	3.067	32.000	45.200
Distance to City Hall (km)	10.918	5.692	1.199	23.732
N		23,135		

Summary statistics at the ward-bill observation level are provided. † The 2016 Canadian Census asked, "What were the ethnic or cultural origins of this person's ancestors?" Statistic Canada then categorizes responses.

4 Empirical Strategy

We first discuss the conceptual framework and the causal relationship between homeownership and councillors' votes based on theoretical work in the literature. We use a directed acyclic graph (DAG) to help us describe the various channels through which homeownership can affect councillors' votes on residential development. We then design an estimating equation to reflect these channels and control for potential confounding factors.

4.1 Conceptual Framework

Our conceptual framework is based on the implications of urban growth theories in the literature, in particular, Ortalo-Magné and Prat (2014). Using an overlapping generation model, Ortalo-Magné and Prat (2014) find that homeowners are more likely to oppose urban expansion and development because new housing reduces future rent faced by all residents, which in turn lowers the property value of homes that owners have invested in.[13] At the same time, some literature has argued that residents, includ-

[13] Molloy et al. (2022) showed that supply restrictions increase both prices and rents, but also increase the price-rent ratio due to the capitalization of future rent increase expectations into current prices. The loosening of supply restrictions could therefore cause a decline in prices that is large relative to

ing renters, display resistance to new housing because of concerns with neighbourhood change or local congestion externalities (Glaeser et al., 2005b; Kahn, 2011). Diamond and McQuade (2019) demonstrate, for the case of affordable housing developments, that new housing can represent an amenity in low-income neighbourhoods but a disamenity in high-income neighbourhoods. Therefore, the particular characteristics of a development as well as the local demographics are important determinants of whether new housing will enjoy public support.

Our conceptual framework also draws inspiration from urban regime theory (Fainstein and Fainstein, 1983; Mossberger and Stoker, 2001). The urban regimes framework conceptualizes local governance as the result of local governments and external actors engaging in a series of complex repeated games, wherein coalitions and rivalries emerge as each actor wields their power in an attempt to secure desirable outcomes (Keiser, 2015; Russo and Scarnato, 2018; Griggs et al., 2020; Bua and Davies, 2022). Particularly in the case of Keiser (2015) and Griggs et al. (2020), work on urban regimes demonstrates how external actors can use voting, lobbying, and community pressure to tilt local decision making in favour of their interests.

We consider three types of agents in our framework: homeowners, renters and city councillors. Homeowners are motivated by maximizing the value of their housing asset (Fischel, 2002). Renters are motivated by lowering rents. Both homeowners and renters are motivated by preserving or improving the level of local amenities. City Councillors are motivated by securing their own reelection by satisfying the motivations of owners and renters within their own ward. When a councillor approves more housing they put downward pressure on prices and rents. When a councillor approves housing within their own ward they risk altering the local amenity level by introducing congestion externalities or changing neighbourhood characteristics. If the new housing is considered a local disamenity, approving the housing will be viewed negatively by local homeowners who bear both financial and local amenity downsides. Local renters have an ambiguous view on new local housing if it presents a disamenity as they enjoy decreased rents, but also experience decreased local amenities. If the new housing is considered a local amenity, the support among homeowners is ambiguous, while renters will support the project. In addition, the political influence of homeowners and renters comes most directly from their ability to vote for councillors of their wards. Both homeowners and renters can vote, though studies have shown that voter turnout for

the rent effect.

renters is lower (Jiang, 2018). Therefore, councillors may be particularly sensitive to the concerns of local homeowners. Beyond voting, residents can lobby their existing councillor to approve or deny new housing by threatening to withhold support in the future.

The scenario described above provides two main predictions. First, conditional on the the new housing characteristics and other local demographics, councillors who represent more homeowners will show more opposition to housing in order to satisfy their constituents.[14] Second, councillors will be less supportive of a new housing project that lowers local amenities if it is located in their own ward, rather than elsewhere, because a reduction in the local amenity level is a concern for all of their constituents. The level of local disamenity caused by a project usually varies by project size: while a small structure often represents an upgrade to local building stock without substantially affecting crowding, demographic change, or neighbourhood character, a large tower project could be seen to contribute significant new local congestion or bring locally unpopular neighbourhood changes. Therefore, councillors' support for housing in their own wards are likely to vary with project size.

We will test these theoretical predictions in our empirical analysis. In particular, we formulate three main hypotheses: (1) councillors who represent more homeowners oppose more housing bills; (2) councillors' support for new housing is affected by whether it is located in their own ward; (3) councillors' support for new housing in their own ward declines as the size of the project increases.

4.2 Data Generating Process in a Directed Acyclic Graph

The directed acyclic graph (DAG) in Figure 6 spells out the channels through which homeownership impacts councillor votes on housing development. Two channels in the DAG are particularly notable: one is lobbying the *existing* councillor to vote against residential development proposals; the other is electing a new councillor whose idiosyncratic preferences towards development represent homeowners' values. Although councillors' personal biases towards development are unobserved, we make use of the discrete change in ward boundaries to employ councillor level fixed effects to control for idiosyncratic beliefs of councillors and isolate the partial effect of homeowner lobbying

[14]This is true regardless of whether the proposed development is in a councillor's own ward, as higher-density development in another ward may be used to justify taller buildings in the councillor's own wards in the future, or may put downward price pressure on the city-wide market.

on councillor votes.

Figure 6: Directed Acyclic Graph

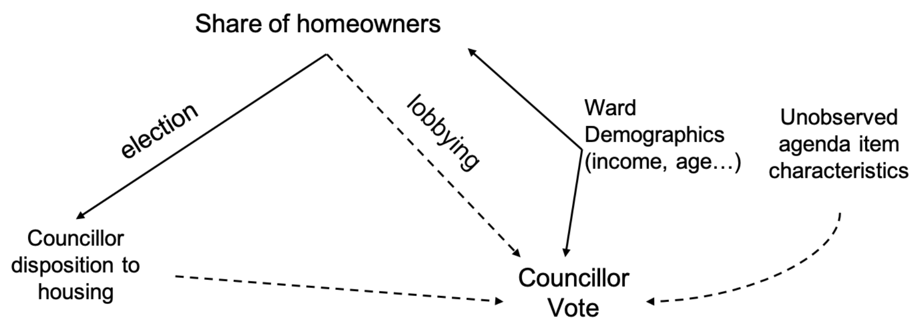

The diagram outlines the connected causal paths that link a ward's homeowner share with the voting behaviour of a councillor. Solid lines indicate observed paths and dashed lines indicate unobserved paths.

An important challenge to the causal interpretation of our analysis is that households endogenously choose to own or rent. Many demographic characteristics, particularly income, affect households' housing tenure choice. The existing literature discusses factors influencing tenure choice. Sinai and Souleles (2013) argued that homeownership reduces the risk of moving because of the positive correlation between the sale price of a household's current house and the purchase price of the next. Sinai and Souleles (2005) argued that households buy as a hedge against market fluctuations in housing cost. Housing is often the largest asset in a person's portfolio and can bring tax advantages (Aizcorbe et al., 2003). In Canada, the primary residence of a household is exempted from capital gains tax. Furthermore, Henderson and Ioannides (1983) argued that households would more likely owner-occupy their homes if their consumption demand for housing was less than their investment demand. Davidoff (2006) showed that if labor income was highly correlated with housing prices, consumers would choose to own less housing. These financial reasons behind homeownership imply that high-income households are more likely to own, although particular high-income households have incentive to rent. In our empirical analysis we control for local income and other demographic characteristics to account for this endogeneity in homeownership choice.

Figure 6 diagrams confounding factors that may jointly affect both the homeownership rate and councillors' votes on development. For example, a ward having older residents could affect both homeownership and councillors' votes, as older people are more likely to own a home, and are more likely to oppose development because they benefit much less from it (Ortalo-Magné and Prat, 2014). We control for confounding demographic conditions directly. In addition, unobserved agenda item characteristics may also affect councillors' votes, for instance, some agenda items may be simply unpopular. As shown in the subsection below, we use agenda item fixed effects to control for these unobservables, which increases the precision of our estimates. While we include a variety of controls, home-ownership may be correlated with other unobservable household characteristics that impact support for new housing. We discuss this limitation to causal inference when interpreting results. The estimated effect of homeownership may include the effect of other correlated demographics that are related to ownership and housing support, but are uncorrelated with the vector of controls we introduce. We therefore estimate the conditional correlation between homeownership and councillor support for housing, rather than a pure causal effect.

4.3 Estimating Equation

The primary econometric model we will estimate is represented by Equation 1:

$$V_{wb} = \beta_0 + \beta_1 H_w + \beta_2 O_{wb} + \theta_w + \Psi_b + \varepsilon_{wb} \tag{1}$$

V_{wb} is a dummy variable that takes a value of one if a councillor from a specific ward (w) voted in favour of a particular housing bill (b). H_w is the share of the ward's residents who are homeowners. O_{wb} is a dummy variable that takes a value of one if a bill (b) concerns a housing project that is within ward w's boundaries. θ_w is a vector of ward level demographic control variables. We include controls for median household income, share of the population with a university education, median age, labour force participation rate, unemployment rate, share of the population that is of European ancestry, share of the population who are immigrants, and share of the households that are single-person households. We also include a control variable for the distance between the ward's centroid and Toronto City Hall, as well as a squared version of this term. City Hall is meant to proxy for the city center. The centrality of the neighbourhood could be correlated with local resident preferences. Ψ_b is a vector of bill fixed effects, which control for the average support for the bill and also controls for any time variation in voting behaviour. We have two parameters of interest, β_1 and β_2. β_1 captures the partial effect of the homeownership rate on the probability the ward's representative voted in favour of a housing bill. β_2 captures the partial effect of a project being within a councillor's own ward on the probability of councillor support.

The statistical variation needed to identify a homeowner effect (β_1) comes from cross-sectional variation in ward characteristics. While we include a battery of control variables to rule out some confounding demographic variation, our coefficient estimate could be impacted by latent demographic variables that jointly affect homeownership and councillor behavior and are not perfectly correlated with our control variables. The estimation of β_2, representing a NIMBYism effect, is derived from variation in the location of specific bills.

We also estimate versions of the model where we control for councillor fixed effects (C_{wb}), shown by Equation 2. Councillor fixed effects absorb councillor level idiosyncratic behaviour. Through the interpretation of the coefficients of councillor dummy variables we can estimate the idiosyncratic disposition towards housing of each councillor. By controlling for councillor effects, we shut down the mechanism that connects homeowners and councillors through elections. Therefore, in Equation 2, β_1 corresponds to only the effect that homeowners wield through the lobbying process, independent of the councillor's idiosyncratic beliefs. If councillors are entirely governed by their consistent idiosyncratic beliefs we would expect β_1 to equal zero in Equation 2.

$$V_{wb} = \beta_0 + \beta_1 H_w + \beta_2 O_{wb} + \theta_w + \Psi_b + C_{wb} + \varepsilon_{wb} \tag{2}$$

While we use a binary dependent variable, we are primarily interested in estimating unbiased partial effects and therefore elect to use OLS estimation, rather than a probit or logit model (Angrist, 2001). In our setting, with a large number of fixed effects, the use of a nonlinear estimator would suffer from an incidental parameter problem which would create biased estimates (Lancaster, 2000). When estimating standard errors we cluster our estimates at the bill level, accounting for potential error correlation in voting behaviour among votes on the same bill.

The 2018 rewarding event provides some additional variation in homeowner exposure among councillors, which enables the estimation of Equation 2. 21 councillors served under both ward systems and the rewarding significantly changed the demographics of their constituencies. Figure 7 shows the changes experienced by these councillors in terms of constituent homeownership share, as well as median household income and European ancestry share. These three variables will be shown to correlate with changing councillor voting behavior. The average councillor who changed wards experienced a 5.3 percentage point change in the homeownership share of their constituency.

Figure 7: Changing Constituencies of Councillors

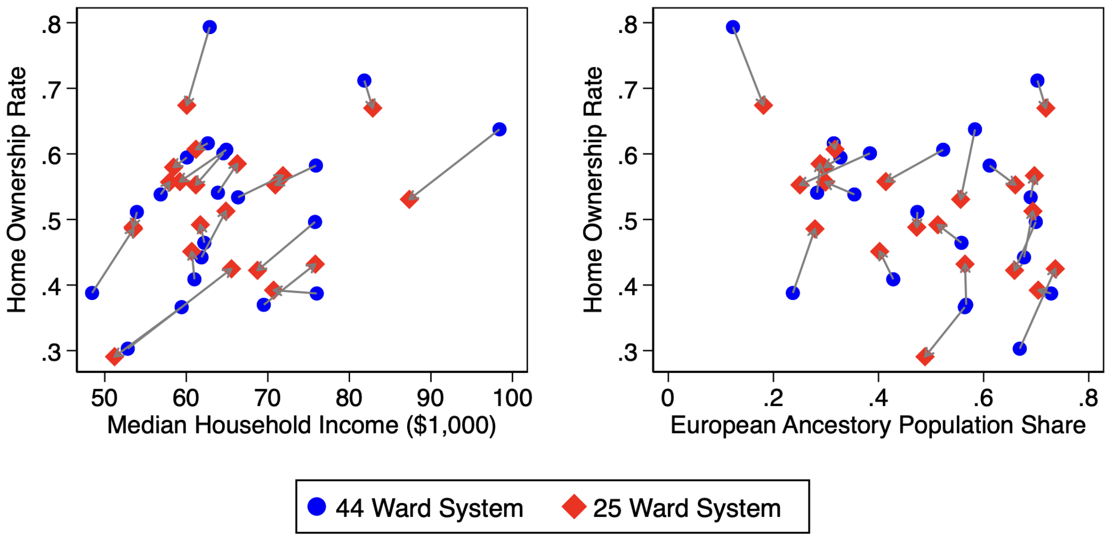

The 21 councillors who served under both ward systems experienced significant changes in the demographics of their constituents. Each connected pair of symbols corresponds to one councillor.

5 Results

We report the regression results in three subsections. Subsection 5.1 tests the first and second hypotheses, discusses the main results from both the Equations 1 and 2 models, and describes the relationship between homeownership and councillors' support for housing. Subsection 5.2 tests the third hypothesis by examining councillors' heterogeneous response to housing development by project type. Results in this section strongly support NIMBYism; that is, councillors are strongly against large projects in their own wards, but not in other wards. Subsection 5.3 looks more closely at the role of local demographic characteristics other than homeownership, showing their effect on average councillor support for housing as well as their effect on the paper's main results. Subsection 5.4 zooms in on the specific councillors who oppose housing consistently and illustrates that voting "no" to housing did benefit some politicians' careers.

5.1 Homeowner Effect

Table 2 provides main regression results, with column 4 corresponding to the specification with a full set of control variables (Equation 1), and column 5 additionally including councillor level fixed effects (Equation 2). Column 1 regresses a dummy variable for whether a councillor voted in support of new housing against the share of the councillor's constituents who are homeowners and includes bill fixed effects. The estimate suggests that a 10 percentage point increase in the local homeownership rate correlates with a significant 1.1 percentage point reduction in the probability that a councillor supports a housing bill. Column 2 estimates the correlation between housing support and whether the proposed housing project is located within a councillor's own ward. The result is not statistically significant. Column 3 includes both independent variables of interest in the same regression, which does not affect point estimates significantly.

Column 4 includes a full set of control variables which reduces the role of confounding variables that might jointly impact councillor voting behaviour and the local homeownership rate. Ultimately, the empirical approach can establish a conditional correlation between the local homeownership rate and councillor voting behavior, but the mechanisms could flow through other confounding variables that are not controlled for.[15] The homeownership rate has a strong relationship with the probability that a councillor supports new housing construction. We estimate that a 10 percentage point increase in the homeownership rate correlates with a 1.3 percentage point reduction in the probability that a councillor supports a given housing bill after controlling for a wide array of demographics. As described in Section 3, 91.9% of votes cast on housing related bills were in support of housing. Therefore, a 1.3 percentage point decrease in support probability translates to a 16.0% increase in the likelihood of opposing a particular housing bill.[16] We therefore find a significant portion of the variation in councillor voting behaviour is tied to the local homeownership rate. These results are consistent with our first hypothesis.

The specific mechanism that links homeownership to councillor opposition could

[15]In Appendix D we provide sensitivity analysis to the omission of particular covariates. We show that our results are robust to the omission of any one covariate, and are generally consistent across dropping any possible combination of covariates. In Appendix D we also provide an alternative analysis methodology by adopting a propensity score matching procedure. As discussed in the Appendix, this method may reduce the influence of omitted, confounding variables. We estimate a very similar homeownership effect using this alternative method.

[16]This is calculated as 1.3% divided by (100-91.9%).

Table 2: Effect of Local Homeownership on Councillor Support for New Housing

	(1)	(2)	(3)	(4)	(5)
Homeowner percentage	-0.114**		-0.115**	-0.130**	-1.856**
	(0.011)		(0.011)	(0.020)	(0.580)
Councillor's own ward (dummy)		-0.005	-0.009	-0.015	-0.017
		(0.009)	(0.009)	(0.009)	(0.008)
Single person household share				-0.143**	2.347**
				(0.041)	(0.794)
Immigrant share				-0.126**	3.065**
				(0.031)	(0.885)
Of European descent share				-0.223**	0.621**
				(0.023)	(0.227)
University ed. share				0.021	-1.384**
				(0.024)	(0.529)
Labour force participation rate				0.275**	-2.199**
				(0.063)	(0.661)
Unemployment rate				-0.425	-2.019
				(0.319)	(1.856)
Median household income (log)				-0.014	2.408**
				(0.022)	(0.833)
Median age				0.004**	-0.029*
				(0.001)	(0.011)
Distance to City Hall (km)				-0.005**	-0.062
				(0.001)	(0.038)
Distance to City Hall squared (km)				0.000	0.003*
				(0.000)	(0.001)
Agenda item fixed effects	Y	Y	Y	Y	Y
Councillor fixed effects	N	N	N	N	Y
R^2	0.612	0.609	0.612	0.616	0.671
N	23135	23135	23135	23135	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis.

flow through additional unobserved omitted variables. In particular, the choice to be a homeowner is endogenous, and homeowners may share characteristics that lead them to oppose housing for reasons other than their financial stake in the housing market or their disliking of the change in amenities caused by the development. We discuss this issue further in Section 5.3.

For the own-ward effect, the point-estimate of column 4 suggests that councillors are 1.5 percentage points less likely to support a housing bill if it is in their own ward, rather than located elsewhere, though the effect is not statistically significant. This result is consistent with our second hypothesis, albeit the link appears weak. In Section 5.2, we explore the heterogeneity of own-ward effects based on project characteristics, finding that the average effect reported in Table 2 masks considerable heterogeneity across project types.

Table 2, column 5 provides results that include councillor level fixed effects, matching Equation 2. When councillor fixed effects are included, all identifiable variation in the homeownership rate is generated by councillors who served as representatives of more than one ward. The redrawing of ward boundaries during our study period provides important statistical variation in the constituency composition of councillors. In the 44 ward system, only two councillors represented more than one unique ward. In the 25 ward system we observe no instances of a councillor representing more than one ward. We observe 21 councillors who served in different wards across the two systems.

As discussed above, there are two mechanisms by which local homeowners may influence councillor voting. First, homeowners might influence the type of councillor who is elected; and second, homeowners may lobby the councillor while they are in office in order to influence their voting behaviour. We find that when the same councillor changes wards, and is exposed to more homeowners, they dramatically reduce their support for housing. For a councillor who moves to a ward with a 10 percentage point increase in homeowners, we estimate they reduce their likelihood of voting in favour of a given housing bill by 18.6 percentage points. Similarly, a councillor exposed to a 10 percentage point reduction in homeowner share would be 18.6 percentage points more likely to support a housing bill.

The large estimate could be driven by a compositional effect. In almost all cases, in order to face a different set of constituents the councillor must have successfully won office in a new ward after the redistricting event. Perhaps this set of councillors is unique. However, we examine the selection issue in Appendix D, by contrasting the characteristics of councillors who survived the redistricting event to the full sample of councillors and find they are affected similarly by the local homeownership rate. The councillors serving under multiple wards do not seem to be inherently different in their voting behavior. Rather, when a councillor gained or lost homeowners due to the rewarding, they significantly changed their voting behavior to cater to the new constituency preferences. The effect of a changing homeownership rate on a specific councillor is large compared to the cross-sectional correlation between homeownership rate and councillor behavior. From this we conclude that a councillor's personal beliefs on housing are often secondary to their willingness to cater to resident housing preferences.

When councillor fixed effects are included we also find a marginally significant (at the 10% level) effect of a project being in the councillor's own ward. We find councillors are 1.7 percentage points less likely to vote for a project if it is in their own ward rather than located elsewhere.

The influence of homeowners on councillor behavior potentially varies by neighbourhood income. In particular, higher income homeowners may be able to exert additional influence on their councillor because they have the capacity to provide significant campaign donations, or because they have disposable time to engage with the public process (Einstein et al., 2019b). Alternatively, high income homeowners might also have stronger opposition to new housing if they place a higher value on natural amenities (Glaeser et al., 2005b; Glaeser and Ward, 2009). Table 3 provides an additional specification that adds an interaction term between the ward's homeownership rate and the ward's logged median household income. After including control variables (column 3), the point estimates imply that homeowner influence is more negative when local incomes are higher. We estimate that the marginal effect of the homeownership share for the ward with the lowest median income ($48,400) is -0.11, whereas in the ward with the highest income ($112,900) the marginal effect of the homeownship share is -0.19. These figures are comparable to the Table 2, column 4 regressions which estimated an overall partial effect of -0.13. The interaction term estimated lacks precision, suggesting the heterogeneous income result should be interpreted with caution.

5.2 Councillor NIMBYism by Project Type

From the text of bills we are able to recover project characteristics and test for heterogeneous responses across different project types. Table 4 shows the correlation between several project characteristics and the probability that a councillor voted in favour of the project. We find that average project support does not vary significantly based on the size of the project, as measured by total proposed units or building height, or based on whether the project has a Section-37 agreement.

We now investigate whether larger projects are more or less likely to be supported by councillors if they are physically located in that councillor's ward. Table 5 provides results from regressions that correspond to Equation 1 and Equation 2 but also include an interaction term between a project characteristic and the dummy variable for the project being within the councillor's ward. Below we focus on discussing the results

Table 3: Effect of Local Homeownership on Councillor Support for New Housing, Heterogeneous Income Effect

	(1)	(2)	(3)
Homeowner percentage	2.149**	2.176**	0.873
	(0.560)	(0.557)	(0.643)
Median household income (log)	0.096**	0.098**	0.042
	(0.028)	(0.028)	(0.047)
Homeowner percentage x			
Median household income (log)	-0.203**	-0.205**	-0.091
	(0.051)	(0.050)	(0.058)
Councillor's own ward (dummy)		-0.009	-0.015
		(0.009)	(0.009)
Agenda item fixed effects	Y	Y	Y
Standard Demographic Controls	N	N	Y
R^2	0.612	0.612	0.616
N	23135	23135	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis. Three regressions are shown. The dependent variable is a dummy variable for the councillor voting in support of housing.

Table 4: Project Characteristics and Councillor Support for New Housing, Bivariate Regressions

	(1)	(2)	(3)
Number of units in project (log)	-0.001		
	(0.001)		
Building height (log meters)		0.003	
		(0.002)	
Section 37 project			-0.004
			(0.004)
R^2	0.000	0.000	0.000
N	23135	23135	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis. Three regressions are shown. The dependent variable is a dummy variable for the councillor voting in support of housing.

that pertain only to the Equation 1 model that does not include councillor level fixed effects. However, Table 5 displays results from both approaches and we find results are robust to the inclusion of councillor level fixed effects. It should be noted that the level effect of the project characteristics on support is absorbed by the bill level fixed effects. Column 1 interacts the own-ward dummy variable with the proposed project's height, in logged meters. We find that taller buildings make the councillor who represents that ward significantly less likely to vote in favour of the bill. Councillors are actually more supportive of housing in their own ward than elsewhere if the project is a low-rise building with a height under 14 meters (or about 4 stories).[17] We do not interpret this positive result to mean that councillors advocate for low-rise buildings in their own wards, but rather that councillors are more open to low-rise development. As the proposed project gains height above 14 meters, councillors become increasingly resistant to the project being located in their ward. For example, a councillor is 1.3 percentage points less likely to support a 20 meter building in their own ward than they would be to support the same building built elsewhere. A councillor is 6.9 percentage points less likely to vote in favour of a 100 meter tower in their own ward relative to the same tower built elsewhere. The estimates identify a NIMBY effect, wherein councillors are more supportive of high-density housing if it is built outside of their own neighbourhood.

Table 5: Interaction of Project Characteristic with Support of Local Council Member

	(1)	(2)	(3)	(4)	(5)	(6)
Homeowner percentage	-0.131**	-1.840**	-0.132**	-1.851**	-0.131**	-1.853**
	(0.020)	(0.577)	(0.020)	(0.577)	(0.020)	(0.579)
Councillor's own ward (dummy)	0.092**	0.071**	0.038**	0.030**	0.001	-0.004
	(0.021)	(0.021)	(0.006)	(0.006)	(0.009)	(0.008)
Councillor's own ward (dummy) x						
Building height (log)	-0.035**	-0.029**				
	(0.008)	(0.008)				
Councillor's own ward (dummy) x						
Number of units in project (log)			-0.019**	-0.016**		
			(0.004)	(0.004)		
Councillor's own ward (dummy) x						
Section 37 project (dummy)					-0.046*	-0.035
					(0.020)	(0.020)
Standard control variables	Y	Y	Y	Y	Y	Y
Agenda item fixed effects	Y	Y	Y	Y	Y	Y
Councillor fixed effects	N	Y	N	Y	N	Y
R^2	0.617	0.671	0.617	0.671	0.617	0.671
N	23135	23135	23135	23135	23135	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis.

In column 3, we find a similar effect when looking at the number of units proposed in the project. Councillors prefer a project to be in their own ward if the project contains seven or fewer units. For projects greater than seven units, the councillor increasingly opposes local construction. For a proposed 100 unit project, a councillor is 4.9 percentage points less likely to vote in favour if the project is in their ward than if the same project is located elsewhere.

[17] The total effect of a project being in a councillor's own ward is 0.092 - 0.035$\times$ log(building height).

Finally, Table 5 tests for own-ward effects of Section-37 projects. For the average non-Section-37 project, we find councillors are just as likely to support a project regardless of whether it is located in their own ward. However, councillors generally resist having a Section-37 project in their own ward, with the probability of support declining by 4.6 percentage points. The finding is very consistent with descriptions of affordable housing supply barriers that focus on NIMBY opposition, wherein new affordable housing experiences fierce local opposition. However, Section-37 projects also tend to be larger, suggesting local Section-37 resistance may be due to resisting large local developments generally.

Overall, the results in Table 5 indicate that large projects face stronger local opposition. These results are consistent with our third hypothesis. According to the financial incentives facing homeowners and renters discussed in our model, homeowners should be more resistant to new large-scale local development relative to renters. In Appendix D, Table D2, we estimate a model where we introduce a triple interaction term to capture the heterogeneous effect of the homeownership rate on resistance to large, locally sited new construction. We provide some suggestive evidence that NIMBY opposition to large, local projects is stronger in wards with high rates of homeownership.

5.3 Dissecting the Impact of Ward Characteristics on Councillors' Votes

The partial effects of ward level demographics shown in Table 2, column 4 are difficult to interpret as they are conditional on controlling for the other demographics. Table 6 provides simple bivariate regressions of the ward characteristics regressed on the councillor's support for housing. The results show that high rates of single-person households, high rates of education, and high rates of labour force participation are all correlated with councillor support for housing. High household income, an older population, and a higher share of ethnically European residents are correlated with less support for new housing. A high homeownership share is also highly correlated with lower housing support in the bivariate regression. While only correlations, the finding that older and higher income neighbourhoods resist housing more strongly is consistent with past theory and qualitative research.

To examine which socioeconomic factors of a ward play the most important role in councillor's voting behaviors, we conduct a standardized regression analysis. The regression specification is exactly the same as that for Column 4 of Table 2 (i.e., Equation

Table 6: Constituent Demographics and Councillor Support for New Housing, Bivariate Regressions

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Homeowner percentage	-0.113**									
	(0.015)									
Single person household share		0.116**								
		(0.018)								
Immigrant share			-0.016							
			(0.014)							
Of European descent share				-0.017*						
				(0.009)						
University ed. share					0.052**					
					(0.012)					
Labour force participation rate						0.120**				
						(0.029)				
Unemployment rate							0.323			
							(0.255)			
Median household income (log)								-0.048**		
								(0.010)		
Median age									-0.004**	
									(0.001)	
Distance to city hall (km)										-0.002**
										(0.000)
R^2	0.002	0.002	0.000	0.000	0.001	0.001	0.000	0.001	0.002	0.002
N	23135	23135	23135	23135	23135	23135	23135	23135	23135	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis.

1), except that all variables are standardized.[18] The coefficients from the standardized regression are independent of the scale of each socioeconomic factor in a ward, and therefore, they can be used to compare the relative contribution of each factor in councillors' voting outcome. The results are shown in Table 7.

As can be seen, of all the socioeconomic factors, the share of European population in the ward is the strongest indicator of councillors' opposition to housing. One standard deviation increase in the share of European population is correlated with a 0.16 standard deviation increase in councillors' voting no to housing projects. Homeowner percentage and labor force participation rate rank the second highest in their contribution to councillors' votes. Consistent with the previous results, greater homeowner percentage indicates stronger opposition from councillors to housing. Higher labor force participation rate, on the other hand, is positively associated with councillor's support for housing. Immigrant share and single-person household rate have the third and fourth strongest relationships with councillors' votes respectively. Both factors exhibit a negative correlation with housing support, conditional on the other demographics. Contrary to the unconditional correlation result from Table 6, median age from the standardized regression has a significantly positive relationship with councillors' voting

[18]To standardize a variable, we demean the variable and then divide it by its standard deviation, such that the standardized variable has a mean of 0 and variance of 1.

Table 7: Standardized Regression Result

	(1)
Homeowner percentage	-0.058**
	(0.009)
Councillor's own ward (dummy)	-0.008
	(0.005)
Single person household share	-0.051**
	(0.015)
Immigrant share	-0.056**
	(0.014)
Of European descent share	-0.162**
	(0.016)
University ed. share	0.011
	(0.012)
Labour force participation rate	0.058**
	(0.013)
Unemployment rate	-0.011
	(0.008)
Median household income (log)	-0.008
	(0.013)
Median age	0.040**
	(0.009)
Distance to City Hall (km)	-0.112**
	(0.020)
Distance to City Hall squared (km)	0.032
	(0.016)
Agenda item fixed effects	Y
Councillor fixed effects	N
R^2	0.616
N	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis.

behaviors, with older wards supporting more development. This difference suggests that although older populations do display lower housing support, it is not due to their age per se, but due to factors correlated with age such as income or homeownership status. None of the other socioeconomic factors are significant indicators of councillors' support for housing.

One interesting coefficient from the table is the ward's distance to City Hall. It is negative, significant, and large in magnitude, suggesting that wards farther away from downtown, i.e. suburban wards, are more likely to oppose housing.

5.4 Who Are the Councillors that Oppose Housing?

Our empirical strategy allows us to estimate the relative position of specific Toronto councillors on their support for housing. The coefficients on the councillor dummy variables estimated from Equation 2 are interpretable as the partial effect of a particular councillor's idiosyncratic behaviour in terms of support for housing bills relative to an omitted councillor. To avoid a dummy variable trap, we omit a dummy variable for the councillor who voted on the fewest items across the study period (Councillor Adam Giambrone). Because ward demographics are controlled for in Equation 2, the estimates represent the councillors' disposition towards housing holding constant the demographics of their constituents. Figure 8 orders all councillors who voted across the study period from the least supportive of housing to the most supportive. Observations to the left of the zero line represent councillors whose idiosyncratic support for housing was less than Councillor Giambrone and those to the right are more supportive.

In Figure 9 we graph the unconditional effects of councillors. The partial effects are calculated using Equation 2, but omitting local demographics. Figure 8 can be interpreted as how far the councillor pulled the voting behaviour of their ward in a particular direction, while Figure 9 better captures the inherent disposition of that councillor towards housing.

The methodology shows that Rob Ford was the councillor with the second strongest idiosyncratic opposition to housing. After Rob Ford was a councillor, he successfully ran for mayor, campaigning as a populist and maintaining the base of his support from suburban, single-family home communities. Also among the councillors who are estimated to most strongly oppose housing are Doug Ford (brother to Rob) and Michael Ford (nephew to Rob), who formed an informal political coalition whose base of support came from single-family zoned, suburban neighbourhoods. After serving as a City councillor, Doug Ford became the leader of the Ontario Progressive Conservative Party in March, 2018 and became the 26th Premier of Ontario in June, 2018. The advances in Ford brothers' political careers are the epitome of motivations behind a councillor's voting behaviour.

Overall, the ranking of councillors produced by the methodology aligns with our expectations regarding how supportive particular councillors were to constructing new housing. The regression methodology of including bill level fixed effects allows councillors to be directly compared in their support for housing even though they may have voted on a different set of bills during their tenure.

Figure 8: Idiosyncratic Councillor Effects, Conditional on Constituent Characteristics

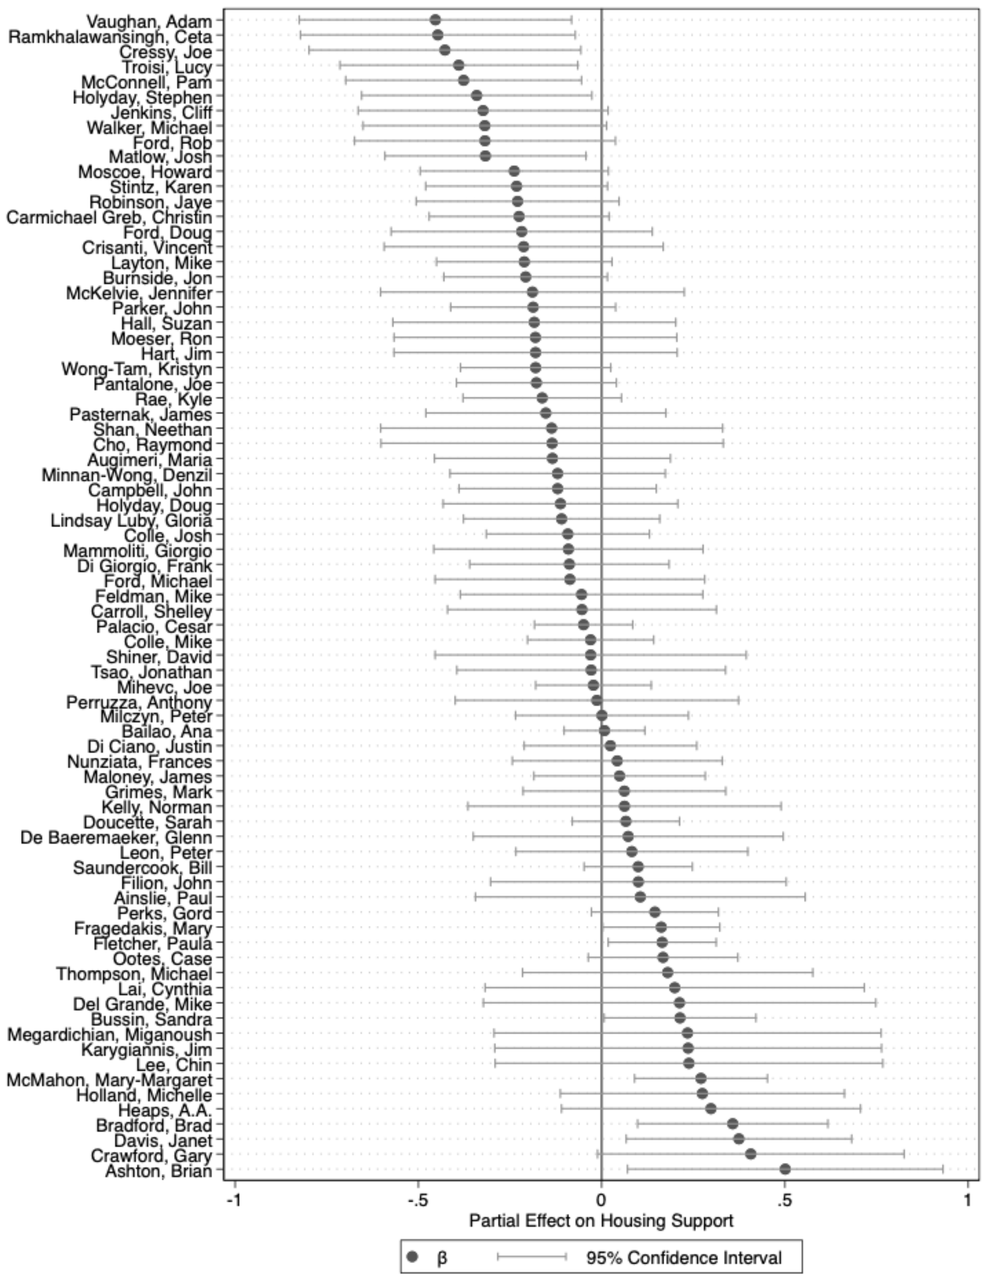

The partial effects of councillor dummy variables, generated by Equation 2 including ward demographic controls, are plotted here. More negative coefficients suggest the councillor was more idiosyncratically opposed to housing, conditional on the demographics of their constituents.

Figure 9: Idiosyncratic Councillor Effects, Unconditional

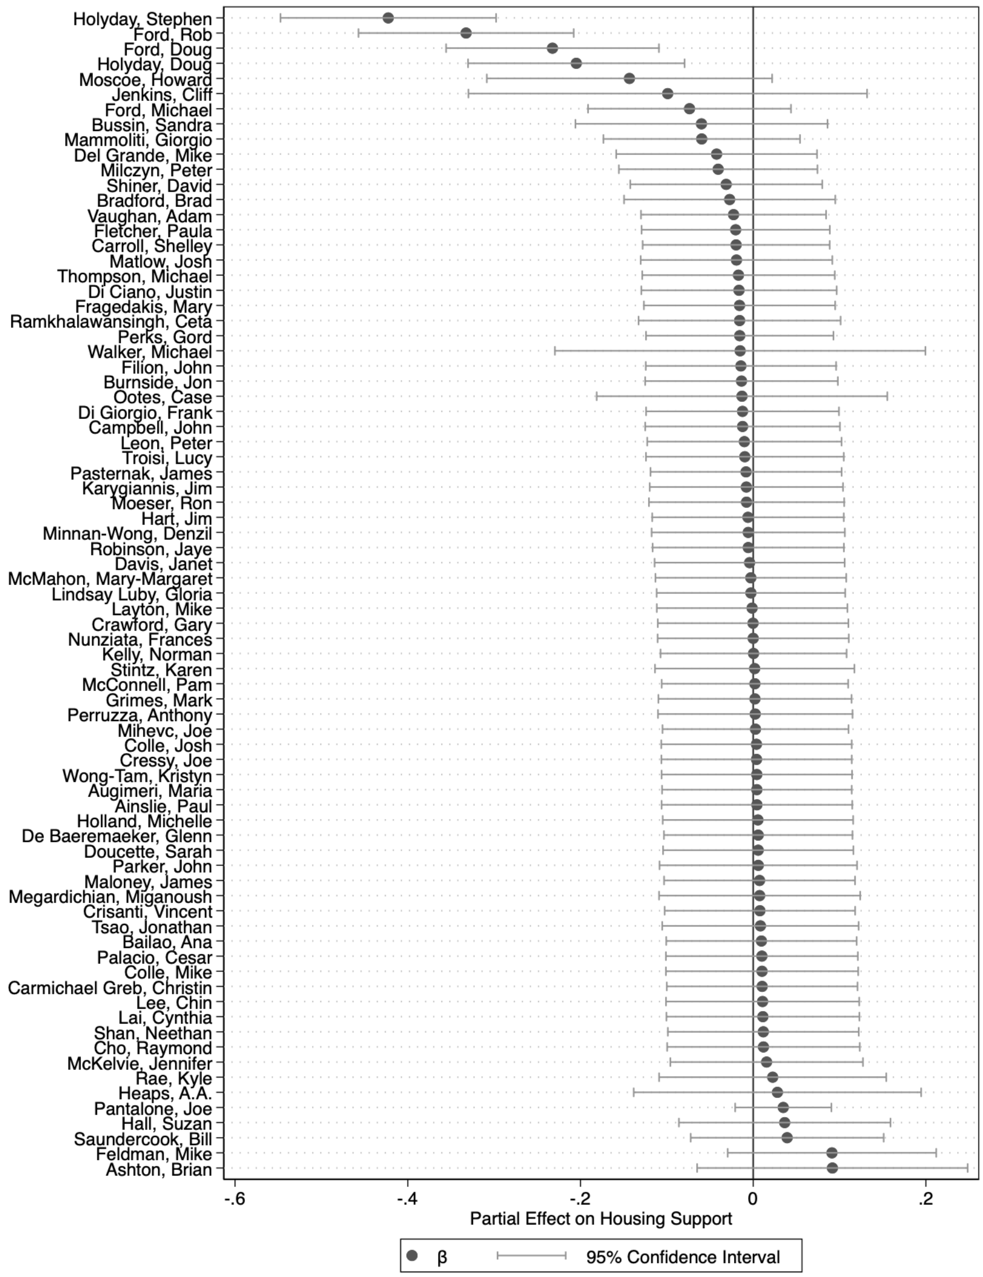

The partial effects of councillor dummy variables, generated by Equation 2 but omitting ward demographic controls, are plotted here. More negative coefficients suggest the councillor was more idiosyncratically opposed to housing.

6 Conclusion

Constraints on new housing impose significant costs on society (Glaeser et al., 2005a; Hsieh and Moretti, 2019; Turner et al., 2014). Local constraints are enacted by elected officials under pressure from voters. Using a machine learning generated data set covering housing bills that came before the Toronto City Council between 2009 and 2020, we estimate the role of local homeownership and NIMBYism on councillor voting behaviour. We find that councillors who face a constituency of more homeowners rather than renters are significantly more likely to oppose the construction of new housing. We directly estimate a NIMBYism effect by testing whether councillors are more likely to block large housing developments if they are in their own ward. Large local housing projects may carry local disamenities in terms of congestion, and may impose a stronger negative effect on local housing values as compared to housing built elsewhere in the city. We find strong evidence that councillors are more resistant of large scale housing development if it is located in their own ward. We find this NIMBY effect is stronger in wards with a high homeowner share.

Fully understanding the political mechanisms that limit housing supply will be important to informing policy that seeks to expand housing supply towards a more socially optimal level. Restricting the supply of new housing imparts large negative externalities on renters, as well as future homeowners who are underrepresented in the local political process. Therefore, pushing housing supply decisions up to the regional, provincial or federal level would allow for these external costs to be better accounted for and would result in a reduced ability of local homeowners to constrain housing supply.

References

Abadie, A. and Imbens, G. W. (2011). Bias-corrected matching estimators for average treatment effects. *Journal of Business & Economic Statistics*, 29(1):1–11.

Ahlfeldt, G. M. (2011). Blessing or curse? Appreciation, amenities and resistance to urban renewal. *Regional Science and Urban Economics*, 41(1):32–45.

Ahlfeldt, G. M. and Maennig, W. (2015). Homevoters vs. leasevoters: A spatial analysis of airport effects. *Journal of Urban Economics*, 87:85–99.

Aizcorbe, A. M., Kennickell, A. B., and Moore, K. B. (2003). Recent changes in us family finances: Evidence from the 1998 and 2001 survey of consumer finances. *Fed. Res. Bull.*, 89:1.

Angrist, J. D. (2001). Estimation of limited dependent variable models with dummy endogenous regressors: simple strategies for empirical practice. *Journal of Business & Economic Statistics*, 19(1):2–28.

Bozikovic, A. (2022). Toronto mayor john toryâs housing plan could bring real change. but is he ready for a fight? *The Globe and Mail.* August 26, 2022.

Bua, A. and Davies, J. S. (2022). Understanding the crisis of New Municipalism in Spain: The struggle for urban regime power in A Coruña and Santiago de Compostela. *Urban Studies.*

Burn-Murdoch, J. (2023). Repeat after me: building any new homes reduces housing costs for all. *Financial Times.* September 14, 2023.

Chen, Y. (2015). Convolutional neural network for sentence classification. Master's thesis, University of Waterloo.

Cheung, R. and Meltzer, R. (2013). Homeowners associations and the demand for local land use regulation. *Journal of Regional Science*, 53(3):511–534.

Clarke, W. and Freedman, M. (2019). The rise and effects of homeowners associations. *Journal of Urban Economics*, 112:1–15.

CREA (2022). MLS home price index methodology. *Canadian Real Estate Association.*

Davidoff, T. (2006). Labor income, housing prices, and homeownership. *Journal of Urban Economics*, 59(2):209–235.

Davidoff, T., Pavlov, A., and Somerville, T. (2022). Not in my neighbour's back yard? Laneway homes and neighbours' property values. *Journal of Urban Economics*, 128:103405.

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

Diamond, R. and McQuade, T. (2019). Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy*, 127(3):1063–1117.

Drutman, L. and Hopkins, D. J. (2013). The inside view: using the enron email archive to understand corporate political attention. *Legislative Studies Quarterly*, 38(1):5–30.

Einstein, K. L., Glick, D. M., and Palmer, M. (2019a). *Neighborhood Defenders: Participatory Politics and America's Housing Crisis*. Cambridge University Press.

Einstein, K. L., Palmer, M., and Glick, D. M. (2019b). Who participates in local government? Evidence from meeting minutes. *Perspectives on politics*, 17(1):28–46.

Fainstein, N. I. and Fainstein, S. S. (1983). *Regime strategies, communal resistance, and economic forces*. Longman Publishing.

Feinerman, E., Finkelshtain, I., and Kan, I. (2004). On a political solution to the NIMBY conflict. *American Economic Review*, 94(1):369–381.

Fischel, W. A. (2001). Homevoters, municipal corporate governance, and the benefit view of the property tax. *National Tax Journal*, 54(1):157–173.

Fischel, W. A. (2002). *The homevoter hypothesis: How home values influence local government taxation, school finance, and land-use policies*. Harvard University Press Cambridge, MA.

Fischel, W. A. (2005). Politics in a dynamic view of land-use regulations: Of interest groups and homevoters. *The Journal of Real Estate Finance and Economics*, 31:397–403.

Glaeser, E. L., Gyourko, J., and Saks, R. (2005a). Why is Manhattan so expensive? Regulation and the rise in housing prices. *The Journal of Law and Economics*, 48(2):331–369.

Glaeser, E. L., Gyourko, J., and Saks, R. E. (2005b). Why have housing prices gone up? *American Economic Review*, 95(2):329–333.

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

Griggs, S., Howarth, D., and Feandeiro, A. (2020). The logics and limits of "collaborative governance" in Nantes: Myth, ideology, and the politics of new urban regimes. *Journal of Urban Affairs*, 42(1):91–108.

Gyourko, J., Hartley, J. S., and Krimmel, J. (2021). The local residential land use regulatory environment across us housing markets: Evidence from a new Wharton Index. *Journal of Urban Economics*, 124:103337.

Gyourko, J. and 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.

Gyourko, J. and Molloy, R. (2015). Regulation and housing supply. In *Handbook of Regional and Urban Economics*, volume 5, pages 1289–1337. Elsevier.

Hankinson, M. (2018). When do renters behave like homeowners? High rent, price anxiety, and NIMBYism. *American Political Science Review*, 112(3):473–493.

Hankinson, M. and Magazinnik, A. (2022). The supply-equity trade-off: The effect of spatial representation on the local housing supply. *The Journal of Politics, forthcoming.*

Henderson, J. V. and Ioannides, Y. M. (1983). A model of housing tenure choice. *The American Economic Review*, 73(1):98–113.

Hilber, C. A. and Robert-Nicoud, F. (2013). On the origins of land use regulations: Theory and evidence from us metro areas. *Journal of Urban Economics*, 75:29–43.

Honnibal, M. (2016). Embed, encode, attend, predict: The new deep learning formula for state-of-the-art NLP models. *Explosion AI*.

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

Hsieh, C.-T., Moretti, E., et al. (2015). *Why Do Cities Matter?: Local Growth and Aggregate Growth*. Number w21154. National Bureau of Economic Research Cambridge, MA.

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

Jiang, B. (2018). Homeownership and voter turnout in US local elections. *Journal of Housing Economics*, 41:168–183.

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

Keiser, R. (2015). Urban regime change: A silver lining for scandals. *Urban Affairs Review*, 51(4):504–532.

Lancaster, T. (2000). The incidental parameter problem since 1948. *Journal of econometrics*, 95(2):391–413.

Lee, D. S., Moretti, E., and Butler, M. J. (2004). Do voters affect or elect policies? Evidence from the us house. *The Quarterly Journal of Economics*, 119(3):807–859.

Mast, E. (2020). Warding off development: Local control, housing supply, and NIMBYs. *The Review of Economics and Statistics*, pages 1–29.

Minaee, S., Kalchbrenner, N., Cambria, E., Nikzad, N., Chenaghlu, M., and Gao, J. (2021). Deep learning–based text classification: a comprehensive review. *ACM Computing Surveys (CSUR)*, 54(3):1–40.

Molloy, R., Nathanson, C. G., and Paciorek, A. (2022). Housing supply and affordability: Evidence from rents, housing consumption and household location. *Journal of Urban Economics*, 129:103427.

Mossberger, K. and Stoker, G. (2001). The evolution of urban regime theory: The challenge of conceptualization. *Urban affairs review*, 36(6):810–835.

Omstead, M. (2020). Historic funding model adopted for toronto community housing: Still not sufficient, councillors say. *The Toronto Observer*. January 7, 2020.

Ortalo-Magné, F. and Prat, A. (2014). On the political economy of urban growth: Homeownership versus affordability. *American Economic Journal: Microeconomics*, 6(1):154–81.

Peksa, S. (2015). Residents and councillor concerned about 22 townhomes proposed on lawrence. *Streets of Toronto*. July 23, 2015.

Quinn, K. M., Monroe, B. L., Colaresi, M., Crespin, M. H., and Radev, D. R. (2010). How to analyze political attention with minimal assumptions and costs. *American Journal of Political Science*, 54(1):209–228.

Rennie, J. D., Shih, L., Teevan, J., and Karger, D. R. (2003). Tackling the poor assumptions of naive bayes text classifiers. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, pages 616–623.

Russo, A. P. and Scarnato, A. (2018). “Barcelona in common”: A new urban regime for the 21st-century tourist city? *Journal of Urban Affairs*, 40(4):455–474.

Schuetz, J. (2020). Teardowns, popups, and renovations: How does housing supply change? *Journal of Regional Science*, 60(3):459–480.

Sinai, T. and Souleles, N. (2013). Can owning a home hedge the risk of moving? *American Economic Journal: Economic Policy*, 5(2):282–312.

Sinai, T. and Souleles, N. S. (2005). Owner-occupied housing as a hedge against rent risk. *The Quarterly Journal of Economics*, 120(2):763–789.

Solé-Ollé, A. and Viladecans-Marsal, E. (2012). Lobbying, political competition, and local land supply: recent evidence from Spain. *Journal of Public economics*, 96(1-2):10–19.

Solé-Ollé, A. and Viladecans-Marsal, E. (2013). Do political parties matter for local land use policies? *Journal of Urban Economics*, 78:42–56.

Stacy, C., Davis, C., Freemark, Y. S., Lo, L., MacDonald, G., Zheng, V., and Pendall, R. (2023). Land-use reforms and housing costs: Does allowing for increased density lead to greater affordability? *Urban Studies*.

Turner, M. A., Haughwout, A., and Van Der Klaauw, W. (2014). Land use regulation and welfare. *Econometrica*, 82(4):1341–1403.

Vyhnak, C. (2023). Triplex conversion a model for housing more Toronto residents. *Toronto Star*. January 20, 2023.

Wilkerson, J. and Casas, A. (2017). Large-scale computerized text analysis in political science: Opportunities and challenges. *Annual Review of Political Science*, 20:529–544.

Willing, J. (2019). OMB back from the dead: Ontario PCs reverting to old rules for planning appeals. *Ottawa Citizen*.

Appendices

A Details of Council Approval Process

This appendix provides additional details on Toronto's process for receiving, reviewing, and approving new development proposals. Figure A1 depicts the City's development approvals processes for proposed housing projects. When a condominium project does not require a rezoning or change to the City's Official Plan, the application proceeds under an alternative process as shown.[19]

Figure A1: Council Approval Process

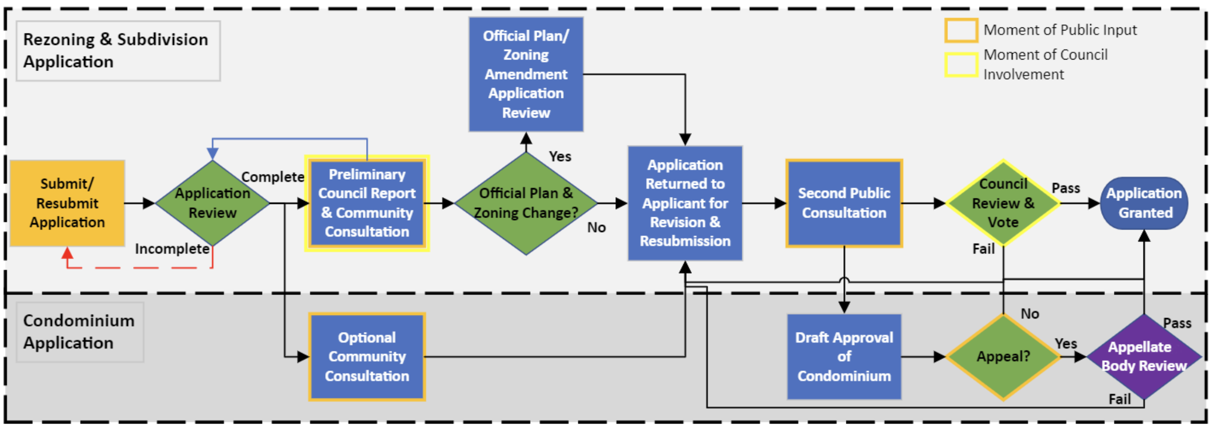

The flow chart summarizes the process of new development proposal and approval in Toronto.

Beginning on the left of the figure, an application is submitted to the planning staff who review the application for completeness and quality. If the application is deemed to be incomplete or of insufficient quality, it is returned to the applicant for revision. If not, it is forwarded to the planning office for review. If the planning office approves the application it will generate a preliminary council report and initiate a preliminary community consultation. If planners find that the application requires a zoning change or community plan amendment, it must pass through an additional planning amendment review process before proceeding.

After the prescribed zoning amendments have been incorporated into the proposal, the applicant is required to subject the application to an additional public hearing. The

[19]For condominium projects that do not require a rezoning, an application does not require a city council vote and is approved after public consultation. The approved application can be challenged by public appeal prior to the initiation of construction, however. Like rezoning and subdivision applications, if an appeal is filed for a condo approval, the final determination is made by the OLT.

applicant must incorporate or provide a response to this feedback prior to resubmitting the proposal.

For rezoning and subdivision applications, the final determination is made through a vote by city council. The council can decide to approve the proposal, to send it back for revisions, to defer the vote, or to deny the project. If the project is denied by council, the developer can choose to appeal the decision to a provincial appellate body, which can choose to approve the project if; 1) it is considered to be consistent with the City's Official Plan; 2) it is consistent with the City's zoning bylaw; and 3) if the appellate body does not consider the reasons for denial compelling.

There are two important moments of council involvement; the first is near the beginning with the creation of the preliminary council report and the second is near the end when council ultimately votes on fully reviewed applications. The analysis of this paper examines only the second instance of council involvement.

The decisions made preceding Council's involvement amount to a strategic, repeated game, played by many parties in the housing development process. The final decision to approve or deny applications for additional housing therefore occurs as a function of the desires of developers, real estate professionals, other development-related professionals, homeowners and renters, and city council.

B Textual Analysis Based on Machine Learning Algorithms

We use textual analysis based on machine learning (ML) algorithms to classify city council bills from the City of Toronto. The classification includes two levels; the first level classifies the bills into those that are housing-relevant and those that are irrelevant. The second level further categorizes housing-relevant bills into those containing pro-housing content and those containing anti-housing content. For example, bills that propose new housing development are labeled as pro-housing, whereas bills that oppose higher-density rezoning are labeled as anti-housing. We partnered with the firm Sigtica (at sigtica.com) to implement the textual analysis. Below we describe the details of the ML method. We start by describing the technical workflow, and then discuss the two ML estimators used.

B.1 Technical Workflow

The following is the technical workflow used to classify bills as irrelevant, housing-relevant, pro-housing, or anti-housing.

Step 1. Web data extraction

We visited the City of Toronto website and downloaded the complete text of all bills and ancilliary documents (i.e., PDF documents) presented before Toronto city council between 2009 and 2020. Automating this step required that all headers linked to bills on the web page were clicked and that enough loading time was allowed so that all the content could load. A web crawler application was used to download these files. All documents crawled from the City of Toronto website were assigned to a separate folder with the identification number of the bill they corresponded to. This process yeilded an initial dataset of more than 8,000 webpages published between 2009 and 2020.

Step 2. Break down bill text into paragraphs

Attempting to classify entire bills using ML can be challenging due to the large the number of words contained in an average bill. The large size can give rise to the curse of dimensionality within algorithmic training due to the large number of words used within the model to predict the likelidhood of a bill being housing relevant. Given that the housing related content of a bill might be limited only to a small portion of the overall bill text, we broke the full text of each bill and accompanying documents into individual paragraphs; in the case of lists, individual bullet points were adopted. We refer to these bodies of text as "excerpts." More technically, whenever a line break appeared in the document, the text followed was classified as a new excerpt. In total, there were over 96,000 excerpts in this granular dataset, which spread across 8,244 bills.

Each excerpt was converted into a matrix of unique word frequencies, then reduced in dimension through text preprocessing algorithms, such as term frequency-inverse document frequency (tfidf), deleting illegal characters (e.g. brackets, asterisks, @, %, ˆ, etc), and lemmatization (accomplished using the Scikit Learn library). All these processes help reduce the size of unique word matrices and to improve the performance of the model. These word matrices were then provided as input to the ML algorithm for prediction.

Step 3. Train a text-classifier using machine learning

Once the individual excerpts are extracted and cleaned, we applied a text-classifier approach to process the bills. This approach falls under supervised ML because we mapped words in excerpts to a labelled outcome variable (housing relevant v.s. irrelevant). To train the model, we first manually read approximately 12,000 excerpts and assigned them to one of three categories: irrelevant, relevant and pro-housing, or relevant and anti-housing. We completed this step with the help of a custom built web browser interface provided by Sigtica. The manually tagged excerpts were then used to train a model, which was subsequently adopted to predict the categories of the excerpts that have not been manually labeled. Given that our classification has two levels, we used a two-stage text classifier approach described in detail as follows:

Step 3.1: Housing relevant/irrelevant text classifier

The vast majority of included excerpts did not contain housing-related content. Because removing irrelevant information from the dataset can help to improve prediction, we first trained a text classifier to dichotomously identify whether an excerpt is housing relevant. The method used here is a two-class model, where the probabilities assigned to "relevant" and "irrelevant" sum to one.

Step 3.2: Pro/anti-housing text classifier

We then ran a second two-class model with the probabilities assigned to "pro" and "anti" housing (also summing to one). Ultimately, we did not make use of this aspect of the ML classification as we were able to manually classify our final set of bills that the ML procedure identified as housing relevant.

Step 4. Use the trained model to predict the master dataset

Once all 96,000 excerpts were assigned a probability of housing relevancy/irrelevancy, the data was then grouped and aggregated back to the bill-level. We generated a database of all bills that contain at least one excerpt of text that contained at least a 10% chance of being housing relevant according to the ML procedure. We found 2,566 such bills.

Finally, we manually read through these 2,566 bills. We found that 631 bills were actually housing relevant, showing that the ML procedure performs well. For example, when the ML procedure estimates a bill contains at least one excerpt that is likely (>50%) housing relevant, our manual reading finds that 45% of those bills are in fact housing relevant. In addition to recording if a bill is relevant to housing, we record whether it is pro- or anti-housing as well as project characteristics such as building height and number of units.

Ultimately, we did not rely completely on the ML approach to classify bills. We used the approach to generate a list of bills that are likely to be housing relevant from the full set of bills, then manually read this set of bills to remove false positives and record bill details. The combination of ML methods and human involvement gives us confidence in the accuracy of the approach.

B.2 Machine Learning Estimators

We investigated two separate ML estimators, described below. In our analysis, we use the probabilities predicted by the Neural Network classifier to filter all housing related bills.

Naive Bayesian approach: Using a popular package called Scikit Learn, we trained a Complement Naive Bayes text classifier. This method implements traditional ML tools, which are not based on neural networks. This method is extremely fast, but under-performs in terms of predictive accuracy. We use this model as a baseline to contrast the performance of other text classification algorithms (Rennie et al., 2003).

Neural Network approach: We use a Parallel Convolution Neural Network (CNN), powered by the Ludwig library, which uses TensorFlow's Neural Network estimation process. The parallel CNN encoder is inspired by Yoon Kim's Convolutional Neural Network for Sentence Classification (Chen, 2015). The overall model architecture is the Encoder-Decoder Network based on Ludwig system, and the CNN falls under the Deep Learning umbrella (Minaee et al., 2021; Honnibal, 2016).

The two approaches discussed above both have limitations. In both cases, a confusion matrix is presented at the end of training to help evaluate the quality of each model. This, in simple terms, shows how well the model can "guess" examples fed to it. The resulting true positives, false positives, true negatives and false negatives together help create the F1 score, which is used to evaluate the quality of these models. Broadly speaking, F1 scores can be interpreted as a type of "accuracy score" in helping users understand how well the trained models work on average. In testing, we found the CNN model outperformed the Naive Bayes model. Therefore, we selected the results from the CNN model.

C Sample Selection and Overall Housing Market Activity

The City Council bills we use to construct our sample may not be representative of all housing production activity in Toronto. As discussed in the main text, the process involved in permitting housing is complex, and not all housing creation requires the involvement of the City Council. In this appendix we compare the spatial distribution of our set of housing bills to other measures of housing production activity.

Figure C1 provides choropleth maps at the 44-ward level, comparing different measures of housing production. Panel A shows the locations of the 631 bills used in the analysis of this paper. Bill activity is concentrated in the central corridor of Toronto. Toronto has a major road (Yonge St.) and a transit corridor running north-south through the geographical center of the city. The ward with the most bills is Toronto Center-Rosedale (Ward 27), with 276 unique bills. The ward covers a high density area with significant condominium construction activity.

Figure C1, panel B shows the distribution of all housing permitted in the city, while Panel C shows permits for multi-family housing only. Panel B shows a concentration of permits in the northern section of the central corridor. The distribution of multi-family housing, shown in Panel C, shows a concentration of permits issued around the southern section of the central corridor. As discussed, our sample of bills is skewed towards multi-family developments, as these are most likely to require council involvement. The maps confirm this pattern, with the spatial distribution of our bills generally correlated with the pattern of permitted housing.

There are several reasons why the bill sample would not capture all the variation in permitted housing. In addition to some housing projects not requiring council involvement, council involvement does not guarantee a permit will be issued. Some bills are rejected, preventing a permit from being issued. Among bills approved, developers could abandon a council approved plan before receiving a permit, for example if it became economically unviable.

Figure C1, Panel D shows the actual net change in housing stock across wards from 2006-2016. The time period is slightly different as the data is only available in Census years. While the other maps only count new unit activity, Panel D will also account for lost units due to demolitions. The largest net increase in units is found in the southern area of the central corridor, which generally mirrors the spatial pattern of multi-family housing permits, and is also an area where a significant share of our bills are located.

Figure C1: Comparison of Bill Sample, Permitted Housing, and Net Housing Stock Change

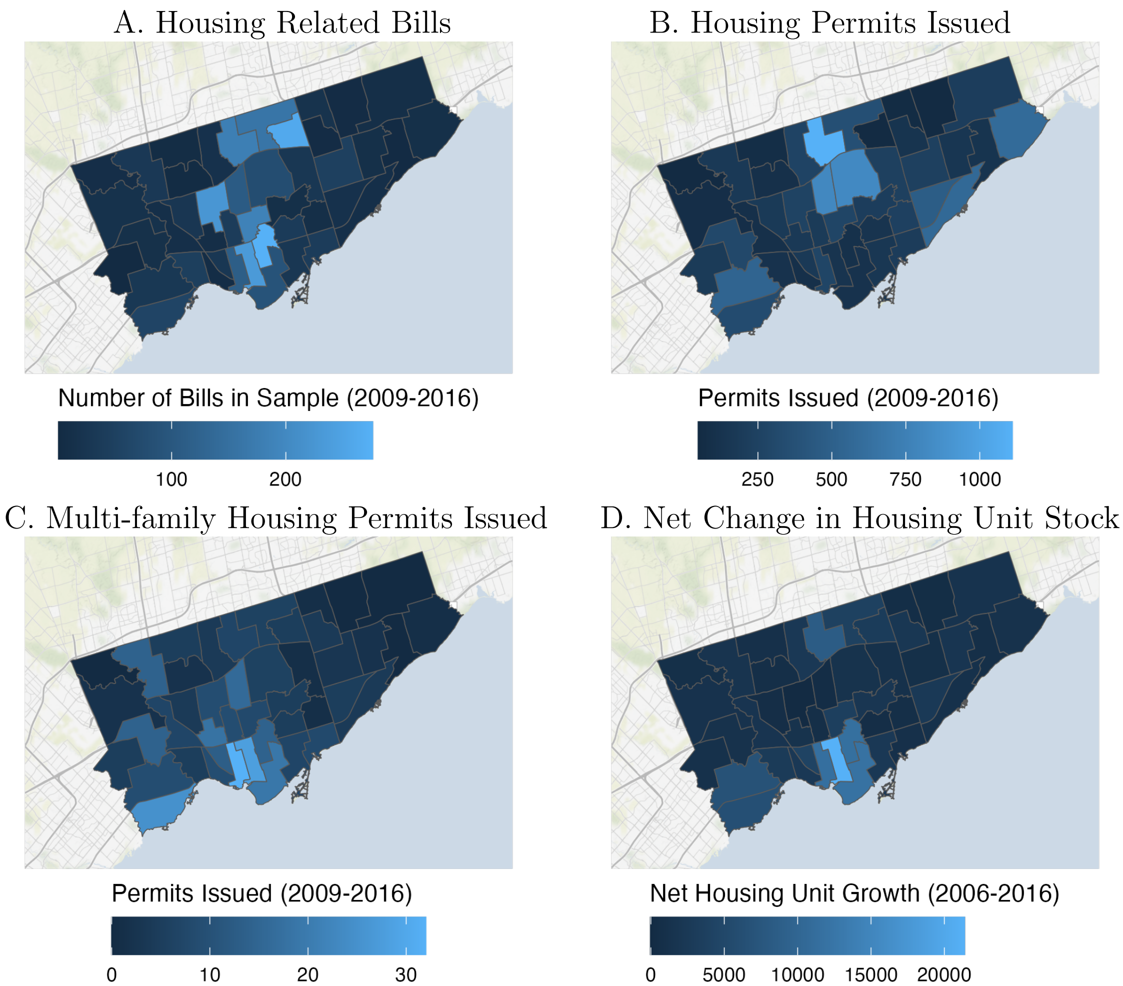

The maps compare housing activity across our set of council bills, permit data from the City of Toronto, and housing stock data from the Canadian Census.

Table C1 provides neighbourhood demographic characteristics for the average housing unit that was voted on in the council data, contrasted with the demographic characteristics of the average housing unit that was actually permitted. Our sample of 641 bills cover 94,687 proposed units of housing. Over the same period, 43,928 units of housing were actually permitted. The local demographics of the two samples are not drastically different. The housing represented in observed council bills are somewhat more likely to come from areas with higher education, higher income, and lower homeowner share. The differences in the data are consistent with oversampling from denser areas where multifamily proposals were more common.

Table C1: Average Neighbourhood Characteristics of Housing Approved by Council vs Permitted

	Mean of Council Approved Housing Units	Mean of Permitted Housing Units
Homeownership share	.456	.498
Single person household share	.410	.369
Immigrant share	.403	.435
European share	.549	.523
University ed. share	.497	.438
Labour force participation rate	.687	.668
Unemployment rate	.051	.052
Median household income	69202	67410
Median age	37.1	38.3
Observations	94,687	43,928

We contrast the average ward demographics for a unit of housing proposed in our sample of housing bills relative to average neighbourhood characteristics of actually permitted housing units.

Though we ultimately analyze bills that pertain to only a subset of overall housing activity, the spatial distribution and local demographic characteristics of our bills are strongly correlated with the characteristics of actual permitting and constructed housing.

D Robustness Tests

Covariate Sensitivity Analysis

The main regression specification (Equation 1) includes control variables based on local demographic information. The inclusion of these covariates is meant to help isolate the influence of local homeownership on the voting behavior of councillors. The choice of particular covariates may impact our parameter estimates of interest. In this appendix we include results from a full set of alternative specifications where we omit some or all of the covariates to test the sensitivity of our results. From the 10 control variables included in the main specification, we construct a list of every possible combination of covariates from among the set of 10. There are 1,024 possible combinations.

Figure D1.i plots the β_1 estimates and Figure D1.ii plots β_2 estimates, with colours indicating whether the result is statistically significant at the 5% level. We find that the main results are generally robust to alternative vectors of covariates. For the home-owership effect, across 1,024 regressions, 95% of estimates are negative and 85% are statistically significantly negative at the 5% level. Results are not dependent on any single control variable as any combination of nine covariates yields a negative and significant result.

For the own-ward effect (β_2), our main specification reported a null effect (Table 2). However the magnitude of the effect is highly robust to omitting covariates, with all 1,024 coefficients estimated as negative. In 23% of regressions, the effect becomes statistically significant.

In the main analysis, we find that the own ward effect becomes stronger as proposed projects become larger. For example, Table 5, column 1 indicates that the likelihood of supporting a bill within a councillor's own ward falls by 0.035 for every one unit increase in the logged height of the project in meters. Figure D1.iii provides results from the same (Table 5, column 1) regression equation, but varies the number of covariates used. We find this result is highly robust, with all 1,024 estiamtes remaining negative and significant. The range of point estimates across the specifications is -0.037 to -0.031.

Propensity Score Matching Method

In the main specifications of this paper we use OLS. In Table D1 we provide additional results by adopting a Propensity Score Matching (PSM) approach. We use a PSM approach to accomplish one-to-one matching among wards. We collapse

Figure D1: Sensitivity of Results to Covariates

i. Homeownership Effect Estimate Across Different Sets of Covariates

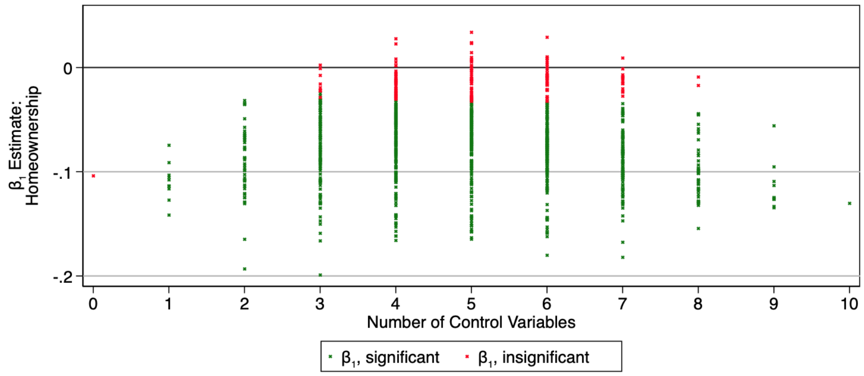

ii. Own Ward Effect Estimate Across Different Sets of Covariates

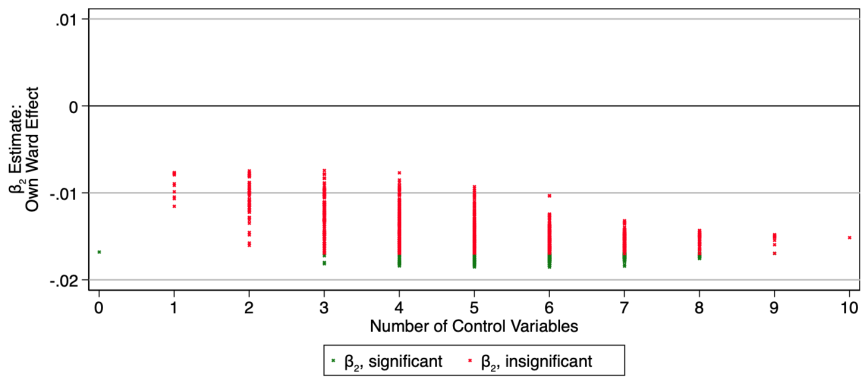

iii. Interaction of Own Ward and Building Height Effect Estimate Across Different Sets of Covariates

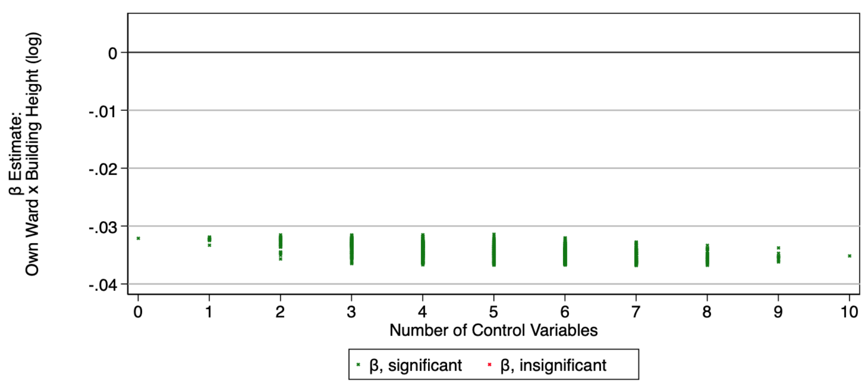

Each subfigure displays 1,024 point estimates, each pertaining to a separate regression. The number of covariates included are varied from left to right.

our data set to the ward level, taking average values across all variables. We divide the sample into low and high homeownership wards, splitting the sample along the median homeownership rate (54.7%). We match each of the 34 high-homeownership wards to a low-homeownership ward by identifying a ward with similar demographic characteristics. We match observations by available control variables. We also add a matching variable for what electoral system the ward was operating under, either the 44 or 25 ward system.

Our main matching result is shown in Table D1, column 4, where we match according to all 11 variables as shown. The result shows that moving from a low to high homeownership ward lowers the probability that a councillor supports a housing bill by 0.019. The OLS result from Table 2, column 4, indicated a 10 percentage point increase in the homeownership rate correlated with a decline of 0.013 in the probability of a councillor supporting housing. In the matching framework, the high-homeownership wards had an average homeownership rate of 63.4% while the low-homeownership wards had a rate of 46.5%. Moving from the low to high homeownership group increases the average homeownership rate by 17.8 percentage points. Therefore, the matching estimates imply that a 10 percentage point increase in the homeownership rate correlates with a decrease of 0.010 in the probability of supporting a housing bill, closely matching our OLS result.

Table D1: Propensity Score Matching: Homeownership Level Effect on Housing Bill Support

	(1)	(2)	(3)	(4)
High homeownership rate	-0.022*	-0.022**	-0.019**	-0.019*
	(0.010)	(.008)	(0.005)	(0.009)
Matching Variables:				
Single person household share	N	N	Y	Y
Immigrant share	N	Y	Y	Y
European share	N	Y	Y	Y
University ed. share	Y	Y	Y	Y
Labour force participation rate	N	Y	Y	Y
Unemployment rate	N	N	N	Y
Median household income (log)	Y	Y	Y	Y
Median age	N	N	Y	Y
Distance to City Hall (km)	N	N	Y	Y
Distance to City Hall squared (km)	N	N	N	Y
Ward system	Y	Y	Y	Y
N	69	69	69	69

Significance levels: $*$: 5% $**$: 1%. Robust standard errors, as described in Abadie and Imbens (2011), are shown in parenthesis.

We also provide results (Columns 1-3) where we use a subset of covariates to match high and low homeownership wards. Results are consistent, generating negative and significant results with a similar magnitude.

Alternative Specification for Estimating Heterogeneous Resistance to Large Developments by Homeownership Level

In Table D2 we interact the homeownership rate with project characteristics. We add control variables as shown in the table. The strength of opposition to large, new housing within a councillor's own ward increases as the local rate of homeownership increases. Columns 1 and 2 show that the partial effect of raising building height within a councillor's own ward on bill support becomes more intense as local homeownership increases (significant at the 10% level). We find similar, marginally significant results when looking at the number of units involved in the project (columns 3 and 4). The coefficient estimates suggest that high-homeownership areas are more amenable to Section-37 projects, but the result is not statistically significant.

Differences Among Councillors Serving in Multiple Wards

In the main analysis we find that adding councillor level fixed effects greatly increases the negative correlation between homeownership rate and councillor support for housing. However, including councillor fixed effects essentially limits analysis to the subset of councillors who served in multiple wards. Potentially, this subset of councillors behaves differently in general. Table D3 examines differences in behavior among councillors who served under both ward systems, relative to the full sample. Comparing columns 2 and 3 suggests that councillors who served under multiple wards are actually less sensitive to the local homeownership rate when voting on housing. Restricting the sample to only examine the period before the rewarding demonstrates a similar pattern (columns 4 and 5).

The result suggests that councillors who served under multiple wards are not uniquely sensitive to the local homeownership rate. Rather, they are sensitive to *changes* in their constituent homeownership rate and are willing to significantly alter their support for housing when they inherit a constituency of more or less homeowners.

Table D2: Interaction of Project Characteristic with Support of Local Council Member: Fully Interacted Specification

	(1)	(2)	(3)	(4)	(5)	(6)
Homeowner percentage	-0.300**	-1.992**	-0.216**	-1.925**	-0.168**	-1.883**
	(0.037)	(0.571)	(0.024)	(0.566)	(0.023)	(0.573)
Councillor's own ward (dummy)	-0.266*	-0.221*	-0.088*	-0.029	-0.065	0.000
	(0.106)	(0.107)	(0.036)	(0.035)	(0.044)	(0.042)
Councillor's own ward (dummy) x Homeowner percentage	0.678**	0.580*	0.231**	0.109	0.123	-0.008
	(0.242)	(0.243)	(0.065)	(0.062)	(0.081)	(0.079)
Councillor's own ward (dummy) x Building height (log)	0.059	0.065				
	(0.041)	(0.042)				
Homeowner percentage x Building height (log)	0.056**	0.055**				
	(0.010)	(0.010)				
Councillor's own ward (dummy) x Building height (log) x Homeowner percentage	-0.181	-0.190				
	(0.100)	(0.101)				
Councillor's own ward (dummy) x Number of units in project (log)			0.004	0.005		
			(0.016)	(0.016)		
Homeowner percentage x Number of units in project (log)			0.030**	0.030**		
			(0.004)	(0.004)		
Councillor's own ward (dummy) x Number of units in project (log) x Homeowner percentage			-0.039	-0.041		
			(0.034)	(0.035)		
Councillor's own ward (dummy) x Section 37 project (dummy)					-0.111	-0.108
					(0.082)	(0.083)
Homeowner percentage x Section 37 project (dummy)					0.106**	0.104**
					(0.020)	(0.021)
Councillor's own ward (dummy) x Section 37 project (dummy) x Homeowner percentage					0.172	0.169
					(0.168)	(0.171)
Standard control variables	Y	Y	Y	Y	Y	Y
Agenda item fixed effects	Y	Y	Y	Y	Y	Y
Councillor fixed effects	N	Y	N	Y	N	Y
R^2	0.618	0.672	0.618	0.672	0.617	0.671
N	23135	23135	23135	23135	23135	23135

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis.

Table D3: Testing for Differences Among Councillors who Served Under Both Ward Systems

	(1)	(2)	(3)	(4)	(5)
Homeowner percentage	-0.130**	-0.343**	-0.156**	-0.340**	-0.146**
	(0.020)	(0.037)	(0.030)	(0.036)	(0.030)
Councillor's own ward (dummy)	-0.015	-0.008	-0.034**	-0.008	-0.031*
	(0.009)	(0.012)	(0.013)	(0.012)	(0.013)
Single person household share	-0.143**	-1.034**	0.398**	-1.040**	0.506**
	(0.041)	(0.100)	(0.055)	(0.100)	(0.058)
Immigrant share	-0.126**	-0.071	-0.193**	-0.076	-0.124*
	(0.031)	(0.049)	(0.060)	(0.050)	(0.060)
Of European descent share	-0.223**	-0.265**	-0.564**	-0.269**	-0.597**
	(0.023)	(0.038)	(0.054)	(0.038)	(0.063)
University ed. share	0.021	0.203**	-0.128**	0.206**	-0.141**
	(0.024)	(0.037)	(0.034)	(0.037)	(0.034)
Labour force participation rate	0.275**	0.912**	0.579**	0.920**	0.801**
	(0.063)	(0.113)	(0.106)	(0.113)	(0.128)
Unemployment rate	-0.425	-4.571**	0.010	-4.586**	0.794
	(0.319)	(0.731)	(0.541)	(0.734)	(0.551)
Median household income (log)	-0.014	-0.088**	0.078**	-0.092**	0.108**
	(0.022)	(0.032)	(0.029)	(0.033)	(0.028)
Median age	0.004**	0.009**	0.003*	0.009**	0.006**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Distance to City Hall (km)	-0.005**	-0.016**	0.012**	-0.016**	0.011**
	(0.001)	(0.002)	(0.002)	(0.002)	(0.003)
Distance to City Hall squared (km)	0.000	0.000**	-0.001**	0.000**	-0.001**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Agenda item fixed effects	Y	Y	Y	Y	Y
Councillor fixed effects	N	N	N	N	N
Councillor sub-sample	Full	Single-ward	Multiple-ward	Single-ward	Multiple-ward
Time period	Full	Full	Full	Pre-rewarding	Pre-rewarding
R^2	0.616	0.638	0.641	0.638	0.669
N	23135	13091	10044	12911	9113

Significance levels: $*$: 5% $**$: 1%. Standard errors are clustered at the bill level and are shown in parenthesis.