

The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement

Alexander W. Bartik, Arpit Gupta, and Daniel Milo*

July 15, 2024

Abstract

We present a novel method called “generative regulatory measurement” that uses Large Language Models (LLMs) to interpret statutes and administrative documents. We demonstrate its effectiveness in analyzing municipal zoning codes, achieving 96% accuracy in binary classification tasks and a 0.9 correlation in predicting minimum lot sizes. Applying this method to U.S. zoning regulations, we establish five facts about American zoning: (1) Housing production disproportionately happens in unincorporated areas without municipal zoning codes. (2) Density in the form of multifamily apartments and small lot single family homes is broadly limited. (3) Zoning follows a monocentric pattern with regional variations, with suburban regulations particularly strict in the Northeast. (4) Housing regulations can be clustered into two main principal components, the first of which corresponds to housing complexity and can be interpreted as extracting value in high demand environments. (5) The second principal component associates with exclusionary zoning.

JEL-Classification: R52, R58, K11, O38, R31, C81

Keywords: housing regulation, zoning codes, large language models, natural language processing, artificial intelligence, municipal ordinances, retrieval augmented generation

*Bartik: Department of Economics, University of Illinois at Urbana-Champaign, abartik@illinois.edu. Gupta: New York University, Stern School of Business, ag5808@stern.nyu.edu. Milo: New York University, Stern School of Business, dm4766@stern.nyu.edu. This paper has benefited from conversations with Jaehee Song (discussant), Tom Cui, Giles Duranton, Ingrid Gould Ellen, Chris Elmendorf, Joseph Gyourko, Stijn Van Nieuwerburgh, Theresa Kuchler, Albert Saiz, Johannes Stroebel, and comments from conference and seminar participants at the NYU Stern Corporate Governance luncheon, NYU Law School, the Online Seminar in Economics and Data Science, Wharton real estate, UNC finance, AREUEA National Conference, the Dallas Fed, and UIUC economics. Preeyonuj Boruah, David Dai, Amber Gao, Grace Getman, Drew Harrington, Karin Hobelsberger, Nagharjun Mariappan, Scott Overbey, Alok Ranjan, Ritul Soni, and Kevin Zhou provided excellent research assistance. We gratefully acknowledge funding from Emergent Ventures and the NYU Stern Center for Global Economy and Business. Any errors or omissions are the responsibility of the authors. A replication package for the project can be found at: <https://github.com/dmilo75/ai-zoning>.

1 Introduction

Housing regulation have a significant impact on housing affordability, patterns of urban development, and productivity growth in the construction sector (Rognlie, 2016; Goolsbee and Syverson, 2023; Glaeser and Gyourko, 2018; Glaeser and Ward, 2009; Gyourko et al., 2008). However, precise measurements of housing regulations remains limited due to the complexity of interpreting diverse zoning ordinances that differ across municipalities. These challenges hinder our ability to grasp the underlying drivers of housing regulation.

Our paper argues that advances in Large language Models (LLMs) enable the scalable and accurate classification of regulatory documents, a task which we refer to as *generative regulatory measurement*. We use state-of-the-art Artificial Intelligence (AI) methodologies to estimate zoning regulations across 63% of the population who lives under a local government subject to zoning ordinances. We validate LLM-generated regulatory categorizations against human-coded measurements from the Pioneer Institute (see Glaeser and Ward (2009)) in the state of Massachusetts, which provides an effective training dataset for our analysis.

Our results suggest that LLMs deployed on the latest models (Chat GPT-4 Turbo) have achieved near-human rates of precision in classifying regulation, with an accuracy rate around 96% for binary questions. LLMs also perform strongly on numerical questions with a correlation of 0.7 between generated data and analyst responses overall, and exceeding 0.90 for measuring minimum residential lot size, an important regulation governing density in single-family areas. We also validate our measures by correlating our generated data against existing surveys of regulation from the Wharton index (Gyourko et al., 2021), finding some association both for specific questions asked across both the survey-based approach as well as our generated data (for affordable housing mandates and minimum lot sizes), as well as for the overall indices.

We use the resulting LLM-produced dataset, alongside other housing data, to establish five key facts about housing regulation in the United States. First, zoning ordinances ap-

pear to broadly bind the new development of housing across the United States. We draw a distinction between incorporated and unincorporated. Incorporated areas are municipalities with established local governments, frequently created to enable more stringent local housing regulations. Over the period from 2000–2020, we find that 42% of the growth in new housing units in the US was concentrated in unincorporated areas, including a majority of new housing units in the South. This is despite the fact that unincorporated regions are home to only 23% of the overall population, and are typically low-demand regions with low prices and rents, and high vacancy rates, far from municipal job centers and amenities. Reliance on greenfield sprawl in far exurban locations for a large share of its housing production suggests meaningful constraints on infill housing production in incorporated areas.

Second, we argue that the average level of housing regulations in incorporated regions are characterized by a low-density bias. Multifamily apartments are explicitly illegal in the entirety of many cities. Using zoning maps for a sample of the largest cities covering 18.2 million people, we find 36% of land area is zoning only for single-family zoning, while multifamily apartments are allowed in only 31% of land area. Among single-family zoned areas in municipalities across the entire country, 66% have minimum lot size requirements above 5,000 square feet, 17% of requirements are above 10,000 square feet, and 7% have requirements exceeding half an acre. While exceptions and changes to these rules are possible through variance processes, these facts highlight how existing regulations mandate low-density housing uses and limit high-density housing construction.

Our third key finding is that housing regulation varies within metropolitan areas in ways broadly consistent with a monocentric city model. Denser building purposes are allowed in the centers of cities, while inner ring suburbs may have greater bulk regulations and lower density requirements. This pattern is particularly more pronounced in the Northeast, which has substantially more strict bulk regulation requirements than other regions of the country.

Fourth, we find that housing regulations can be summarized along two key dimensions: the first of which is associated with housing *complexity*, and proxies for housing demand. To do so, we generate a large set of housing regulation questions, beginning with a set of questions initially generated by the Pioneer Institute, and augment with additional questions on the process of housing regulation. Across this set of regulatory questions, we find that the first principal component is high in regions that have high construction and high prices, and is low in regions with low construction and low prices. This pattern is indicative of a factor generally associated with demand-side pressures. A typical and representative question, which receives high loading in this factor, is the presence of inclusionary zoning mandates: requirements that housing contain a fraction of units with affordability requirements. We find that these and other regulations are generally favored by municipalities with high housing demand, consistent with the notion that these are regulations which extract some housing surplus by municipalities.

The fifth and final fact focuses on the second principal component of our regulatory dataset, which appears to proxy for exclusionary zoning regulations and we interpret as indicating housing exclusion. This principal component loads heavily on bulk regulations such as minimum lot sizes, as well as process requirements in ways that can be broadly characterized as exclusionary zoning. These appear to be regulatory requirements drawn especially in prosperous suburbs around metropolitan areas to further reduce local housing supply and low cost units specifically. This factor, as a consequence, is strongly and negatively associated with housing supply elasticity. We link these exclusionary motives to educational sorting specifically, finding an association between these regulations and local educational test scores.

We interpret these facts in a framework in which municipalities pick housing regulations in the context of competition. Regions with high demand, such as the centers of metropolitan areas and high-housing cost areas, are able to extract housing value. The regulatory tools they use to do so can be interpreted as regulations on the shadow price of

housing. Other regions, particularly suburban entities around major metropolitan areas, instead employ exclusionary zoning to exclude the presence of lower-income residents. These practices lower housing supply and density across municipalities in the United States, while pushing housing production to far exurban areas outside of zoned municipalities.

The key contribution of our paper is the development of a methodology which is able to scalably and precisely measure the content of textual documents. Our focus is creating a national database with granular information on local zoning codes, due to their importance in shaping housing markets. However, challenges in appropriately interpreting and analyzing textual databases are common across multiple domains (in building codes, other regulations, court cases, earnings call transcripts, newspapers, etc.) and so our approach also has broader applicability in suggesting possible approaches towards the analysis of such texts more broadly. Developing such approaches has become increasingly important as the quantity and complexity of regulation has risen over time ([Singla, 2023](#)).

Our results serve as an initial proof of concept towards the use of LLMs in the systematic generation of content in regulatory and legal documents, and suggest these models are a promising tool to measure and generate data in these contexts. We estimate high accuracy rates for such an approach, close to what we might expect for a non-expert but skilled human. While expert humans still have an edge in precise classification, LLM approaches bring several advantages. First, they are far more scalable than human approaches: we are able to deploy our regulatory classification measure across a sample of thousands of municipalities, a task which would be challenging if not prohibitively expensive for humans. Our approach therefore opens up the prospect of scalable and accurate regulatory classifications across multiple domains. Second, the LLM classifications also have the benefit of verifiability and auditability. We prompt the LLM to provide the precise text in the regulatory document which supports the categorization, enabling

other humans (or AI agents) to verify and check the reason for classification. Third, our approach is also flexible to changes in researcher determinations of definitions and subsequent advances in AI models. Researchers using such approaches, therefore, can easily adapt and replicate models over time, accommodating those who prefer alternate specifications or more accurate models.

The granularity of our regulatory estimates across municipal areas also allows us to contribute to the literature on housing regulation more broadly. While our findings on exclusionary zoning are broadly consistent with other hypotheses for the origin of housing regulation such as the homevoter hypothesis (Fischel, 2002); we provide a complementary rationale based on excluding low-income residents from local public goods such as schools. Additionally, we emphasize an extractive component of housing regulation associated with housing demand. The resulting two-factor model differs from prior research, such as Gyourko et al. (2008), which instead emphasizes a single component of regulatory strictness. Our results therefore provide both a more detailed analysis of regulatory codes, as well as an alternate narrative for their economic functions.

Importantly, our results should be seen as illustrating a base level of performance using widely accessible tools, and have considerable scope for improvement along several dimensions. We perform some refinements of question prompting,¹ question background information, and multi-step processing.² Further pre-processing of documents to focus LLMs on relevant text is also likely able to improve model accuracy. Additionally, we use the highest quality LLM available at the time of writing (Chat GPT-4 and Claude 3 Opus), but these models are likely to improve over time. We also plan to expand the scope of this work to examine changes in zoning codes over time, in analyzing housing regulations across countries (including in other languages), as well as in analyzing building codes in

¹This entails rephrasing questions for the LLM through strategies like breaking multi-part questions into different components, and breaking compound questions into individual clauses (i.e., if the question asks about whether multi-family housing is allowed either by right or through a special overlay, we ask about those two possibilities separately).

²Multi-step processing entails breaking a task into multiple steps and querying the LLM separately for each step.

conjunction with zoning codes. Combined, the promise of these efforts suggest that LLMs are likely to fundamentally reshape our ability to understand the content and impact of regulations broadly.³

Contributions to Literature The central contribution of our project is the creation of a standardized, comprehensive dataset of zoning across the United States. Much of the existing literature on housing regulations has used either indirect measures or proxies for zoning regulation. The first strand of this literature has focused on survey-based approaches to measuring housing regulations. One of the most heavily used such nationwide measures of housing regulation includes the Wharton Regulatory Index (Gyourko et al., 2008, 2021; Huang and Tang, 2012). This pioneering approach to measuring housing regulations was based on surveys sent to 2,649 distinct municipalities (there are 19,488 municipalities in the United States in total), asking for information on the regulatory process, details of local land use regulations, and outcomes of the permitting and regulatory process. The survey itself builds on earlier work which surveyed a smaller number of municipalities (Mayer and Somerville, 2000), and other research has focused on surveys given to local officials and planners (Saks, 2008). We complement this survey approach through a direct measurement of housing regulations drawn from municipal regulations. Relative to surveys, this has the advantage of being comprehensive, rather than being limited by low or biased survey response rates. Our approach is also scalable and easy to augment with new questions, while surveys are inherently limited to the set of questions which were asked and which respondents are willing to answer.⁴

The second strand of this literature includes wedge-based approaches, which instead aim to impute housing regulations by examining the expected spatial macroeconomic

³Replication code, which can be adapted to other use cases, can be found at: <https://github.com/dmilo75/ai-zoning>.

⁴There is also, obviously, information captured by surveys that cannot be captured in the text of municipal zoning codes. For example, perhaps certain aspects of the zoning code are never or rarely enforced, or perhaps the zoning commission never approves particular kinds of projects even if they are legally permissible.

distortions resulting from zoning. Examples in this literature [Hsieh and Moretti \(2019\)](#), [Glaeser et al. \(2005\)](#), [Herkenhoff et al. \(2018\)](#), and [Duranton and Puga \(2019\)](#). [Babalievsky et al. \(2021\)](#) apply a similar production function based approach to impute the impact of commercial zoning impacts.

Third, other national approaches have examined textual data, but in more limited ways. [Ganong and Shoag \(2017\)](#) focus a scaled count of judicial decisions on “land use.” While this is surely a proxy for regulatory strictness, it leaves open the question of precisely which housing regulations are driving housing litigation. In a similar spirit, [Stacy et al. \(2023\)](#) use machine learning tools to identify newspaper articles discussing changes to zoning restrictions in eight metropolitan areas and classify them as either loosening and tightening zoning restrictions and then analyze the effects of these changes in regulation on housing supply and rents. Our approach, by contrast, is able to establish more cleanly the precise nature of housing regulations across a broad sample of jurisdictions in the United States.

Another literature has attempted to address the limitations in national-level approaches through more detailed analysis of specific regulations at the state level. Most prominent is the approach by the Pioneer Institute, which has engaged in explicit classification of zoning rules for 187 municipalities in the state of Massachusetts. Prior work by [Glaeser and Ward \(2009\)](#) establishes that regulatory intensity measured in this dataset does indeed associate with higher costs and lower construction. [Gyourko et al. \(2008\)](#) mention both the importance of this kind of detailed local analysis, as well as the challenges in scaling this approach to the national level:

“The proliferation of barriers and hurdles to development has made the local regulatory environment so complex that it is now virtually impossible to describe or map in its entirety. [Glaeser et al. \(2006\)](#) come closest to doing so. For a subset of the Boston metropolitan area, they conducted a detailed analysis of local zoning codes, permitting precise calculations of potential housing supply across communities. However, the enor-

mity of that effort prevents it from being replicated in other markets by a single research team.”

We argue that the practical difficulties behind the scaling up of this approach have now been addressed through the development of modern AI LLMs, providing both the granularity of the state-based approaches along with the scale of the national regulatory studies. Indeed, the Pioneer Institute data—the most comprehensive of these state based approaches—is a crucial test for our approach. We begin our analysis by first analyzing data in Massachusetts using the same data source for municipal documents identified by the Pioneer Institute team, which allows for a cross-validation of the accuracy of our AI-led approach against the existing housing regulation classification. This serves as an important validation check of our approach. Other detailed state-level analyses of housing regulation include [Shanks \(2021\)](#) which also focuses on Massachusetts and uses Machine Learning tools (Latent Dirichlet Allocation). California has also been the subject of detailed and specific analysis, focusing in particular on growth limitations ([Quigley and Raphael, 2005](#); [Jackson, 2016](#)), as has Florida ([Ihlanfeldt, 2007](#)). The closest paper to ours in terms of measurement is [Mleczo and Desmond \(2023\)](#), which uses Natural Language Processing methods applied to municipalities in the San Francisco and Houston areas. We differ in our use of methods, employing LLMs rather than filtering and keyword matching, enabling greater accuracy across a greater list of regulatory variables.

These studies leave important gaps in our understanding of housing regulations under both the national and state-level analyses. While the national approaches establish that housing regulations appear to drive important variation across the country in housing costs and construction activity, they have less to say about which specific regulations are the key drivers. Isolating specific regulatory impacts is essential for policy seeking to remedy possible impacts of regulatory driven housing cost increases. Alternatively, more detailed state-level data offers the potential to isolate the specific aspects of housing regulation that are most binding. These approaches, however, are limited in their geo-

graphic scope outside the unique states of Massachusetts, California, and Florida. Consequently, the extent to which specific housing regulations drive costs and construction activity across the country are unclear. Both line of research are also not able to contrast costs with potential benefits or amenities, making it impossible to disentangle supply and demand side effects which are crucial to establishing the cost-benefit tradeoffs of housing regulation.

Relative to this literature, our contribution is to construct a more comprehensive and detailed measure of how zoning regulations and building codes vary across the United States. We provide the most detailed assessment to date of all relevant housing regulations (i.e., minimum lot sizes, whether multifamily apartments can be constructed, inclusionary zoning mandates, setback rules, etc.) that apply to construction in local areas.

Additionally, we also contribute to the literature by testing the accuracy and usefulness of LLMs in creating novel regulatory and policy datasets. Existing research on AI models emphasizes both their promise in analyzing textual data ([Zhao et al., 2023](#)), as well as challenges with undesirable AI features such as “hallucination” and manufactured model output ([Azamfirei et al., 2023](#)). Verifying whether LLMs can accurately parse large legal documents—and for which questions—is therefore a crucial step towards our understanding of the capacities of these models, with the promise of opening up the large-scale use of textual documents for quantitative research. A broader contribution of our project is therefore a large-scale application of large language models to a complex regulatory and policy dataset generation task. This serves as a critical test case for the efficacy and reliability of LLMs in not only understanding and processing complex legal and regulatory language but also in discovering and extracting novel, actionable insights from a vast array of documents. Prior literature has used textual data to extract information, particularly sentiment, from text ([Hassan et al., 2019](#); [Romer and Romer, 2004](#); [Tetlock, 2007](#); [Lopez-Lira and Tang, 2023](#)); a few papers have begun to use LLMs for generative data purposes in existing textual, financial, and regulatory documents ([Giesecke, 2023](#);

Jha et al., 2023; Yang, 2023; Bybee, 2023; Hansen and Kazinnik, 2023). Hoffman and Arbel (2023) argues for the use of LLMs in “generative interpretation” in estimating the meaning of legal contracts.

The central contribution of this project is to establish a solid groundwork for the ongoing application and advancement of large language models (LLMs) in the field of legal and regulatory research. By demonstrating how these advanced models can optimize data generation, improve information accessibility, and facilitate predictive analysis, we argue for incorporating LLMs into the wider research, regulatory, and policy ecosystem.

2 Construction of National Housing Regulatory Database

2.1 Municipal Codes and Zoning

In the United States, local governments are “creatures of the state” subordinate to state control. Municipal corporations are authorized, subject to state law, to organize local government, and refer to cities, towns, villages, and other government units which function in that capacity. This concept largely overlaps with the Census definition of “incorporated place” which we use to organize our analysis.⁵

In most states, one of the powers granted to municipalities by the state government is control over local zoning decisions; indeed, the desire to control local zoning is a common reason to incorporate in the first place. Zoning, broadly, consists of two key sets of regulations: land use regulations, which partition local land into distinct use classes, and bulk regulations, which restrict the density of buildings in different land use classes. Examples of bulk regulations include: coverage, setbacks, height restrictions, and floor area ratio caps. Other mandates and requirements, such as parking minimums, further constrain both commercial and residential development in different areas.⁶

⁵In several states the “Township” form of government also has jurisdiction in zoning which aligns with the Census County Subdivision definition.

⁶States and municipalities also enact building codes, which govern the building and safety standards

Municipalities enforce laws by issuing municipal codes which outline local regulation in different domains. Zoning codes outline permitted uses for different classes of land as well as relevant housing regulations. Some regulations apply broadly to all land within a jurisdiction; other regulations (such as minimum lot sizes) typically vary depending on the specific use class and district (i.e., single-family zoning, commonly referred to as R-1, or commercial or industrial). These ordinances are typically updated over time to reflect changes in local regulations, and are aggregated by different companies. Table 1 illustrates the breadth of our sample coverage. In total, we cover 25% of all municipalities in the US and 6% of all townships. This coverage is skewed to larger cities, and so of the 76% of of the population in the US that live in either a municipality or a township, we cover 63%. Panel B shows our underlying sources for the municipal codes in our sample. American Legal Publishing provides significant numbers of records in the Northeast and Midwest, Municode provides especially good coverage in the South as well as in the Midwest, and [Ordinance.com](https://www.ordinance.com) provides substantial coverage of the West and Northeast.

The primary dataset for our analysis consists of the full-text of zoning documents. At the municipality-level, we also draw on information on building permits data from the Census Building Permits Survey. We also connect to rent and price data drawn from the American Community Survey (ACS) at the municipality level.

2.2 Large Language Models

Large Language Models (LLMs) are a form of artificial intelligence that primarily handle sequential data such as sequences of words in textual data. LLMs are based on the deep learning “transformer” architecture as introduced in [Vaswani et al. \(2017\)](#). The key innovation is the “attention mechanism,” enabling the model to focus on multiple words of the input text at once. This helps the model understand words in context, such as sentences or paragraphs. Transformers also represents a significant advancement in terms of

that new construction needs to adhere to.

both accuracy and runtime over previous models like Recurrent Neural Networks, which processed sequences linearly. LLMs are trained with semi-supervised learning, first pre-training the model on a large corpus of text and subsequently fine-tuning the model with human feedback. After training, LLMs can generate human-like text, answer questions, summarize text, and generalize from their training to perform tasks they were never explicitly trained for, a concept known as zero-shot learning. This means the model does not need as an input explicit examples of additional training to perform well in an out-of-sample exercise, a key advantage we use in our analysis.

LLMs have several advantages and disadvantages relevant for our setting in applying to housing regulatory textual analysis. The central advantage is scalability: we are able to load large quantities of municipal code data for classification and analysis, which far exceeds the capacity of any human team to analyze. Other advantages include the prospect for additional training, allowing for increased accuracy over time as LLMs improve in accuracy and additional training data is incorporated into the analysis.

Potential drawbacks in using LLMs for this purpose center on the inaccuracy of measurement and classification. This can happen either through limitations in the context window used to identify relevant text from the sample corpus, or the content and lack thereof of similar questions and related texts in the underlying training sample. Legal interpretation requires many assumptions and nuances, and even though LLMs are likely exposed to legal interpretation in their training, they may need to be reprompted on them to ensure greater focus for the questions at hand. Even current state-of-the-art LLMs may inadvertently produce incorrect information, produce information with an incorrect degree of certitude, and potentially manufacture data output (“hallucination”). Possible biases in the responses are linked to the quality of training data and the prompting and multi-step processing steps, and so measurement error may or may not be classical depending on the explanatory variable of interest. Finally, relevant information to answer zoning regulation questions may be outside the domain (i.e., in the form of state regula-

tion not contained within our ordinance sample). We attempt to measure these drawbacks through comparison of LLM-generated output against human defined categorizations of regulation.

2.3 Processing Municipal Codes Using LLMs

To conduct our analysis, we use a standard framework known as “retrieval-augmented generation” (Lewis et al., 2020). The basic objective of this approach is to combine a large pre-trained language model with external information retrieval, in order to give the LLM the ability to “look up” information from a vast corpus of text during the generation process. We outline our general procedure in Figure 1.

The first step of our process is to download and scrape the sources of municipal codes listed in Table 1, which provides us with a large corpus of zoning documents relevant for our analysis. These municipal codes contain detailed housing and zoning regulations relevant for our study, and we filter out ordinances which do not contain zoning information by searching for key phrases, like common table headers (i.e. “Table of Uses”) or zoning district names (i.e., R-1 for the first residential zoning district). We scrape each section within an ordinance separately, and partition sections so that they contain between 50 and one thousand tokens of text.⁷ Any images in the tables are transcribed using Amazon Textract. We then use text embeddings, which are vector representations of the text’s semantic meaning. This enables efficient search through zoning documents. The basic intuition behind embedding is to represent words with vectors which represent a dimension in embedding space, such that words with similar semantic meaning are closer in this space. For our zoning document, this ensures that we are able to retrieve components of the document relevant for our specific questions. Different embedding algorithms conduct this task in distinct ways; we use the `text-embedding-ada-002` algorithm from OpenAI for the national sample, and a newer algorithm `text-embedding-3-large`

⁷We use the OpenAI tokenizer where one token is roughly four characters of text.

for the testing sample comparison with Pioneer.

We similarly embed the questions we want answered from the documents, which for ease of comparison we limit to the question base already used by the Pioneer Institute (i.e., “Is multifamily zoning allowed in this area as-of-right?”). We rephrase these questions from the original wording provided by the Pioneer Institute in order to produce a more simplified version which is easier for the LLM to parse. This primarily consists of breaking down compound questions.

With two separate embedded vectors in hand, the zoning documents from a particular municipality and a question we would like answered, we then isolate the parts of an ordinance most relevant to answer the question. The length of typical zoning documents exceeds the context windows currently usable by LLMs, so we need to select specific sections of text that are most likely to be relevant to the question. We use cosine similarity, a standard measure of distance between two vectors, to rank sections of text by how likely they are to be relevant to the question. We then refine this ranking by using a cross-encoder reranking model⁸ on the top 50 sections of text, which processes the question and section text pairs simultaneously to determine the most semantically similar sections.⁹ We then select text to show the LLM in order of highest relevance until a threshold of four thousand tokens is reached.

We include three key pieces of information to provide the LLMs. First, we include 4,000 tokens of relevant text to the LLMs. Second, we provide rephrased zoning question, as described above to simplify model parsing. Third, we also provide additional background information and assumptions. The background information and model assumptions were taken directly from the Pioneer study (their “Issue Overview” and “Research Coding” sections for each question) and were based on trial and error for what information was most relevant to improve model performance. Appendix ?? contains

⁸We specifically use the Cohere reranking model for this step.

⁹When double checking answers on select questions we instead use keyword inclusion to re-rank section text.

full information the original Pioneer questions, our rephrased questions, as well as the additional background information and assumptions provided.

All three pieces of information are provided in a single call to the LLM, in order to produce model output which is our answer. In many cases, to answer a specific question, we chain together multiple calls. Some pieces of information are queried prior to asking the question, which are called subtasks, to provide pre-processing or background research. For instance, when asking about the largest frontage requirement for all single family residential districts, we first ask the LLM to name all districts which allow single family housing. We do this as a separate step because the relevant text defining allowable uses in a district, and the text defining frontage requirements for districts are typically in different sections of the ordinance under different embedding vectors. Additionally, LLM performance is enhanced when it is only required to answer a direct single step question in each call. Finally, we provide a “system prompt” where we tell the LLM that it is a municipal zoning expert, detail what the structure of the prompts for particular questions will be, and tell the LLM to think ‘step by step’ to induce chain of thought reasoning.

We also engage in post-processing of certain questions, which functions to double-check answers. For instance, an affirmative “Yes” to a question about whether townhouses/attached housing is allowed typically means the LLM has likely found affirmative evidence that such housing typologies are allowed, while an answer of “No” signifies either a lack of approval, or a lack of sufficient context for the LLM to answer the question. In such cases where an answer could indicate lack of information, we reprompt the LLM and directly use keywords like “townhouse” or “attached” to refine and rerank our search (instead of the reranking algorithm).

The key takeaway from our approach towards generative regulatory parsing is that, at least with models available at time of writing, model accuracy improves substantially above simple “zero shot learning” examples given additional human input. We provide substantial human input in the areas of prompt engineering and providing background

information as well as assumptions, which helps to focus the LLM on the relevant focus of the text. Additionally, we design a multi-step reasoning chain for each question to simplify the tasks required of the LLM in each sub-step. Such additional human processing is likely necessary in other contexts as well, at least until further advances in LLMs are made.

2.4 Model Validation with Pioneer Data

Performance analysis is a crucial step in validating the effectiveness and reliability of LLMs for tasks such as zoning ordinance interpretation. By comparing the accuracy of different LLM approaches against a ground truth dataset, we can assess their ability to provide consistent and correct answers to zoning-related questions. This analysis helps identify the strengths and weaknesses of each model, as well as any discrepancies between the model outputs and the reference data.

To do so requires a high-quality reference dataset. The Pioneer dataset serves as an excellent starting dataset for our purposes, as previously mentioned, due to the expert classification of a large number of municipalities. The main drawback in using this dataset is the staleness of responses—with responses categorized as of 2004. Many regulations have changed in the intervening twenty years, and we have access only to the most recent zoning ordinances, not the ones that prevailed in that time period. Additionally, the Pioneer Institute relied on some outside information (i.e., directly contacting local regulatory bodies) in addition to the text. To address these issues, we construct a testing dataset based on 30 randomly chosen municipalities from the Pioneer Institute dataset, and 1) exclude question responses which relied on outside context, and 2) correct inaccuracies in the original classification.¹⁰

¹⁰Due to the time-intensive nature of the expert correction step, we only check responses in which our LLM approach disagrees with the Pioneer Institute classification. This means that we potentially overstate model accuracy in cases in which the LLM agrees with the Pioneer Institute original classification; but that original classification was wrong. We are currently expanding our error-correction process to adjust for these cases as well.

Table 2 shows the performance results of our baseline Chat GPT-4 Turbo model against the testing sample in Massachusetts. Among continuous questions, our generated answers have an overall correlation of 0.67 against the ground truth of expert classifications, after winsorization of our model at the 1% level and corrections of errors in the Pioneer sample. This represents a quite high benchmark, and also incorporates substantial heterogeneity. When asking about the number of zoning districts in the municipality, we obtain a correlation of 0.98. When asking about the minimum of residential min lot sizes (i.e., the lot size requirement for R-1 zoned single family homes, an important zoning question determining allowable density), we find a quite high 0.92 correlation. These results suggest we are able to reach quite high model performance when matching against continuous numerical outcomes.

We find even higher model accuracy when measuring binary questions (i.e., those with a yes or no answer like whether “multi-family housing is allowed” which we measure perfectly across all municipalities). There, we observe a model accuracy of 96% across all questions. Because the raw accuracy measure may be biased depending on the base rate of answers, we also provide a Relative Squared Error (RSE) which compares each model’s result compared to a naive model which guesses the sample model. We observe quite small RSEs as well.

In Figure 2, we visualize the average results across questions in Table 2. In dark blue, we plot the percent correct for each model using the percent accuracy for binary variables, the correlation for continuous variables, and adjusted percent correct for categorical questions. We also plot the frequency each model says “I don’t know” in grey, which varies across each model and question type. Finally, we attribute the remainder as the incorrect percent for each model (shown in light blue).

2.5 Heterogeneity Across Models

While our benchmark results appear quite accurate, we also contrast them with estimates drawn from other models. This analysis helps identify the strengths and weaknesses of each model, as well as any discrepancies between the model outputs and the reference data. Furthermore, performance analysis allows researchers to make informed decisions about which LLM is best suited for their specific use case and to identify areas for improvement in the models' knowledge and reasoning capabilities.

In Figure 2, we contrast model performance across GPT-4 Turbo (the benchmark model), Claude 3 Opus, and GPT-3.5 Turbo. For binary questions, we find that GPT-4 Turbo is the highest performer, followed by Claude 3 and then by GPT-3.5 Turbo (which has an accurate rate of around 80% for binary questions).¹¹ Interestingly, this model order is not preserved in continuous questions, for which we actually observe the highest model performance in GPT-3.5 Turbo, followed by GPT-4 and Claude 3. However, this difference is mostly driven by differences in performance on one question, the minimum lot size question, which can tend to have extreme outliers because of districts within jurisdictions with particularly large minimum lot sizes.

2.6 Understanding Model Errors

To better diagnose reasons for model error, in Figure 3 we provide a complete decomposition of all of the reasons for disagreement between GPT-4 Turbo and the original Pioneer Study on binary questions. We manually reviewed each question that Chat GPT-4 Turbo disagreed with the Pioneer Institute, and present the reasons for discrepancies in a figure. We outline, for each of the questions, the specific reason for disagreement: whether the pioneer study was itself outdated or inaccurate and subsequently corrected, whether

¹¹This performance may reflect fundamental features of GPT-4 Turbo versus Claude, but it could also reflect the fact that we fine-tuned our prompting and chaining strategies to optimize performance on Chat GPT-4 Turbo and it is possible that if we had instead fine-tuned to maximize performance on Claude 3 that Claude 3 would have performed better.

the LLM misinterpreted context (i.e., it was provided the correct information, and simply provided an inaccurate answer), whether the LLM missed the context, and whether the answer itself was coded as incorrect but the true classification appears somewhat ambiguous.

Largely, answers from the Pioneer Institute that our model did not match were due to changes in the underlying ordinance since the Pioneer Institute study roughly 20 years ago. LLMs missed the context in two cases, while in four cases the answer itself was ambiguous. The most important category for our purposes are cases in which the LLM misinterpreted the context—this happens in nine cases, most often with respect to whether townhouses are allowed and with permit caps or phasing. Six questions do not have this type of error happen at all. When considered over a large sample, these results appear promising in suggesting that errors are typically quite rare.

Importantly, the errors also appear balanced across false positives as well as false negatives. Table 3 provides a confusion matrix comparing our baseline GPT-4 Turbo model against the Pioneer classifications, separating true positives, false positives, true negatives, and false negatives. Our errors are equally represented among false positives as well as false negatives (six each), suggesting no obvious bias in our classification.

2.7 Comparison Against Wharton Regulatory Index

To further validate our results, we compare our answers to another commonly used dataset of national housing regulation: the Wharton Index of [Gyourko et al. \(2021\)](#). To do so, we scale up our generative regulatory measurement approach up to the national level, asking the same set of questions in the Pioneer Institute data for a large sample of national municipalities.

In Panel A of Table 4, we first compare our questions with the Wharton approach on two questions which find overlap across the question bases: on affordable housing and minimum lot sizes. Unfortunately, there are small nuances which do not permit

a completely clean comparison. We use the Pioneer Institute wording which classifies both mandates and incentives as constituting affordable housing (question 17), while the Wharton study only considers affordable housing mandates (question 9a). For minimum lot sizes, we currently consider minimum lot sizes across all districts, while the Wharton study (question 7b) only considers residential districts, and categorizes these into four bins (whereas we use the precise minimum lot size).

Despite these limitations, we find a sizable correlation between our measure of affordable housing and the one measured in the Wharton study of 0.36. We observe smaller, but still sizable correlations, between 0.11–0.29 when examining the minimum lot size questions.

3 Spatial Patterns of Housing Regulation and Development

3.1 The Role of Unincorporated Areas in Housing Production

The first basic distinction we draw in the data is across incorporated areas, which are regions governed by a municipal corporations, and unincorporated areas outside of the jurisdiction of such governments. Typically, local control over zoning decisions is a key factor behind the decision to incorporate. Unincorporated areas are still subject to housing regulations at the county and state levels, but typically face far lower housing regulations. Unincorporated regions are generally lower inhabited regions located either in the exurban fringes of cities or in more distant rural areas.

To illustrate the nature of municipal boundaries, we show in Figure 4 incorporation maps around four large and representative metropolitan areas: Philadelphia, San Francisco Bay Area, Atlanta, and Houston. We highlight incorporated areas in green and unincorporated areas in red, and shade the areas to reflect local density. We see that the dense coastal areas of San Francisco and Philadelphia have comparatively little unincorporated lands around their metro areas, while Houston and Atlanta have comparatively

more unincorporated land within the metropolitan area. Perhaps surprisingly, some of this land is fairly dense, reflecting built up activity, which is concentrated in the fringes of this unincorporated zone closest to the city itself.

We in Figure 5 the concentration of housing growth in unincorporated areas. Panel A of this figure shows the percentage change in housing units from 2000–2020 as a fraction of the land incorporated in 2000 at the census block level. Fully incorporated areas had zero net changes in housing growth over this period; the entirety of housing growth in this period was in areas at least partially unincorporated. We show in Panel B the amount of new housing units added specifically in incorporated and unincorporated areas, with 42% of housing production taking place in unincorporated areas, despite these areas having only 23% of the total population. This trend is especially true in the South, where a majority of housing production takes place in unincorporated regions.

The fact that a large fraction of housing development in the United States consists of greenfield construction in unincorporated regions is potentially surprising given their remote nature and the consequent distance from local job centers and local amenities. We show in Table 5 associations between incorporated and unincorporated areas in the U.S. Unincorporated areas have substantially cheaper homes in prices and rents; higher vacancy rates, and residents have longer commutes.

Several factors contribute to this phenomenon. Unincorporated areas often have less restrictive zoning regulations, making it easier and less costly for developers to build new housing. Land costs tend to be lower in these areas, further incentivizing development. Additionally, many incorporated municipalities have adopted growth control measures that limit new construction, pushing development to the fringes.

While supply constraints push development to the the fringes of American cities, we show that this is in contrast to actual demand patterns. House prices and rents are lower in unincorporated areas, which are also typically far from local jobs and amenities. Commute lengths tend to be higher, especially very long commute times above an hour.

This pattern of development has significant implications for urban form and sustainability. While it allows for continued housing production in the face of restrictive zoning in incorporated areas, it also contributes to urban sprawl. This can lead to increased infrastructure costs, as new developments in unincorporated areas often require extensions of roads, utilities, and public services. Residents also frequently face higher commuting lengths in traveling to jobs, limiting employment opportunities and generating environmental pollutants. These factors broadly suggest that development frictions in incorporated areas push housing supply to greenfield development in unincorporated areas, motivating our first fact.

Fact 1. *Housing development in the United States is concentrated in unincorporated areas.*

3.2 Density Restrictions and Their Impact on Urban Form

Next, we turn to the nature of regulations within incorporated regions. We focus first on density regulations, which directly limit the ability for infill development. In Panel B of Table 9, we highlight a number of density restrictions on multifamily apartments across the United States. 5% of municipalities prohibit multi-family housing entirely, while 38% ban mixed use developments (apartments above commercial units). An overwhelming majority of jurisdictions (86%) limit multi-family housing as a conversion from single family or non-residential buildings. These estimates are typically lower when weighting by population, suggesting fewer constraints in more populated areas, but suggest broad hostility to apartment construction across a substantial part of the United States.

Further restrictions in land use restrictions regarding where dense housing can be built. To understand this set of regulations, we turn to land use zoning maps, which show allowable densities and housing typologies in different municipalities. We collect zoning maps for a sample of 31 large municipalities covering 18.2 million people, including Chicago, Seattle, Kansas City, Detroit, San Francisco, Austin, San Antonio, Tampa, Los Angeles, San Diego. Our results, shown in Figure 6, highlight limitations on allow-

able densities even in some of the nation's largest cities. 36% of land area is zoning only for single-family zoning, while 41% of land area is zoned for single-family or duplex, and multifamily apartments are allowed in only 31% of land area.

Density restrictions also apply to single-family housing units through bulk regulations which limit lot and building size. Figure 10 shows the distribution of four different housing regulations across the US: number of zoning districts, largest frontage requirement, mean minimum lot size (across all zoning districts), and minimum minimum lot sizes (across all zoning districts). The figure shows that these regulations vary substantially. For example, a large mass of municipalities has no minimum lot size requirement at all, while a non-trivial share of municipalities have minimum minimum lot sizes in excess of ten thousand square feet.

We show more precise distributions in Appendix Figure A2. Among single-family zoned areas in municipalities across the entire country, 66% have minimum lot size requirements above 5,000 square feet, 17% of requirements are above 10,000 square feet, and 7% have requirements exceeding half an acre. In contrast, the average size of a new home built in 2023 was 2,411 square feet, indicating that many jurisdictions have minimum lot size requirements considerably larger than typical of new housing construction across the United States. We summarize these various restrictions on density as our second fact.

Fact 2. *Housing density is limited across the United States through regulations on multifamily housing and small lot single-family housing.*

We next interpret municipal regulations in the context of the monocentric city (Alonso, 1964; Mills, 1967; Muth, 1971).

We show various regulatory variables along the dimension of distance to city center in Figure 7. Affordable housing mandates are decreasing in distance from the center of the city, illustrating that these regulations are most commonly found at the centers of cities. Minimum lot size requirements show a different pattern, and vary markedly across re-

gions. In much of the country outside the northeast, these regulations are fairly flat across the distance to city center. In some areas, these regulations are more commonly found in closer suburban regions to cities, while becoming less strict further away. The northeast is a marked exception to this pattern, featuring an increasing relationship between these regulations and distance to the city center.

To further illustrate these patterns at the metropolitan level, Figures 11 and 12 show maps of minimum lot sizes and affordable housing incentives, respectively, for jurisdictions within the metropolitan areas surrounding four select cities in the United States, San Francisco, Chicago, Atlanta, and Boston.

These graphs document substantial variation in both minimum lot sizes and affordable housing mandates and incentives within metropolitan areas across municipalities, with the central city and inner suburbs having lower minimum lot sizes and higher rates of affordable housing mandates than in jurisdictions farther from the central city. This figure illustrates a key advantage of our approach: the ability to construct measures of zoning ordinances at the level of the municipality across a wide variety of municipalities and regions in the United States. Appendix Figure A3 also shows a heatmap of correlations between regulations at the national level.

We also show these associations in Table 6. The number of zoning districts decreases robustly across distance to city center. Some components of allowable density decrease with distance, especially whether townhouses and mixed-use development is allowed. Unincorporated areas also increase with distance from city center, as shown in Appendix Figure A1.

Fact 3. *Housing regulation increases in land gradients away from city centers. The main exception to this pattern is the Northeast.*

4 Characterizing Housing Regulations

4.1 Principal Component Analysis of Housing Regulatory Dataset

We also attempt to construct a nationwide index of our questions to better benchmark against the Wharton study. We focus on a PCA analysis to ensure greater comparability with the Wharton Index, which engages in dimension reduction across questions to provide an omnibus index consisting of the first principal component across sub-indices that group similar questions together. We similarly examine a principal components of all of our questions at the national level, finding two key principal components which appear to drive the bulk of the cross-sectional variation in zoning answers. Appendix Table ?? provides the loadings of each question on the two principal components, and Figure 8 map the two principal components across the nation.¹²

In Panel B of Table 4, we find positive correlations of both PCAs against the composite Wharton Index at the CBSA level. The first PC correlates at 0.28 against the Wharton Index, while the second PC correlates at 0.10. These findings suggest that our regulatory measures overlap somewhat with existing measures of regulation, providing some reassurance of basic fit, but also seem to provide somewhat distinct information as reflected in the correlation being less than one.

4.2 Regulatory Complexity and Housing Construction

Table 9 shows the association of housing regulations across income and urban categories across the United States. We observe, for instance, that affordable housing mandates are found much more often in higher income and urban areas.

To disentangle the relative roles of demand and supply in housing production, we show in Figure 9 the associations between building permits, median house prices, and

¹²Tables ?? and ?? highlight correlations of these principal components against housing market outcomes and socioeconomic determinants.

our two key principal components. Panel A highlights that areas with a high value for the first principal component generally have high house prices as well as construction, while areas low in this dimension typically have low prices as well as building activity.

This association suggests the role of housing demand and the first principal component of housing regulation. A plausible mechanism to connect these reflects the role of this set of housing regulations on various kinds of housing exactions which are only possible in high-demand environments. For instance, affordable housing mandates and incentives are policy tools which will be in greater demand in high-demand areas, and also impose additional costs which are also only able to be met by developers in high-demand environments. We show loadings on this factor in Appendix Table [A2](#). Another key variable which loads heavily on this regulatory dimension is the number of zoning districts, which proxies for the relative complexity of housing regulations, and for this reason we refer to the entire principal component as reflecting a “complexity” dimension of zoning.

We show associations of the first principal component in Table [7](#). We focus on column 1, which controls for MSA fixed effects and compares municipalities within metropolitan areas. Consistent with largely demand interpretation of this principal component: areas that are high on this dimension have a high college share, are higher in job density, lower poverty rates, have substantially higher shares of Democrats. These are high-demand areas which may be prone to extraction of value by local governments ([Diamond, 2017](#)).

To be sure, such regulations may also effect housing supply. We explore this variable in more detail in columns 5–6 of this table, additionally adding additional topographical and land availability controls such as the fraction of land developed in 2001, the squared fraction of land developed in 2001, and the fraction of land with a flat topography. We find these controls are essential to isolating the relevant elasticity explainable by regulatory factors, rather than background land availability, and after including them we find that areas higher in the housing complexity variable also feature lower housing supply elasticities.

Fact 4. *The first principal component of housing regulation captures a dimension of housing complexity which varies with housing demand. This factor captures aspects of extraction of housing value.*

4.3 The Role of Exclusionary Zoning

Exclusionary zoning refers to land use regulations that limit housing density and types, often with the effect of excluding lower-income residents from certain areas. These regulations typically include measures such as large minimum lot sizes, restrictions on multi-family housing, and other bulk regulations that effectively increase the cost of housing in a given area. Our analysis reveals that the second principal component of housing regulations identified in this study correlates strongly with these measures of exclusionary zoning, particularly minimum lot sizes and other bulk regulations that limit density, the loadings for which we show in Appendix Table A2.

The spatial distribution of exclusionary zoning practices is not uniform across the United States. Our findings indicate that these practices are more commonly found in suburban areas, with a particularly strong presence in the Northeast region. This pattern suggests a spatial segregation effect, where more affluent suburbs use zoning regulations to maintain socioeconomic homogeneity and limit the influx of lower-income residents.

We show associations of the second principal component in Table 7. We focus on column 1, which controls for MSA fixed effects. Areas characterized by high scores on our second principal component exhibit distinct socioeconomic patterns. These localities tend to have higher housing values, higher rents, and higher income levels. They also show a higher proportion of white residents, and lower poverty rates. Average math test scores are substantially higher, while the fraction of children eligible for free lunch is substantially lower. Density patterns favor single-family residents, with owner-occupied fractions substantially higher and housing unit density substantially lower. These correlations underscore the role of exclusionary zoning in maintaining economic and racial

segregation in residential patterns.

Fact 5. *The second principal component of housing regulation captures exclusionary zoning practices.*

5 Discussion and Framework

Having discussed the construction of a national housing regulation dataset and five key facts which arise from the data, we discuss in this section an interpretive framework to put these facts in context. The key question that arises is: how can we characterize the tools of housing regulation in the hands of municipalities, and what accounts for the observed spatial variation in these regulations?

The first principal component of our housing regulation dataset, corresponding to higher complexity, can be seen roughly as a set of regulations on the shadow price of construction. This is because regulations such as mandatory inclusionary zoning can be seen as an additional implicit tax on new development, with the proceeds either redistributed to other residents (as in the case of affordable housing units), or else extracted by the government for the purpose of public goods or private benefits. While these are not regulations explicitly on the price of development, they may be interpreted as a shadow price or cost on housing.

By contrast, the second principal component, associated with exclusionary zoning, primarily consists of regulations on the quantity of housing built in the form of various bulk regulations, particularly on minimum lot sizes. These regulations, as discussed above, effectively truncate the housing quality distribution on the left side of the distribution, reducing the number of lower-income residents able to live locally. We connect such exclusionary motives to educational sorting at local levels.

The complexity or extraction motive for housing regulation shows up most strongly in the centers of American cities, where demand for living is highest and so the ability

for local governments to extract value is correspondingly higher as well. By contrast, exclusionary regulations are commonly found in suburban areas in the peripheries of those cities across the United States, but are particularly pronounced in the Northeast.

This spatial regulatory pattern contributes to the sorting of Americans along the dimensions of age and income. Cities are home to poorer residents to access smaller housing, wealthy residents, as well as younger working households. By contrast, suburban areas are more typically home to young families with

The combination of these two sets of regulations combine to limit housing production and density across American cities, and thereby pushes out additional housing production to the far exurban fringes of cities where demand is low, but supply constraints are less binding.

Our results help to connect classic previous theories of zoning. [Fischel \(1987\)](#) and [Hamilton \(1975, 1976\)](#) argue that housing regulation and property taxes can create efficient public goods provision in the context of [Tiebout \(1956\)](#) sorting. By contrast, [Zodrow and Mieszkowski \(1986\)](#) argue instead that local taxes are distortionary and function like excise taxes, while more recent scholarship has emphasized the segregation motives of zoning, especially exclusionary zoning ([Rothstein, 2017](#); [Cui, 2024](#)). Our two sets of regulatory controls by municipalities spans this prior literature, and helps to explain the circumstances under which housing regulation can appear extractive, and the conditions under which they sustain local public goods investment under exclusionary environments.

6 Conclusion

This study makes significant progress in using large language models (LLMs) to accurately measure and analyze complex zoning regulations across a broad sample of U.S. municipalities. Our results demonstrate that state-of-the-art LLMs can achieve near-human

levels of accuracy in classifying zoning rules from textual documents, with accuracy levels of 96% for binary questions and correlations as high as 0.92 for continuous questions like minimum lots sizes. Our approach also correlates with existing measures of regulation from the Wharton Index. This generative regulatory measurement approach enables the creation of a comprehensive, nationwide dataset of municipal zoning regulations.

We make all collected data and the associated replication code publicly available. Our AI-driven approach is scalable, auditable, and allows for refinement as LLMs continue to advance. With further development, this generative regulatory measurement framework can be extended to building codes, regulations in other domains, across different countries and languages, and to other regulatory contexts.

We use the resulting dataset on housing regulations across the United States to establish five key facts about housing regulation. Our results point to a tradeoff between two typologies of housing regulation: one which is extractive in nature, and another which is instead primarily focused on exclusion within neighborhoods. These can be thought of as regulations on the shadow price and quantity of housing, respectively. Spatially, we find that extractive regulations are found in high-demand areas, such as the centers of cities, while exclusionary zoning practices are found in the suburban peripheries around city centers. These two sets of regulations combine to reduce allowable density within incorporated areas, pushing incremental development in the United States outside these jurisdictions into unincorporated regions.

References

- Alonso, William**, *Location and land use: Toward a general theory of land rent*, Harvard university press, 1964.
- Azamfirei, R, SR Kudchadkar, and J Fackler**, “Large language models and the perils of their hallucinations,” *Critical Care*, 2023, 27 (1), 1–2.
- Babalievsky, F, K Herkenhoff, LE Ohanian, and EC Prescott**, “The Sky is Not the Limit: The Macroeconomic Consequences of Commercial Zoning Laws,” Technical Report 2021.
- Baum-Snow, Nathaniel and Lu Han**, “The microgeography of housing supply,” *Journal of Political Economy*, 2024, 132 (6), 1897–1946.
- Bybee, J Leland**, “The ghost in the machine: Generating beliefs with large language models,” Technical Report, Technical report, Working Paper 2023.
- Cui, Tianfang**, “Did Race Fence Off The American City? The Great Migration and the Evolution of Exclusionary Zoning,” Technical Report 2024.
- Diamond, Rebecca**, “Housing supply elasticity and rent extraction by state and local governments,” *American Economic Journal: Economic Policy*, 2017, 9 (1), 74–111.
- Duranton, Gilles and Diego Puga**, “Urban growth and its aggregate implications,” Technical Report, National Bureau of Economic Research 2019.
- Fischel, William A**, *The economics of zoning laws: A property rights approach to American land use controls*, JHU Press, 1987.
- , *The homevoter hypothesis: How home values influence local government taxation, school finance, and land-use policies*, Harvard University Press Cambridge, MA, 2002.

Ganong, Peter and Daniel Shoag, “Why has regional income convergence in the US declined?,” *Journal of Urban Economics*, 2017, 102, 76–90.

Giesecke, Oliver, “State Government Debt Structure,” *Available at SSRN*, 2023.

Glaeser, Edward and Joseph Gyourko, “The economic implications of housing supply,” *Journal of Economic Perspectives*, 2018, 32 (1), 3–30.

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

– , **Joseph Gyourko, and Raven Saks**, “Why is Manhattan so expensive? Regulation and the rise in housing prices,” *The Journal of Law and Economics*, 2005, 48 (2), 331–369.

Goolsbee, Austan and Chad Syverson, “The Strange and Awful Path of Productivity in the US Construction Sector,” Technical Report, National Bureau of Economic Research 2023.

Gyourko, Joseph, Albert Saiz, and Anita Summers, “A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index,” *Urban Studies*, 2008, 45 (3), 693–729.

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

Hamilton, Bruce W, “Zoning and property taxation in a system of local governments,” *Urban Studies*, 1975, 12 (2), 205–211.

– , “Capitalization of intrajurisdictional differences in local tax prices,” *The American Economic Review*, 1976, 66 (5), 743–753.

Hansen, Anne Lundgaard and Sophia Kazinnik, “Can ChatGPT Decipher FedSpeak?,” *Available at SSRN*, 2023.

- Hassan, Tarek A, Stephan Hollander, Laurence Van Lent, and Ahmed Tahoun**, “Firm-level political risk: Measurement and effects,” *The Quarterly Journal of Economics*, 2019, 134 (4), 2135–2202.
- Herkenhoff, Kyle F, Lee E Ohanian, and Edward C Prescott**, “Tarnishing the golden and empire states: Land-use restrictions and the US economic slowdown,” *Journal of Monetary Economics*, 2018, 93, 89–109.
- Hoffman, David A and Yonathan A Arbel**, “Generative interpretation,” *Available at SSRN*, 2023.
- Hsieh, Chang-Tai and Enrico Moretti**, “Housing constraints and spatial misallocation,” *American Economic Journal: Macroeconomics*, 2019, 11 (2), 1–39.
- Huang, Haifang and Yao Tang**, “Residential land use regulation and the US housing price cycle between 2000 and 2009,” *Journal of Urban Economics*, 2012, 71 (1), 93–99.
- Ihlanfeldt, Keith R**, “The effect of land use regulation on housing and land prices,” *Journal of Urban Economics*, 2007, 61 (3), 420–435.
- Jackson, Kristoffer**, “Do land use regulations stifle residential development? Evidence from California cities,” *Journal of Urban Economics*, 2016, 91, 45–56.
- Jha, Manish, Jialin Qian, Michael Weber, and Baozhong Yang**, “Chatgpt and corporate policies,” *Chicago Booth Research Paper*, 2023, (23-15).
- Lewis, Patrick, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela**, “Retrieval-augmented generation for knowledge-intensive nlp tasks,” *Advances in Neural Information Processing Systems*, 2020, 33, 9459–9474.

- Lopez-Lira, Alejandro and Yuehua Tang**, “Can chatgpt forecast stock price movements? return predictability and large language models,” *arXiv preprint arXiv:2304.07619*, 2023.
- Mayer, Christopher J and C Tsurriel Somerville**, “Land use regulation and new construction,” *Regional Science and Urban Economics*, 2000, 30 (6), 639–662.
- Mills, Edwin S**, “An aggregative model of resource allocation in a metropolitan area,” *The American Economic Review*, 1967, 57 (2), 197–210.
- Mleczko, Matthew and Matthew Desmond**, “Using natural language processing to construct a National Zoning and Land Use Database,” *Urban Studies*, 2023, 60 (13), 2564–2584.
- Muth, Richard F**, “The derived demand for urban residential land,” *Urban studies*, 1971, 8 (3), 243–254.
- Quigley, John M and Steven Raphael**, “Regulation and the high cost of housing in California,” *American Economic Review*, 2005, 95 (2), 323–328.
- Rognlie, Matthew**, “Deciphering the fall and rise in the net capital share: accumulation or scarcity?,” *Brookings Papers on Economic Activity*, 2016, 2015 (1), 1–69.
- Romer, Christina D and David H Romer**, “A new measure of monetary shocks: Derivation and implications,” *American economic review*, 2004, 94 (4), 1055–1084.
- Rothstein, Richard**, *The color of law: A forgotten history of how our government segregated America*, Liveright Publishing, 2017.
- Saks, Raven E**, “Job creation and housing construction: Constraints on metropolitan area employment growth,” *Journal of Urban Economics*, 2008, 64 (1), 178–195.
- Shanks, Brendan**, “Land Use Regulations and Housing Development,” 2021. manuscript.
- Singla, Shikhar**, “Regulatory costs and market power,” Technical Report 2023.

- Stacy, C, C Davis, YS Freemark, L Lo, G MacDonald, V Zheng, and R Pendall**, “Land-use reforms and housing costs: Does allowing for increased density lead to greater affordability?,” *Urban Studies*, 2023. 00420980231159500.
- Tetlock, Paul C**, “Giving content to investor sentiment: The role of media in the stock market,” *The Journal of finance*, 2007, 62 (3), 1139–1168.
- Tiebout, Charles M**, “A pure theory of local expenditures,” *Journal of political economy*, 1956, 64 (5), 416–424.
- Vaswani, Ashish, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin**, “Attention is all you need,” *Advances in neural information processing systems*, 2017, 30.
- Yang, Stephen**, “Predictive Patentomics: Forecasting Innovation Success and Valuation with ChatGPT,” *Available at SSRN 4482536*, 2023.
- Zhao, Wayne Xin, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong et al.**, “A survey of large language models,” *arXiv preprint arXiv:2303.18223*, 2023.
- Zodrow, George R and Peter Mieszkowski**, “Pigou, Tiebout, property taxation, and the underprovision of local public goods,” *Journal of urban economics*, 1986, 19 (3), 356–370.

Tables

Table 1: Sample Coverage Metrics

Panel A: Sample and Local Government Coverage Metrics					
Coverage Metrics	National	Northeast	Midwest	South	West
Total Munis	19,488	2,101	8,481	6,587	2,319
% of Munis In Sample	25	32	19	22	48
Total Townships	16,213	4,111	12,102	0	0
% of Townships In Sample	6	23	0	-	-
Total Pop. (Millions)	331	57	69	127	77
% of Pop. Under Local Gov.	76	100	95	55	78

Panel B: % of Pop. Under Local Gov. Covered By Sample					
Ordinance Aggregator	National	Northeast	Midwest	South	West
American Legal Publishing	11	15	15	6	8
Municode	23	1	19	54	12
Ordinance.com	30	52	12	1	60
Total	63	68	46	61	80

Note: For local governments available in multiple datasets, we prioritize using Ordinance.com and then Municode and reflect that in the population count. We also adjust for geographical overlap between townships and municipalities in tallying population by using census block level population data and corresponding shape files. We use population estimates from the 2022 Census of Governments for municipality population, 2022 State-Level Census Population Data for census region and national population, and 2022 MSA-Level Census Population for MSA population.

Links to data sources are [American Legal Publishing](#), [Municode](#), and [Ordinance.com](#).

Table 2: Performance Results of Chat GPT-4 Turbo on Testing Sample of 30 Municipalities

Panel A: Continuous Questions

Question	RSE	Correlation
How many zoning districts, including overlays, are in the municipality?	0.06	0.98
What is the longest frontage requirement for single family residential development in any district?	1.16	0.70
Minimum of Min Lot Sizes (Square Feet)	0.73	0.61
Mean of Min Lot Sizes (Square Feet)	14.77	0.39
Minimum of Residential Min Lot Sizes (Square Feet)	0.16	0.92
Mean of Residential Min Lot Sizes (Square Feet)	11.80	0.44
Cumulative Average	4.78	0.67
Cumulative Median	1.16	0.67

Note: We calculate performance metrics and sample means (for RSE) only on the set of question municipality pairs that Chat GPT-4 Turbo does not say "I don't know". For Relative Squared Error we compare the model's results to the naive model that guesses the sample mean. The correlation column is the correlation between the model answer and the Pioneer Institute answer. We winsorize data from our models at the 1% level but do not winsorize data from the Pioneer Institute. The Cumulative Average and Cumulative Median are calculated across questions giving equal weight to each question.

Panel B: Binary Questions

Question	RSE	% Accuracy
Is multi-family housing allowed, either by right or special permit (including through overlays or cluster zoning)?	0.00	100%
Are apartments above commercial (mixed use) allowed in any district?	0.07	96%
Is multi-family housing listed as allowed through conversion (of either single family houses or non residential buildings)?	0.08	96%
Are attached single family houses (townhouses, 3+ units) listed as an allowed use (by right or special permit)?	0.30	90%
Does zoning include any provisions for housing that is restricted by age?	0.14	96%
Are accessory or in-law apartments allowed (by right or special permit) in any district?	0.09	96%
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?	0.00	100%
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by special permit?	0.00	100%
Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?	0.00	100%
Is there a town-wide annual or biannual cap on residential permits issued, and/or is project phasing required?	0.33	90%
Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?	0.14	96%
Cumulative Average	0.11	96%
Cumulative Median	0.09	96%

Note: For Relative Squared Error we compare each model's results to the naive model that guesses the sample mode. The accuracy column is calculated as the percent of municipalities that the model matches the adjusted Pioneer Institute answer for each question.

Table 3: Confusion Matrix For Chat GPT-4 Turbo

Question	True Positive	False Positive	True Negative	False Negative	True Positive Rate	False Positive Rate	Precision
Multifamily Allowed	28	0	2	0	1.00	0.00	1.00
Mixed-Use Buildings	15	0	14	1	0.94	0.00	1.00
Conversion To Multifamily	12	1	17	0	1.00	0.06	0.92
Townhouses Allowed	18	1	9	2	0.90	0.10	0.95
Age-Restricted Provisions	22	0	7	1	0.96	0.00	1.00
Accessory Apartments Allowed	18	0	11	1	0.95	0.00	1.00
Flexible Zoning By Right	1	1	27	0	1.00	0.04	0.50
Flexible Zoning By Permit	26	0	3	0	1.00	0.00	1.00
Affordable Housing	22	0	7	0	1.00	0.00	1.00
Permit Cap Or Phasing	8	2	19	1	0.89	0.10	0.80
Wetlands Restricted in Lot Size Calc	23	1	6	0	1.00	0.14	0.96
Total	193	6	122	6	0.97	0.05	0.97

Note: This confusion matrix is generated using the Chat GPT-4 Turbo model on the testing sample of 30 municipalities from the Pioneer study. Observations where the model responds "I don't know" or observations we categorized as ambiguous are excluded. True Positive refers to an outcome where the model correctly predicts the positive class. False Positive is an outcome where the model incorrectly predicts the positive class. True Negative denotes an outcome where the model correctly predicts the negative class. False Negative represents an outcome where the model incorrectly predicts the negative class. The true positive rate (also known as sensitivity or recall) is the proportion of actual positive cases correctly identified by the model. The false positive rate (also known as the false alarm rate or fall-out) is the proportion of actual negative cases incorrectly identified as positive by the model. Precision (also known as positive predictive value) is the proportion of positive identifications that are actually correct.

Table 4: Relationship between Our PCA-derived Indices and Wharton Residential Land Use Regulatory Index

Panel A: Averages and Correlation For Individual Questions

	Wharton Average	Our Average	Correlation
Affordable Housing	0.20	0.24	0.36
Minimum Lot Size	Less than 1/2 acre	0.48	0.39
	1/2 to 1 acre	0.16	0.10
	1 to under 2 acres	0.13	0.15
	2 acres or more	0.23	0.19

Panel B: Correlation Matrix of Index

	Wharton Index	Our Index PC 1	Our Index PC 2
Wharton Index	1.00	0.31	0.09
Our Index PC 1	0.31	1.00	0.10
Our Index PC 2	0.09	0.10	1.00

Note: The sample overlap between our study and [Gyourko et al. \(2021\)](#) is 1,283 municipalities. The question on affordable housing in our study (question 17 from the Pioneer study) considers both mandates and incentives, whereas the Wharton study (question 9a) only considers affordable housing mandates. For minimum lot sizes, our study considers minimum lot sizes across all districts, while the Wharton study (question 7b) only considers residential districts. We drop municipalities that do not have any minimum lot size requirements. We follow the Wharton methodology to aggregate our index to the CBSA level by taking a simple average of all municipalities in that CBSA (only those that are in both our and the Wharton dataset).

Table 5: Associates of Unincorporated Areas

	Bivariate	Metro FE	Distance FE	Metro and Distance FE
Median Home Value	-6.0*** (0.1)	-2.6*** (0.1)	-1.9*** (0.1)	-0.3*** (0.1)
Median Year Built	13.4*** (0.1)	7.2*** (0.1)	11.9*** (0.1)	5.6*** (0.1)
Median Gross Rent	-98.2*** (3.0)	-81.1*** (2.7)	-8.8*** (2.8)	-37.4*** (2.6)
Vacancy Rate	2.7*** (0.0)	2.2*** (0.0)	1.9*** (0.0)	1.6*** (0.0)
Rental Rate	-12.6*** (0.1)	-15.1*** (0.1)	-9.9*** (0.1)	-12.5*** (0.1)
Percent Commute Over 60	1.2*** (0.0)	2.9*** (0.0)	1.6*** (0.0)	2.8*** (0.0)
Percent Over 65	2.0*** (0.0)	2.5*** (0.0)	1.3*** (0.0)	1.7*** (0.0)

Note: For each block group, we determine the percent of its area that is contained within a local government, either a municipality or township, based on data from the 2022 Census of Governments. Block groups that are entirely unincorporated are assigned an indicator value of 1, while those that are entirely incorporated are assigned a value of 0. Partially incorporated block groups are split into both an incorporated and an unincorporated observation, with weights corresponding to the percent incorporated. We use regression weights that are the product of the block group population and the percent are unincorporated, or percent area incorporated for incorporated areas. For each block group, we identify the nearest metropolitan center, measure the distance from this center, and record the corresponding metro area. Distance fixed effects are defined as quintiles of the distance from the city center to account for potential non-linear effects of distance on the dependent variables.

Table 6: Coefficients on Normalized Zoning Questions Regressed Against Log Distance to Metro Center

Variable	US Census Region				All Regions
	West	South	Midwest	Northeast	
Accessory Apartments Allowed	-0.08*** (0.03)	-0.05** (0.02)	-0.09*** (0.03)	0.10*** (0.02)	-0.01 (0.01)
Flexible Zoning By Right	-0.03 (0.03)	-0.02 (0.02)	-0.04* (0.02)	0.03** (0.02)	-0.00 (0.01)
Flexible Zoning By Permit	0.01 (0.03)	-0.02 (0.03)	-0.09*** (0.03)	0.05*** (0.01)	0.00 (0.01)
Affordable Incentive	-0.03 (0.03)	-0.18*** (0.03)	-0.32*** (0.04)	0.00 (0.01)	-0.06*** (0.01)
Affordable Mandate	-0.00 (0.02)	-0.06 (0.06)	-0.19** (0.08)	-0.01 (0.01)	-0.01 (0.01)
Zoning District Count	-0.13*** (0.03)	-0.19*** (0.03)	-0.09*** (0.02)	-0.06*** (0.02)	-0.11*** (0.01)
Permit Cap Or Phasing	-0.01 (0.03)	-0.00 (0.03)	-0.05** (0.02)	0.01 (0.02)	-0.01 (0.01)
Wetlands Restricted in Lot Size Calc	-0.03 (0.04)	-0.03 (0.03)	0.00 (0.03)	0.06*** (0.01)	0.03** (0.01)
Longest Frontage Requirement	0.03 (0.04)	-0.03 (0.03)	-0.02 (0.03)	0.11*** (0.02)	0.05*** (0.01)
Maximum Res Min Lot Size	0.00 (0.02)	-0.01 (0.03)	0.00 (0.02)	0.12*** (0.02)	0.04*** (0.01)
Mean Res Min Lot Size	0.03 (0.03)	-0.00 (0.04)	0.01 (0.02)	0.15*** (0.01)	0.08*** (0.01)
Minimum Res Min Lot Size	0.08 (0.08)	-0.02 (0.04)	0.05 (0.04)	0.09*** (0.01)	0.08*** (0.01)
Mandatory Approval Steps	-0.02 (0.03)	0.01 (0.02)	0.02 (0.02)	0.01 (0.02)	0.01 (0.01)
Distinct Approval Bodies	0.02 (0.03)	0.02 (0.02)	0.01 (0.02)	0.02 (0.02)	0.02 (0.01)
Public Hearing Requirements	0.01 (0.03)	-0.04 (0.03)	-0.01 (0.02)	0.06*** (0.02)	0.02 (0.01)
Max Review Waiting Time	-0.06** (0.03)	-0.00 (0.02)	-0.08*** (0.02)	0.03 (0.02)	-0.02* (0.01)
Multifamily Allowed	0.10*** (0.03)	0.08** (0.03)	0.06** (0.03)	-0.02* (0.01)	0.02 (0.01)
Mixed-Use Buildings	-0.04 (0.03)	-0.10*** (0.03)	-0.06** (0.02)	-0.02 (0.02)	-0.05*** (0.01)
Conversion To Multifamily	-0.06 (0.05)	0.03 (0.03)	0.06*** (0.02)	-0.02* (0.01)	0.00 (0.01)
Townhouses Allowed	-0.04 (0.03)	-0.04 (0.03)	-0.08*** (0.02)	-0.07*** (0.02)	-0.06*** (0.01)
Age-Restricted Provisions	-0.06* (0.03)	-0.09*** (0.03)	-0.11*** (0.02)	-0.02 (0.02)	-0.06*** (0.01)
First Principal Component (Complexity)	-0.08*** (0.03)	-0.17*** (0.03)	-0.24*** (0.03)	0.04*** (0.01)	-0.06*** (0.01)
Second Principal Component (Strictness)	0.05 (0.04)	0.03 (0.03)	0.06** (0.03)	0.14*** (0.01)	0.10*** (0.01)
MSA Fixed Effects	Yes	Yes	Yes	Yes	Yes

Note: We perform this regression on the set of municipalities within 100 miles of the center of a metropolitan area which accounts for 4,819 observations in our sample. The dependent variable is log distance to metro center. A positive coefficient indicates that the variable increases with log distance from the metro center and a negative coefficient means that the variable decreases with log distance from the metro center. See Appendix Table A1 for full definitions of zoning questions.

Table 7: Regression Table for PC1

Variable	Bivariate FE	Bivariate	All	All FE	Elasticity	Elasticity FE
Foreign Born Share	0.04*** (0.02)	0.16*** (0.01)	0.06*** (0.02)	0.00 (0.03)		
Median Household Income	0.03** (0.02)	0.21*** (0.01)	-0.05 (0.05)	-0.04 (0.05)		
Share Population 65 and Over	-0.09*** (0.01)	-0.07*** (0.01)	-0.01 (0.03)	-0.03 (0.03)		
Median Gross Rent	0.09*** (0.02)	0.28*** (0.01)	0.03 (0.03)	-0.05 (0.03)		
Median Home Value	-0.07*** (0.02)	0.17*** (0.01)	-0.01 (0.03)	-0.14*** (0.04)		
Share Units Owner Occupied	-0.06*** (0.01)	-0.00 (0.01)	-0.20*** (0.04)	-0.13*** (0.04)		
Share Population Under 18	-0.00 (0.01)	-0.06*** (0.01)	-0.05* (0.02)	-0.03 (0.02)		
White Share	-0.02 (0.01)	-0.04*** (0.01)	0.18*** (0.03)	0.18*** (0.03)		
Poverty Rate	-0.07*** (0.01)	-0.19*** (0.01)	-0.08** (0.03)	-0.08** (0.03)		
College Share	0.14*** (0.01)	0.26*** (0.01)	-0.02 (0.04)	0.09** (0.04)		
Share Structures Built Before 1970	-0.27*** (0.01)	-0.19*** (0.01)	-0.05** (0.02)	-0.12*** (0.02)		
Share Structures with 2 or More Units	0.11*** (0.01)	0.14*** (0.01)	0.03 (0.03)	0.02 (0.03)		
Vacancy Rate	-0.13*** (0.01)	-0.14*** (0.01)	-0.09*** (0.03)	-0.04 (0.03)		
Share with Commute Over 30 Minutes	-0.09*** (0.02)	0.12*** (0.01)	0.06** (0.02)	0.01 (0.02)		
Job Density	0.05*** (0.01)	0.13*** (0.01)	-0.02 (0.02)	-0.01 (0.02)		
Opportunity Index	-0.01 (0.02)	0.06*** (0.01)	-0.04 (0.03)	-0.05* (0.03)		
Average Math Test Scores	0.14*** (0.01)	0.21*** (0.01)	-0.03 (0.03)	0.07* (0.04)		
Math Learning Rate	0.07*** (0.01)	0.09*** (0.01)	0.00 (0.02)	0.04** (0.02)		
Percent Eligible for Free Lunch	-0.14*** (0.01)	-0.22*** (0.01)	-0.17*** (0.04)	-0.03 (0.04)		
Property Tax Rate	-0.03** (0.01)	0.06*** (0.01)	0.04** (0.02)	0.02 (0.02)		
Total Expenditure Per Capita (2017)	0.04*** (0.01)	0.12*** (0.01)	-0.02 (0.02)	0.00 (0.02)		
Building Permits All Units 2021	0.05*** (0.01)	0.08*** (0.01)	0.03 (0.02)	0.03 (0.02)		
Year of Incorporation	-0.08*** (0.01)	0.01 (0.01)	-0.00 (0.02)	0.01 (0.02)		
Percent Democrat	0.15*** (0.02)	0.27*** (0.01)	0.21*** (0.03)	0.20*** (0.04)		
New Housing Unit Elasticity	0.05** (0.02)	-0.06*** (0.02)	-0.04 (0.02)	-0.01 (0.03)	-0.38*** (0.05)	-0.12* (0.07)
Log Land Area	0.45*** (0.01)	0.46*** (0.01)	0.38*** (0.02)	0.37*** (0.02)		
Log Neighbors within 25 Miles	0.21*** (0.03)	0.09*** (0.01)	-0.10*** (0.02)	0.13** (0.05)		
Housing Unit Density	-0.07*** (0.01)	0.01 (0.01)	0.00 (0.02)	-0.01 (0.02)		
Miles to Metro Center	-0.00 (0.02)	-0.14*** (0.01)	0.06*** (0.02)	0.06** (0.03)	-0.10*** (0.02)	-0.06*** (0.02)
Share Units Affordable	-0.16*** (0.02)	-0.31*** (0.01)	-0.12*** (0.04)	-0.02 (0.05)		
MSA Fixed Effects	Yes	No	No	Yes	No	Yes
Housing Elasticity Controls	No	No	No	No	Yes	Yes
R-squared			0.30	0.40	0.05	0.20
N			2577	2577	3891	3891

Note: All right-hand side variables are in Z-scores. Asterisks denote significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are shown in parentheses. Housing elasticity controls follow [Baum-Snow and Han \(2024\)](#) and include fraction of land developed in 2001, squared fraction of land developed in 2001, and the fraction of land with a flat topography.

Note: All right-hand side variables are in Z-scores. Asterisks denote significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are shown in parentheses. We use State FE for municipalities not within an MSA. For variable definitions please see Appendix Table A3.

Table 8: Regression Table for PC2

Variable	Bivariate FE	Bivariate	All	All FE	Elasticity	Elasticity FE
Foreign Born Share	-0.20*** (0.01)	-0.04*** (0.01)	0.06*** (0.02)	0.02 (0.02)		
Median Household Income	0.31*** (0.01)	0.40*** (0.01)	0.25*** (0.04)	0.17*** (0.04)		
Share Population 65 and Over	0.13*** (0.01)	0.16*** (0.01)	0.07*** (0.02)	0.08*** (0.02)		
Median Gross Rent	0.14*** (0.02)	0.23*** (0.01)	0.10*** (0.03)	0.07*** (0.03)		
Median Home Value	0.24*** (0.01)	0.28*** (0.01)	0.08*** (0.03)	0.12*** (0.03)		
Share Units Owner Occupied	0.33*** (0.01)	0.41*** (0.01)	0.01 (0.03)	-0.01 (0.04)		
Share Population Under 18	-0.02 (0.01)	-0.11*** (0.01)	-0.05** (0.02)	-0.04* (0.02)		
White Share	0.22*** (0.01)	0.22*** (0.01)	0.04 (0.03)	0.02 (0.03)		
Poverty Rate	-0.14*** (0.01)	-0.27*** (0.01)	0.13*** (0.03)	0.10*** (0.03)		
College Share	0.12*** (0.01)	0.28*** (0.01)	-0.03 (0.03)	0.02 (0.04)		
Share Structures Built Before 1970	-0.11*** (0.01)	0.05*** (0.01)	0.05*** (0.02)	-0.03* (0.02)		
Share Structures with 2 or More Units	-0.35*** (0.01)	-0.32*** (0.01)	-0.07*** (0.03)	-0.13*** (0.03)		
Vacancy Rate	0.06*** (0.01)	0.03** (0.01)	-0.01 (0.02)	0.00 (0.03)		
Share with Commute Over 30 Minutes	0.12*** (0.01)	0.24*** (0.01)	0.03* (0.02)	-0.00 (0.02)		
Job Density	-0.21*** (0.01)	-0.14*** (0.01)	-0.04** (0.02)	-0.02 (0.02)		
Opportunity Index	0.14*** (0.01)	0.23*** (0.01)	0.03 (0.02)	-0.01 (0.03)		
Average Math Test Scores	0.15*** (0.01)	0.30*** (0.01)	-0.02 (0.03)	0.01 (0.03)		
Math Learning Rate	0.08*** (0.01)	0.11*** (0.01)	-0.03* (0.01)	-0.01 (0.02)		
Percent Eligible for Free Lunch	-0.19*** (0.01)	-0.34*** (0.01)	-0.00 (0.03)	0.06 (0.04)		
Property Tax Rate	-0.14*** (0.01)	0.07*** (0.01)	0.08*** (0.01)	0.02 (0.02)		
Total Expenditure Per Capita (2017)	-0.06*** (0.01)	0.12*** (0.01)	-0.00 (0.02)	-0.02 (0.02)		
Building Permits All Units 2021	-0.01 (0.01)	-0.03** (0.01)	-0.04** (0.02)	-0.04*** (0.02)		
Year of Incorporation	0.13*** (0.01)	0.10*** (0.01)	0.03* (0.02)	0.05*** (0.02)		
Percent Democrat	-0.21*** (0.01)	0.03** (0.01)	-0.05* (0.02)	-0.07** (0.03)		
New Housing Unit Elasticity	0.38*** (0.02)	0.16*** (0.02)	0.13*** (0.02)	0.14*** (0.02)	-0.75*** (0.05)	-0.03 (0.06)
Log Land Area	-0.02* (0.01)	0.09*** (0.01)	-0.08*** (0.02)	-0.06*** (0.02)		
Log Neighbors within 25 Miles	-0.11*** (0.02)	0.17*** (0.01)	0.11*** (0.02)	0.05 (0.04)		
Housing Unit Density	-0.24*** (0.01)	-0.18*** (0.01)	-0.06*** (0.02)	-0.05*** (0.02)		
Miles to Metro Center	0.15*** (0.02)	-0.04** (0.02)	-0.00 (0.02)	-0.05** (0.02)	-0.16*** (0.02)	-0.10*** (0.02)
Share Units Affordable	-0.21*** (0.01)	-0.31*** (0.01)	0.02 (0.03)	0.09** (0.04)		
Intercept			-0.68*** (0.09)	-0.82*** (0.25)	0.07*** (0.02)	-0.79*** (0.26)
MSA Fixed Effects	Yes	No	No	Yes	No	Yes
Housing Elasticity Controls	No	No	No	No	Yes	Yes
R-squared			0.32	0.40	0.18	0.43
N			2577	2577	3891	3891

Note: All right-hand side variables are in Z-scores. Asterisks denote significance levels: * p < 0.1, ** p < 0.05, *** p < 0.01. Standard errors are shown in parentheses. Housing elasticity controls follow [Baum-Snow and Han \(2024\)](#) and include fraction of land developed in 2001, squared fraction of land developed in 2001, and the fraction of land with a flat topography. We use State FE for municipalities not within an MSA. For variable definitions please see Appendix Table A3.

Table 9: National Sample Question Means

Panel A: Continuous Questions

Question	Mean	National		Income Tercile			Urban/Rural		
		Weight	Count	Low	Mid	High	Rural	Mix	Urban
How many zoning districts, including overlays, are in the municipality?	14	20	5,463	13	14	14	10	16	14
What is the longest frontage requirement for single family residential development in any district?	92	69	5,198	74	85	117	93	97	80
Mean of Residential Min Lot Sizes	24631	17564	5,409	17136	21771	34914	31029	25706	15557
Min of Residential Min Lot Sizes	10102	5885	5,426	6970	8982	14302	12385	10054	7984
How many mandatory steps are involved in the approval process for a typical new multi-family building?	4	4	5,782	4	4	5	5	4	5
For a typical new multi-family building project in this jurisdiction, how many distinct governing bodies or agencies must give mandatory approval before construction can begin?	3.1	3.0	5,749	3.2	3.1	3.1	3.1	3.2	3.1
What is the maximum potential waiting time (in days) for government review of a typical new multi-family building?	217	212	5,093	195	223	235	200	221	225

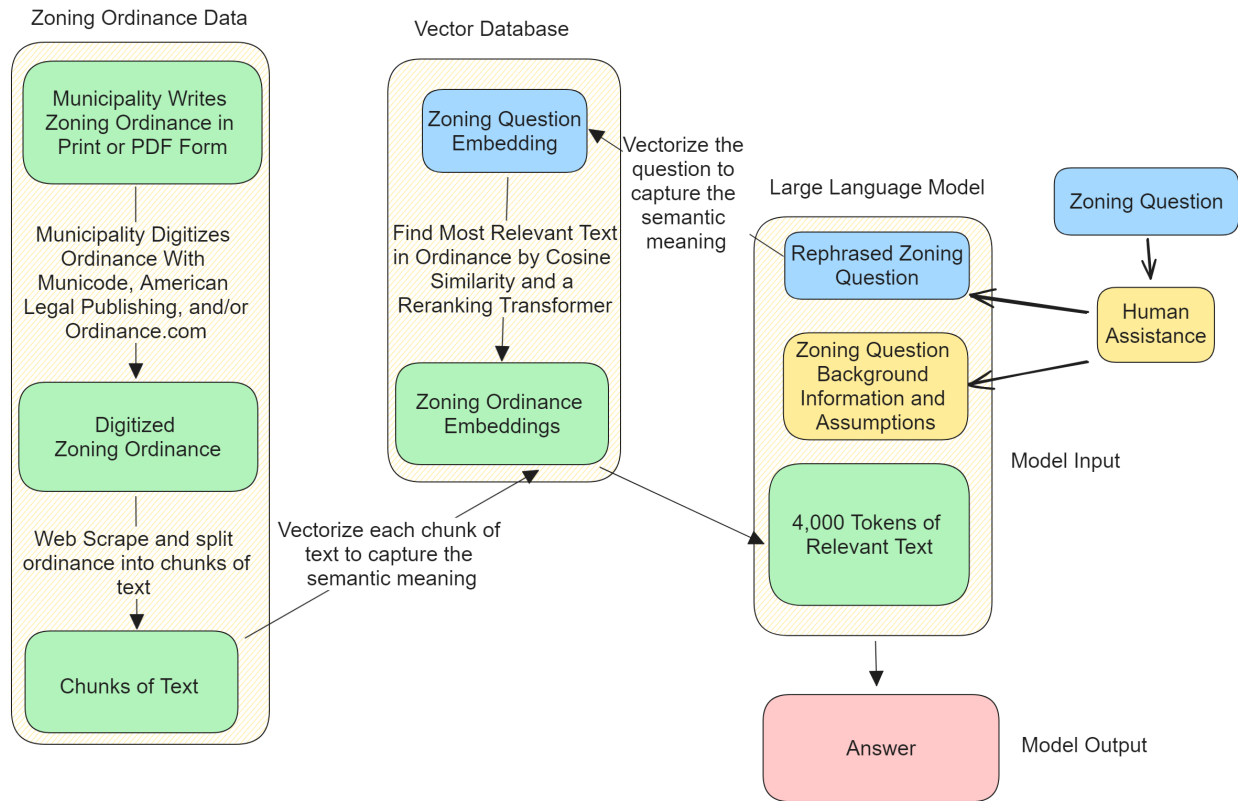
Panel B: Binary Questions

Question	National			Income Tercile			Urban/Rural		
	Mean	Weight	Count	Low	Mid	High	Rural	Mix	Urban
Is multi-family housing allowed, either by right or special permit (including through overlays or cluster zoning)?	95	98	5,686	99	97	90	95	96	93
Are apartments above commercial (mixed use) allowed in any district?	62	70	5,708	64	67	56	55	66	61
Is multi-family housing listed as allowed through conversion (of either single family homes or non residential buildings)?	14	20	5,766	13	14	14	10	15	13
Are attached single family houses (townhouses, 3+ units) listed as an allowed use (by right or special permit)?	80	89	5,788	80	82	78	65	84	83
Does zoning include any provisions for housing that is restricted by age?	45	60	5,115	34	42	58	24	50	51
Are accessory or in-law apartments allowed (by right or special permit) in any district?	33	38	5,771	27	33	39	23	40	22
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?	9	9	5,788	8	8	10	5	10	8
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by special permit?	79	80	5,667	79	80	80	69	86	73
Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?	24	50	5,535	10	22	40	9	28	27
Is there a town-wide annual or biannual cap on residential permits issued, and/or is project phasing required?	17	18	5,794	11	18	21	10	19	16
Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?	9	6	4,602	4	9	17	7	11	7
Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?	7	10	5,775	1	4	15	2	7	10
Are there townwide requirements for public hearings on any type of multi-family residential projects?	30	32	5,703	23	31	36	26	30	31

Note: We define the count (sample size) as the number of municipalities where the model does not say “I don’t know” as the answer. The ‘Weight’ column weights each municipality by its population in the 2022 census of governments. We designate Urban/Rural using the percent overlap of the 2022 shape file for the municipality with the 2020 shape file for urban areas. Specifically, we define Urban as a municipality being 100% in an urban area, Mix as a municipality being partially in an urban area, and Rural as a municipality being 0% in an urban area. From the 2021 Five-Year American Community Survey we use median household income (B19013_001E).


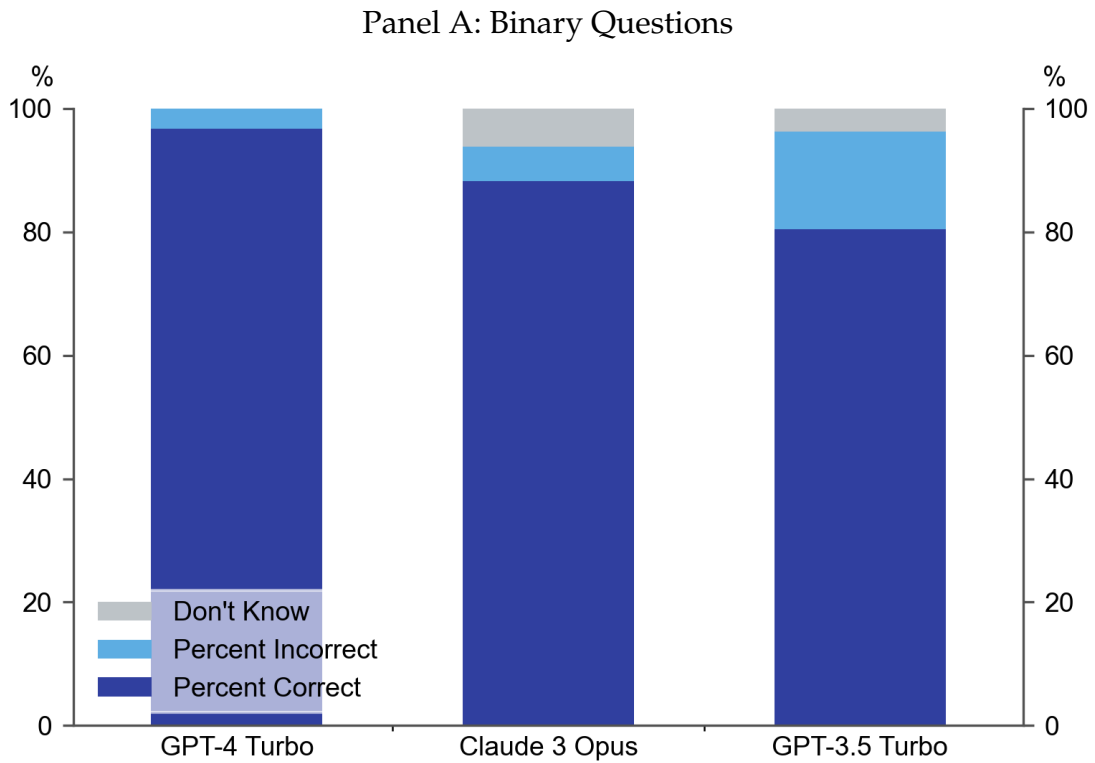
Figures

Figure 1: Model Overview




Note: We scrape each section within a zoning ordinance separately. We split up sections that are longer than one thousand tokens into chunks of at most one thousand tokens. We also combine adjacent sections of less than 50 tokens. So, each section of text varies in length but is between 50 and one thousand tokens. We vectorize each chunk of text using OpenAI embeddings models ([link](#)). Specifically, we used 'text-embedding-3-large'. Sometimes digital aggregators leave tables in image form, especially the aggregator Ordinance.com. So that the model can still read the table, we transcribe images of tables using [Amazon Textract](#). We elicit an open-ended response to each question and then use [function calling](#) to parse out a structured answer (i.e., to ascertain whether an answer is "Yes", "No", or "I don't know" to a binary question). Question background information and model assumptions are based on a combination of the 'Issue Overview' and the 'Research Coding' sections for each question from the [Pioneer study](#) as well as from trial and error in the training sample of municipalities. Rephrased zoning questions came entirely from trial and error on the training sample. Ordinances from digital aggregators (Municode, American Legal Publishing, and Ordinance.com) are either entirely about zoning, partially about zoning (i.e., have one or more sections about zoning), or not about zoning at all. We filter out ordinances not at all about zoning by searching through key phrases, table headers, and zoning district names (i.e., R-1 for the first residential zoning district).

Figure 2: Comparison of Average Performance Across Models



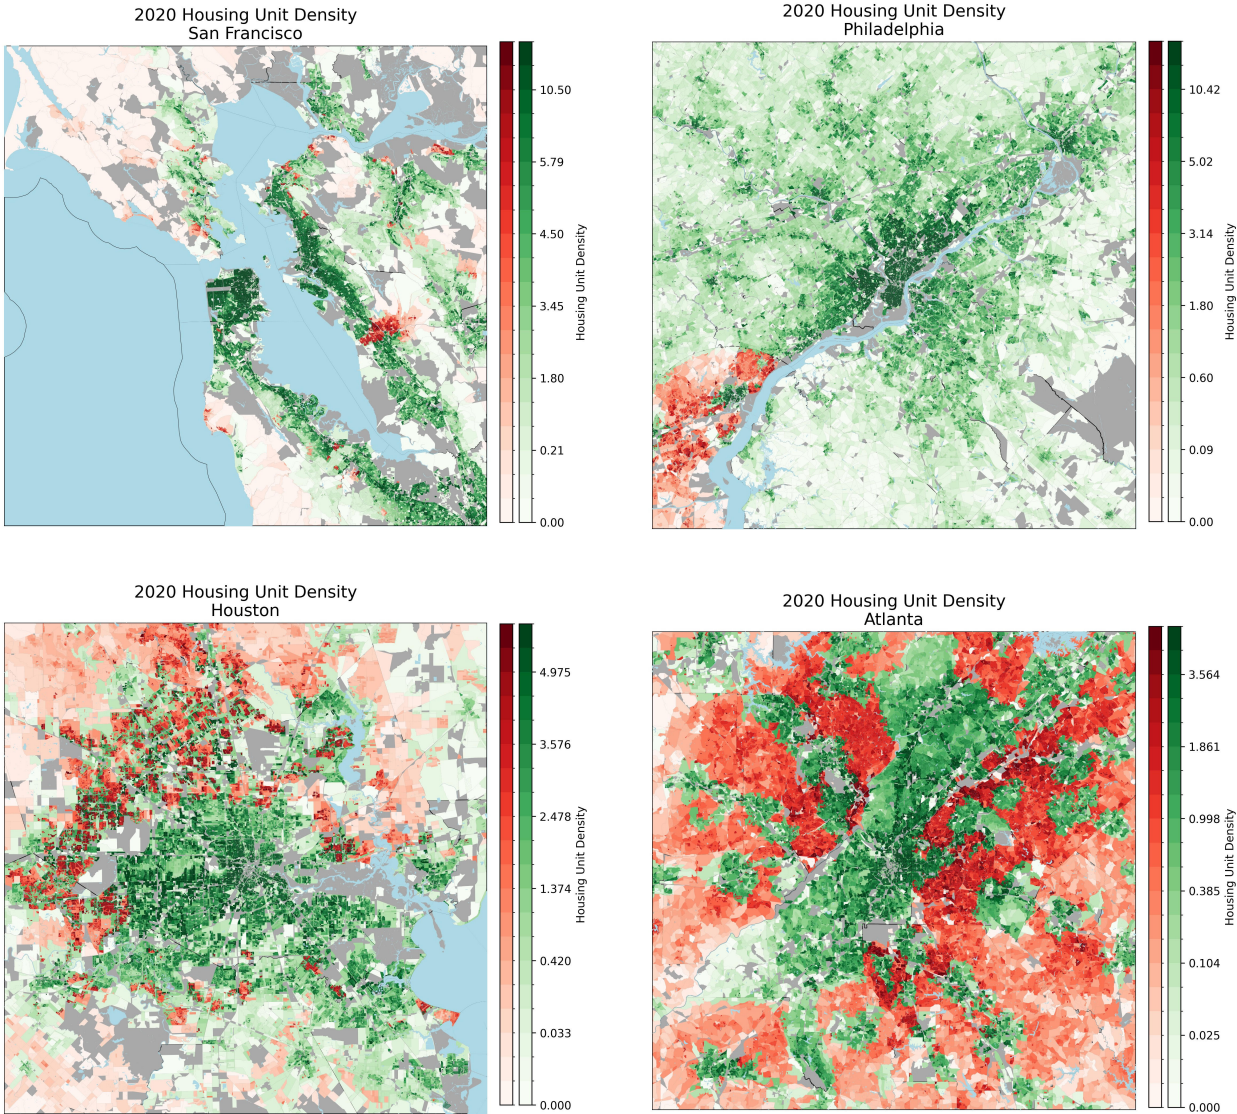
Note: For binary questions we use the percent accuracy and for continuous questions we use the correlation. We drop four question-muni pairs, which we manually categorized as ambiguous answers.

Figure 3: Reasons For Disagreement Between Chat GPT-4 Turbo and Pioneer Study on Binary Questions



Note: We first ran ChatGPT-4 Turbo on the testing sample of 30 randomly selected municipalities that were included in the Pioneer Institute’s study but were not used to train our model. We then identified the binary questions where the model responses disagreed with the Pioneer study. A law student reviewed each of these disagreements individually to determine the reason for the discrepancy, classifying them into the categories shown in the chart.


Figure 4: Unincorporated vs. Incorporated Land in Select Metropolitan Areas




Note: We plot housing unit density in four metropolitan regions: San Francisco Bay Area, Philadelphia, Houston, and Atlanta. Incorporated areas are colored in green, while unincorporated areas are colored in red. The darkness of coloring denotes the degree of density.

Figure 5: Housing Growth in Unincorporated Areas

Panel A: Growth in Housing Units From 2000–2020, by Fraction Incorporated




Panel B: Growth in Housing Units From 2000–2020, Incorporated vs Unincorporated




Note: We count the number of housing units from the 2000 and 2020 census at the census block level. For each block we determine the percent of the block that is incorporated by calculating the percent of its area that overlaps with a municipality or township government in the 2002 Census of Governments. We then multiply the percent of the block that is incorporated by the number of housing units in the block to get the incorporated and unincorporated share of housing units for both 2000 and 2020. We focus on the South and West Census regions since nearly all of the Northeast and Midwest regions are incorporated.

Figure 6: Allowable Zoning Typologies in American Cities

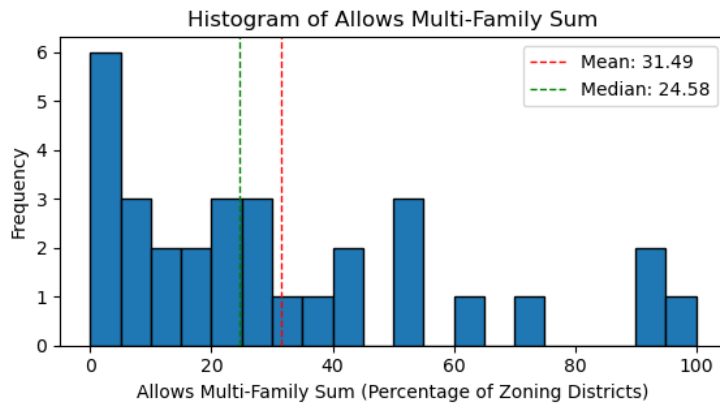
Panel A: Single-Family Only



Panel B: Single-Family or Duplex

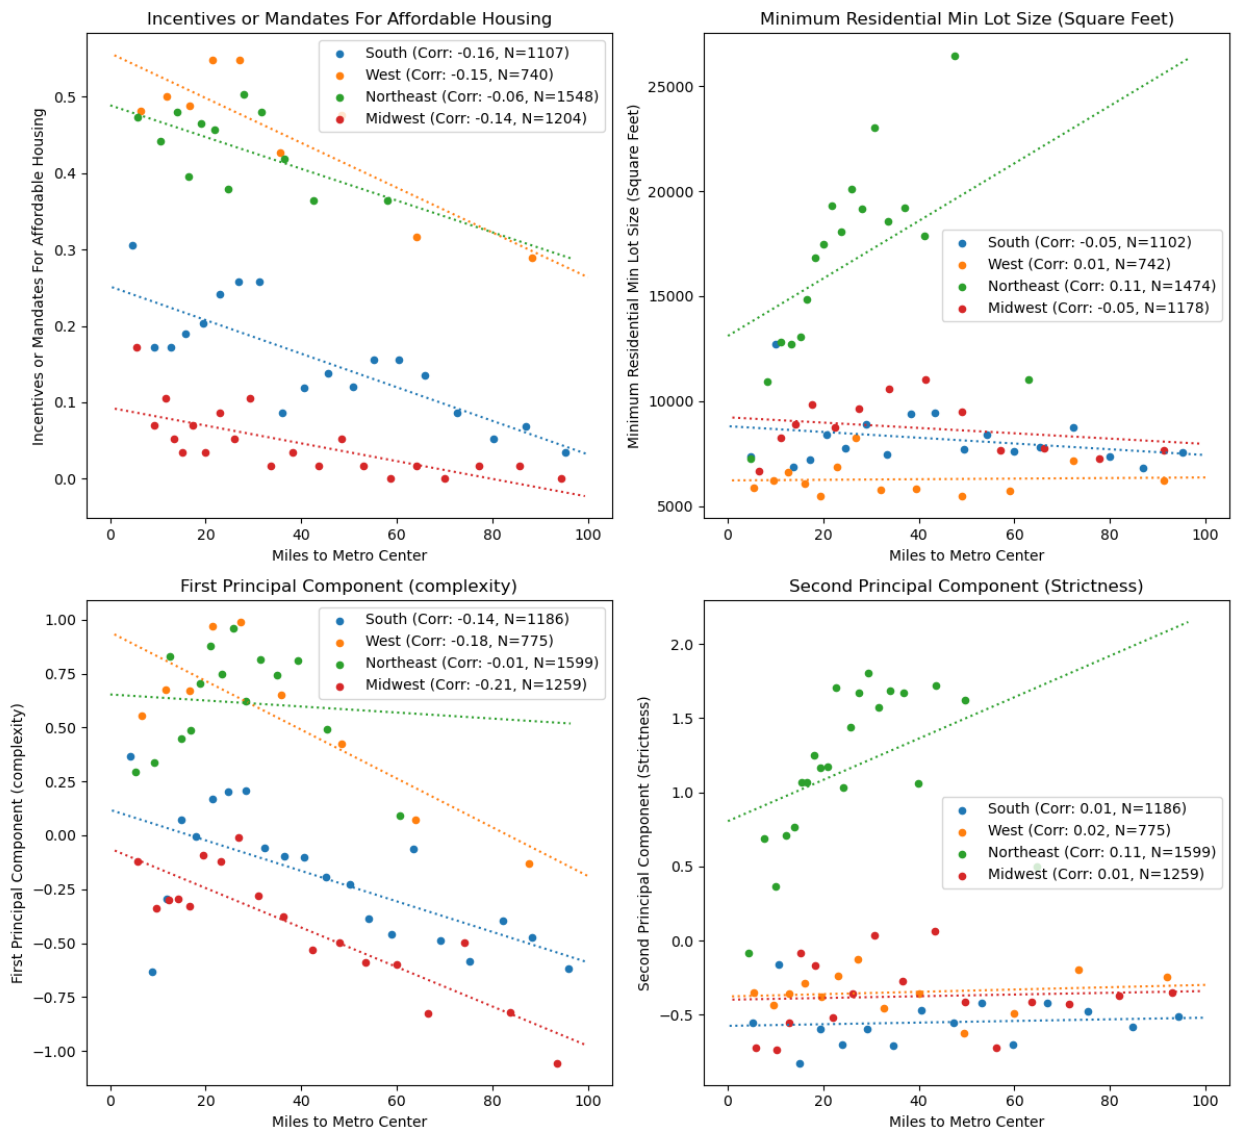


Panel C: Multi-Family



Note: We show the distribution of overall land area zoned for three different uses: single-family only, single-family or duplex, and multi-family. We show results for 31 municipalities with 18.2 million in population, covering Chicago, Seattle, Kansas City, Detroit, San Francisco, Austin, San Antonio, Tampa, Los Angeles, and San Diego.


Figure 7: Regulatory Differences by Distance to City Center




Note: We plot regulatory variables at the municipality level based on the distance from the center of the city for the 73 largest MSAs. We define city hall as the center of the city. We show whether a city has an affordable housing mandate; the minimum lot size; the first principal component of housing regulation, and the second principle component of housing regulations.

Figure 8: Nationwide Maps of Zoning Indices

Panel A: First Principal Component



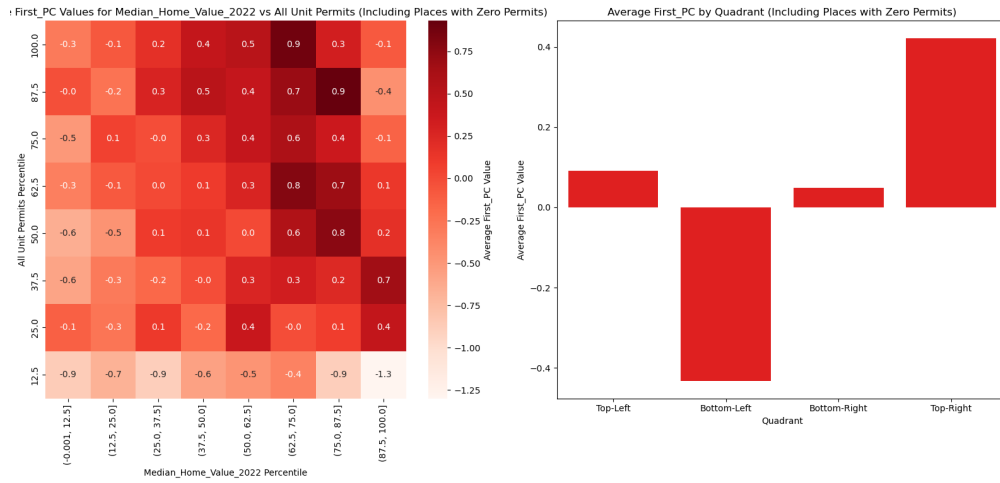
Panel B: Second Principal Component



53
Note: The state-level zoning index value is calculated as the simple average of the index values for all municipalities and townships with available data in our dataset for each state. States shaded in grey have fewer than 10 observations and their index values are not plotted.

Figure 9: Heatmap of Principal Components Vs Housing Cost and Construction

Panel A: First Principal Component



Panel B: Second Principal Component

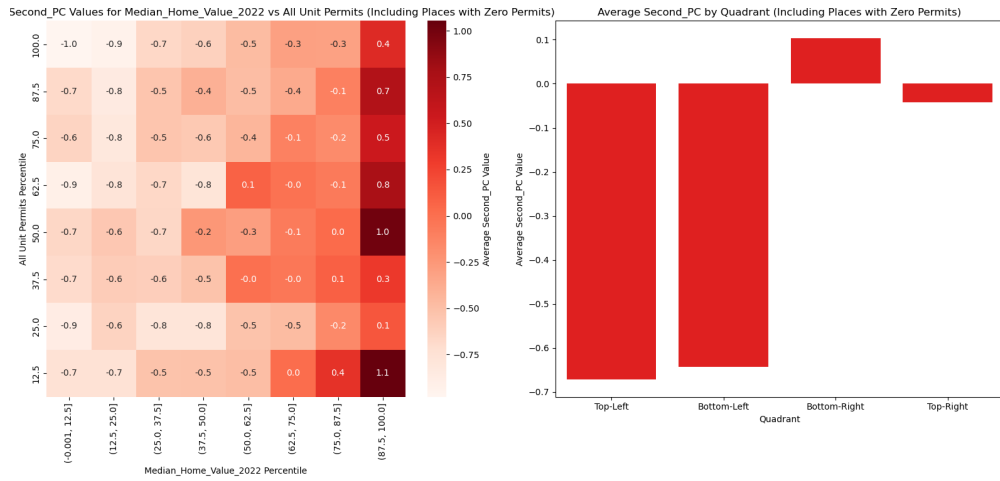
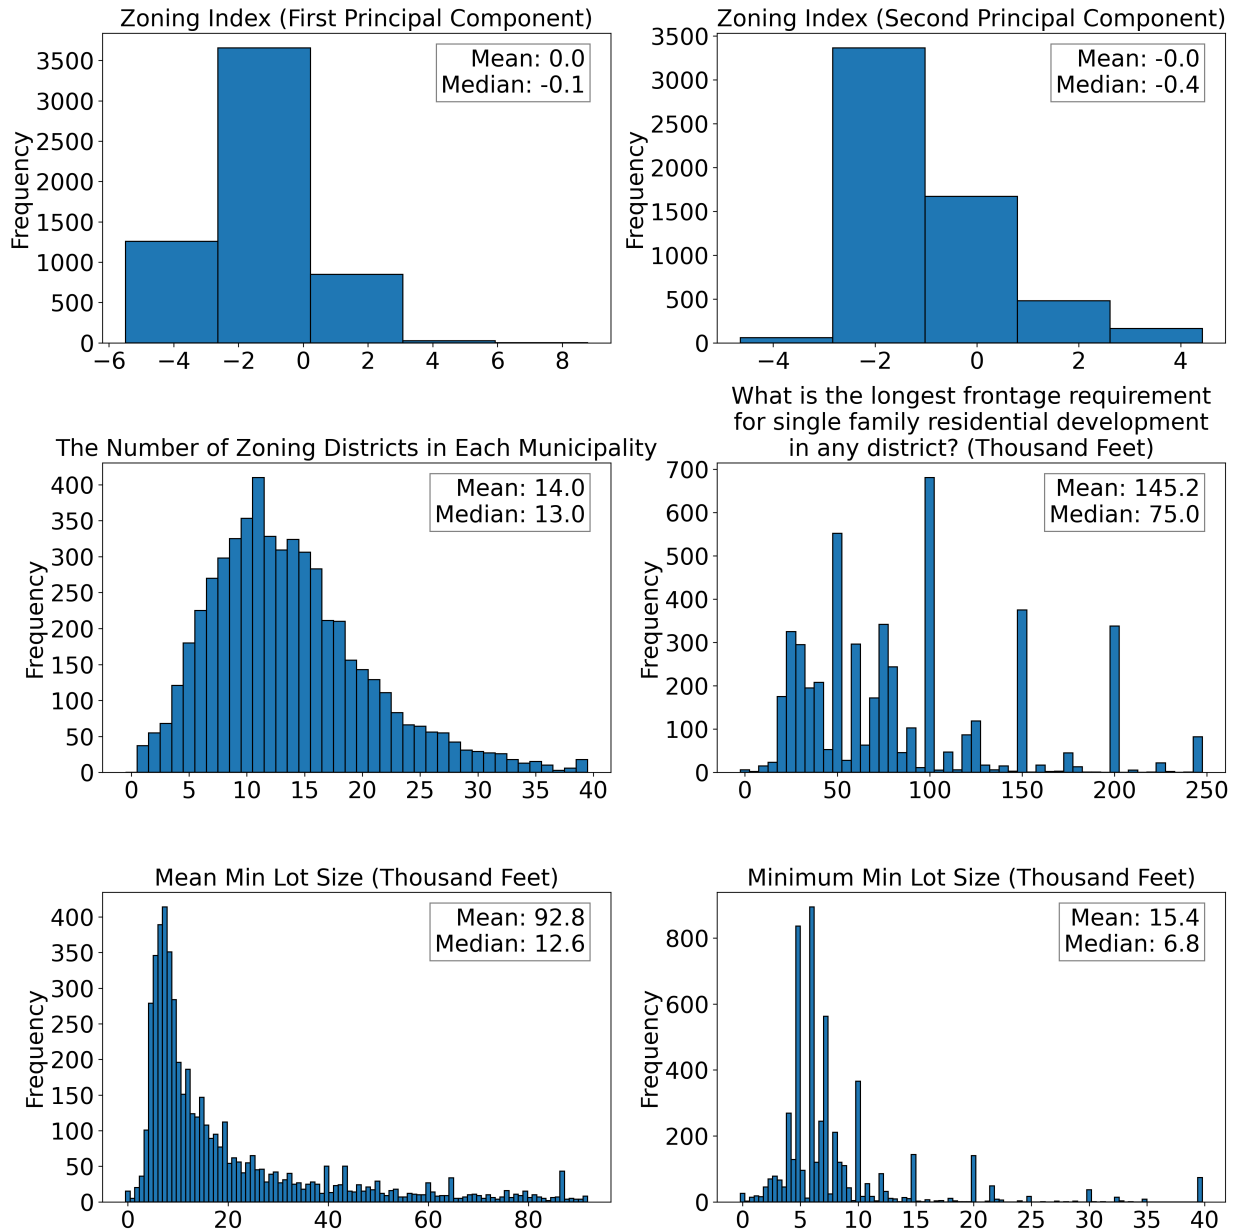
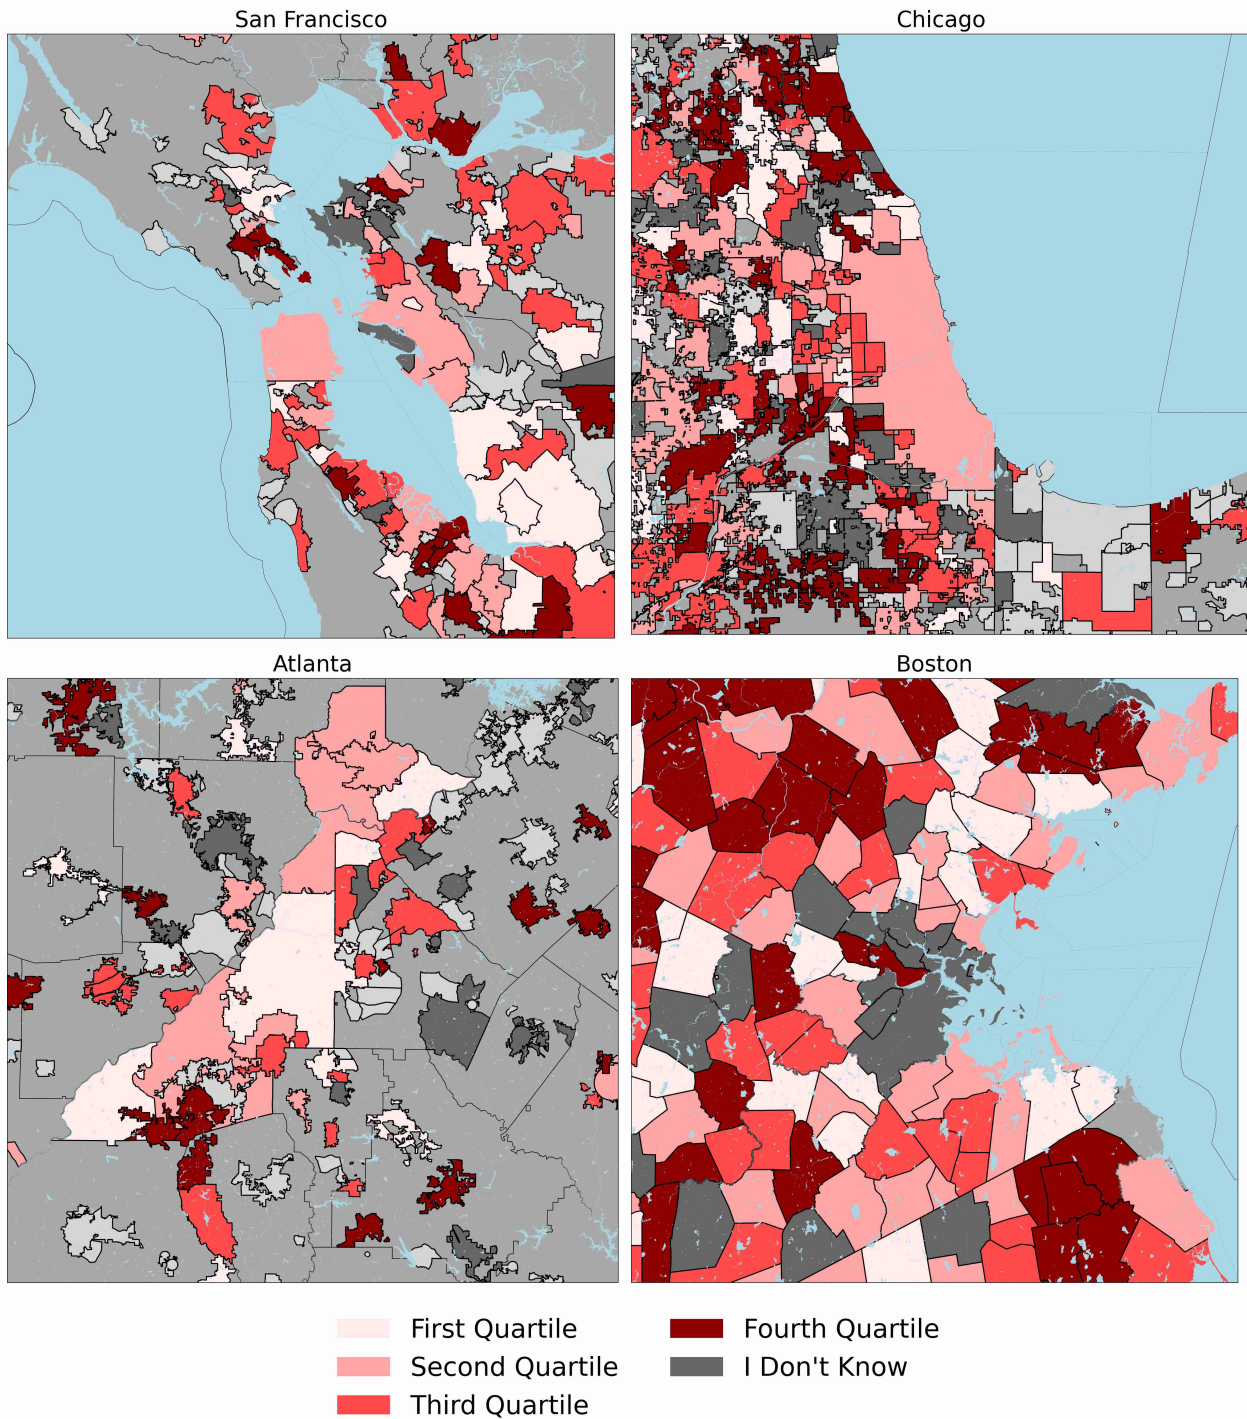


Figure 10: Distribution of Zoning Indices and Housing Regulations



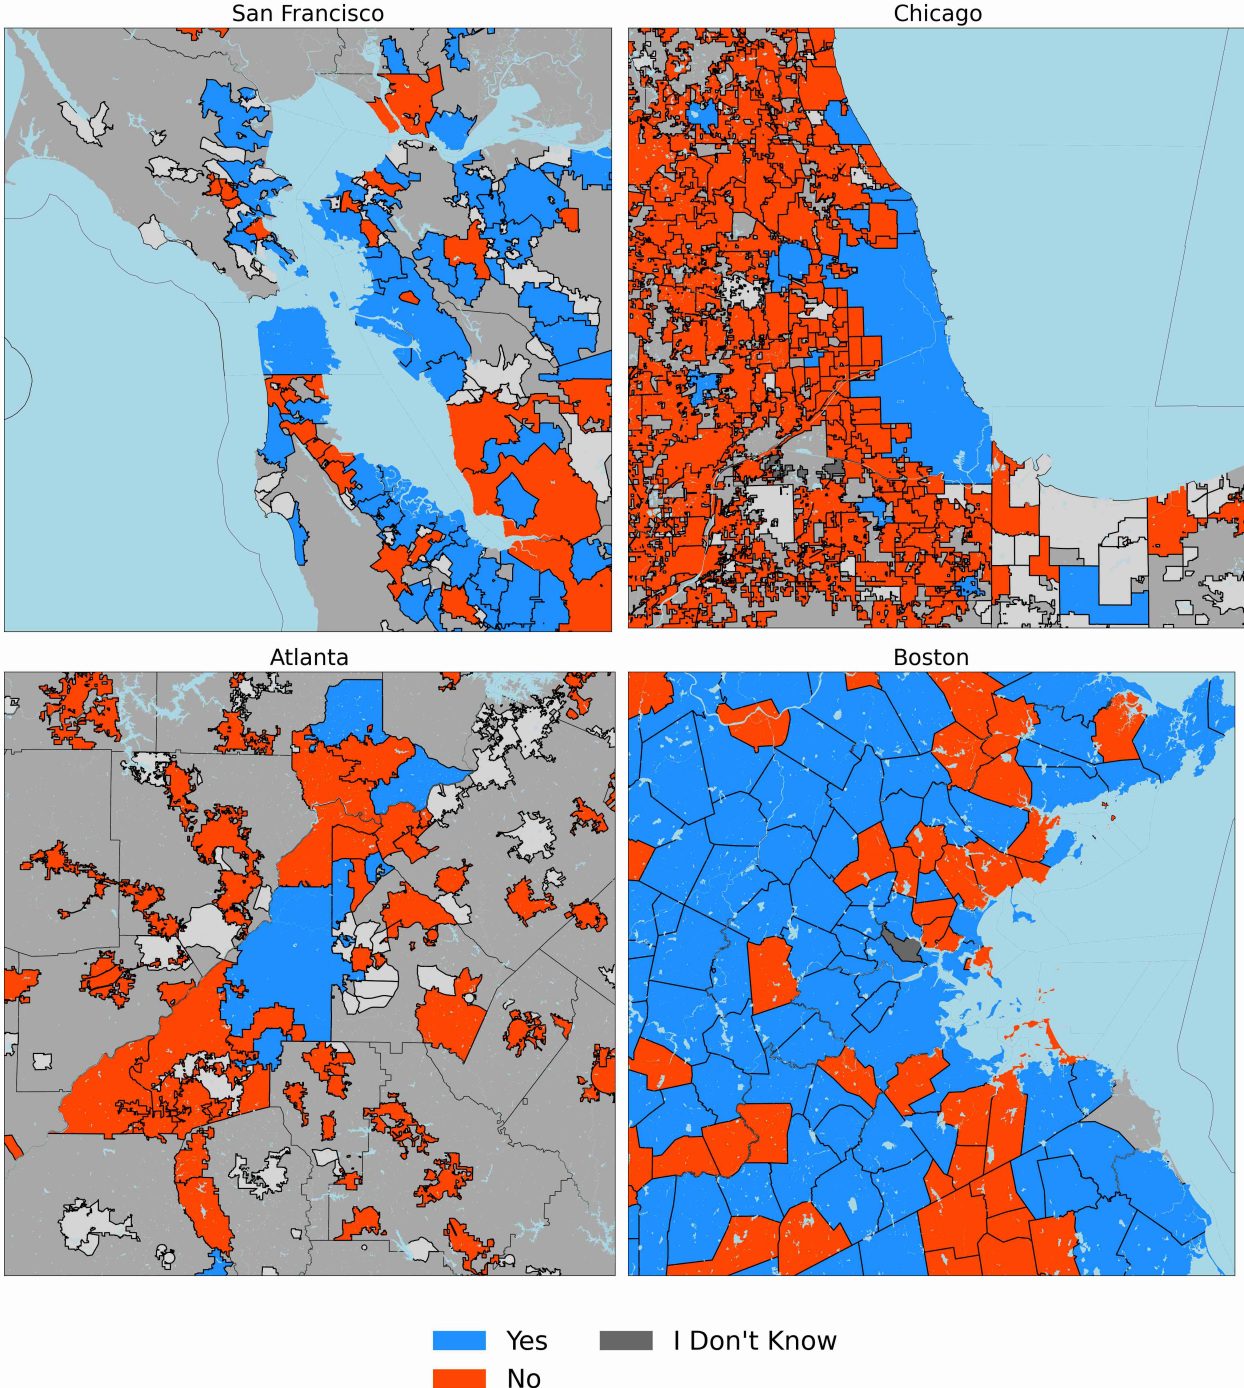
Note: See table 9 footnote for details on the sample. We cut the x-axis at the 99th percentile for the number of districts as well as the second principal component zoning index and at the 95th percentile for the minimum lot size and frontage questions. Mean and median include all data.

Figure 11: Minimum Minimum Lot Size Quartiles For Select Metropolitan Areas



Note: Each map shows roughly a 100km × 100km square area, except for Boston where we show a 75km × 75km square area. Within each map we plot all Census-designated places, except for Boston where we also plot Census county subdivisions that correspond with townships. Both Census-designated place and Census county subdivisions data comes from the 2022 Census TIGER/Line shape files.

Figure 12: Whether There Are Mandates or Incentives For The Development of Affordable Units in Select Metropolitan Areas



Note: Each map shows roughly a 100km x 100km square area, except for Boston where we show a 75km x 75km square area. Within each map we plot all Census-designated places, except for Boston where we also plot Census county subdivisions that correspond with townships. Both Census-designated place and Census county subdivisions data comes from the 2022 Census TIGER/Line shape files.

Figure 13: Correlation Between Median Gross Rents, Median Home Values, Building Permits Per Capita and Zoning Regulations



Note: Univariate correlations are calculated over all valid municipality question pairs (i.e. where the model does not say "I don't know") with a valid outcome variable (i.e. not missing). We winsorize continuous variable answers from our model at the 5% level, but do not winsorize housing outcomes data. Median Gross Rent data comes from both the 2021 and 2010 Five-Year American Community Surveys we use median gross rent (B25064_001E). Median Home Value data comes from both the 2021 and 2010 Five-Year American Community Surveys we use median home value (B25077_001E). Building permits data comes from the 2022 Census Building Permits Survey we use the estimated number of units permitted in 2022. Multi-Unit covers any building with 2-units or more.


Figure 14: Correlations Between Changes in Median Gross Rents, Changes in Median Home Value, Single-Family Building Permits, Multi-Family Building Permit Units and Zoning Regulations



Note: Univariate correlations are calculated over all valid municipality question pairs (i.e. where the model does not say “I don’t know”) with a valid outcome variable (i.e. not missing). We winsorize continuous variable answers from our model at the 5% level, but do not winsorize housing outcomes data. Median Gross Rent data comes from both the 2021 and 2010 Five-Year American Community Surveys we use median gross rent (B25064_001E). Median Home Value data comes from both the 2021 and 2010 Five-Year American Community Surveys we use median home value (B25077_001E). Building permits data comes from the 2022 Census Building Permits Survey we use the estimated number of units permitted in 2022. Multi-Unit covers any building with 2-units or more.

A Appendix

Figure A1: Percent of Land Unincorporated Vs Distance From Metro Center



Note: Percent incorporated is measured as the percent of the block-group that overlaps with a local government (either municipality or township) from the 2022 Census of Governments.

Figure A2: Minimum Lot Size Distribution



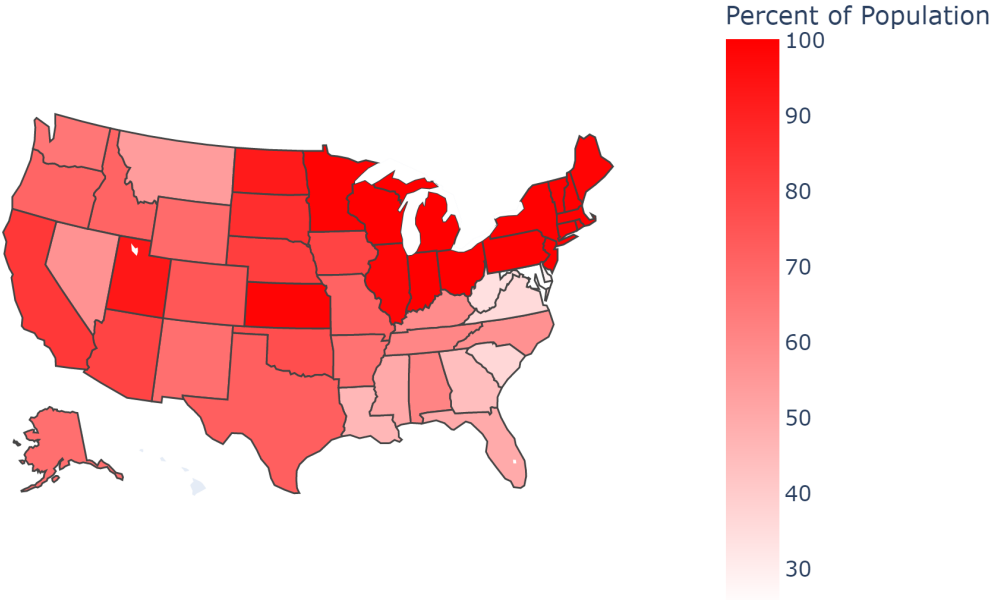


Figure A3: Heatmap of Pairwise Correlations Between Zoning Questions




Note: Please see Appendix Table A1 for full question names. We drop observations where the model says "I don't know".

Figure A4: Percent of the Population Living in Either a Municipality or Township Government By State



Note: See Table 1 footnote for more details on sample coverage

Figure A5: Our Sample Percent of Coverage of Population That Lives Under a Municipality or Township By State



Note: See Table 1 footnote for more details on sample coverage

Table A1: Mapping of Full Pioneer Institute Study Questions to Short Names

Full Question	Short Question
How many zoning districts, including overlays, are in the municipality?	Zoning District Count
Is multi-family housing allowed, either by right or special permit (including through overlays or cluster zoning)?	Multifamily Allowed
Are apartments above commercial (mixed use) allowed in any district?	Mixed-Use Buildings
Is multi-family housing listed as allowed through conversion (of either single family homes or non residential buildings)?	Conversion To Multifamily
Are attached single family houses (townhouses, 3+ units) listed as an allowed use (by right or special permit)?	Townhouses Allowed
Does zoning include any provisions for housing that is restricted by age?	Age-Restricted Provisions
Are accessory or in-law apartments allowed (by right or special permit) in any district?	Accessory Apartments Allowed
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?	Flexible Zoning By Right
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by special permit?	Flexible Zoning By Permit
Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?	Affordable Housing
Is there a town-wide annual or biannual cap on residential permits issued, and/or is project phasing required?	Permit Cap Or Phasing
Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?	Wetlands Restricted in Lot Size Calc
What is the longest frontage requirement for single family residential development in any district?	Longest Frontage Requirement
What is the minimum lot size for single-family homes in each residential district?	Mean Min Lot Size
What is the minimum lot size for single-family homes in each residential district?	Minimum Min Lot Size
How many mandatory steps are involved in the approval process for a typical new multi-family building?	Mandatory Approval Steps
For a typical new multi-family building project in this jurisdiction, how many distinct governing bodies or agencies must give mandatory approval before construction can begin?	Distinct Approval Bodies
Are there townwide requirements for public hearings on any type of multi-family residential projects?	Public Hearing Requirements
What is the maximum potential waiting time (in days) for government review of a typical new multi-family building?	Max Review Waiting Time

Note: "Full Question" refers to how each question was phrased in the Pioneer Institute study and "Short Question" refers to how we abbreviate the question in parts of the paper. Note that the Pioneer Institute study drew on external sources for information on minimum lot sizes, we create additional questions to mimic those variables.

Table A2: Loadings on Principal Components

	First Principal Component	Second Principal Component
Does the zoning bylaw / ordinance include any mandates or incentives for development of affordable units?	0.41	0.11
Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?	0.31	0.11
Does zoning include any provisions for housing that is restricted by age?	0.31	0.01
How many zoning districts, including overlays, are in the municipality?	0.30	-0.19
Is there a town-wide annual or biannual cap on residential permits issued, and/or is project phasing required?	0.22	0.03
Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?	0.22	0.21
Maximum of Residential Min Lot Sizes	0.18	0.38
Are there townwide requirements for public hearings on any type of multi-family residential projects?	0.15	0.12
What is the longest frontage requirement for single family residential development in any district?	0.15	0.41
What is the maximum potential waiting time (in days) for government review of a typical new multi-family building?	0.12	0.03
Minimum of Residential Min Lot Sizes	0.04	0.47
For a typical new multi-family building project in this jurisdiction, how many distinct governing bodies or agencies must give mandatory approval before construction can begin?	0.02	0.02
How many mandatory steps are involved in the approval process for a typical new multi-family building?	0.02	0.10
Is multi-family housing listed as allowed through conversion (of either single family homes or non residential buildings)?	-0.10	-0.00
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?	-0.14	0.03
Is multi-family housing allowed, either by right or special permit (including through overlays or cluster zoning)?	-0.16	0.34
Are apartments above commercial (mixed use) allowed in any district?	-0.22	0.28
Are attached single family houses (townhouses, 3+ units) listed as an allowed use (by right or special permit)?	-0.24	0.32
Are accessory or in-law apartments allowed (by right or special permit) in any district?	-0.30	-0.03
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by special permit?	-0.32	0.19

Note: Prior to performing principal component analysis, all variables were normalized into z-scores. Missing data, where the model output "I don't know," were imputed with k-nearest neighbors. Additionally, each variable was expressed in terms of its expected univariate association with stricter zoning policies, such that more positive values indicate a greater degree of restrictiveness. For example, the variable representing the allowance of multi-family housing was inverted, so that a more positive value indicates that multi-family housing is not permitted, while a more negative value suggests that it is not.

Table A3: Variables, Sources, and Definitions

Variable	Source	Definition
Foreign Born Share	2022 American Community Survey	The percentage of the population that is foreign-born (B05002_013E / B05002_001E).
Median Household Income	2022 American Community Survey	The median income of all households (B19013_001E).
Share Population 65 and Over	2022 American Community Survey	The percentage of the population aged 65 and over (B01001_020E to B01001_025E and B01001_044E to B01001_049E / B01001_001E).
Median Gross Rent	2022 American Community Survey	The median gross rent for rental units (B25064_001E).
Median Home Value	2022 American Community Survey	The median value of owner-occupied housing units (B25077_001E).
Share Units Owner Occupied	2022 American Community Survey	The percentage of housing units that are owner-occupied (B25003_002E / B25003_001E).
Share Population Under 18	2022 American Community Survey	The percentage of the population under 18 years old (B01001_003E to B01001_006E and B01001_027E to B01001_030E / B01001_001E).
White Share	2022 American Community Survey	The percentage of the population identifying as White (B02001_002E / B02001_001E).
Poverty Rate	2022 American Community Survey	The percentage of the population living below the poverty line (B17001_002E / B17001_001E).
College Share	2022 American Community Survey	The percentage of the population aged 25 and over with a bachelor's degree or higher (B15003_022E, B15003_023E, B15003_024E, B15003_025E / B15003_001E).
Share Structures Built Before 1970	2022 American Community Survey	The percentage of housing structures built before 1970 (B25034_008E, B25034_009E, B25034_010E, B25034_011E / B25034_001E).
Share Structures with 2 or More Units	2022 American Community Survey	The percentage of housing structures with 2 or more units (B25024_004E to B25024_009E / B25024_001E).

Continued on next page

Table A3 continued from previous page

Variable	Source	Definition
Vacancy Rate	2022 American Community Survey	The percentage of vacant housing units (B25002_003E / B25002_001E).
Share with Commute Over 30 Minutes	2022 American Community Survey	The percentage of workers with a commute time over 30 minutes (B08303_008E to B08303_013E / B08303_001E).
Housing Unit Density	2022 American Community Survey	The number of housing units in a local government divided by the area from its shape file.
Share Units Affordable	2022 American Community Survey	The percentage of housing units affordable to households earning the state median income. This measure combines rental and owner-occupied housing affordability, determined using the state median income. Rental units are affordable if the monthly rent does not exceed 30% of the monthly median household income, and owner-occupied units are affordable if their value is less than three times the annual median household income. The total number of affordable rental and owner-occupied units is summed and divided by the total number of housing units to determine the share of units that are affordable.
Job Density	Opportunity Insights	Number of jobs per square mile in each census tract in 2013. 'job_density_2013' from the Opportunity Atlas neighborhood characteristics dataset.
Opportunity Index	Opportunity Insights	The kid family rank, a measure of economic mobility.
Average Math Test Scores	The Education Opportunity Project at Stanford University	The average math test score pooled across grades (3rd-8th) and years (2008-2019) (cs_mn_avg_mth_ol).
Math Learning Rate	The Education Opportunity Project at Stanford University	The slope of the increase in math test scores from 3rd to 8th grade pooled across years (2008-2019) (cs_mn_grd_mth_ol).

Continued on next page

Table A3 continued from previous page

Variable	Source	Definition
Percent Eligible for Free Lunch	The Education Opportunity Project at Stanford University	The percentage of students eligible for free lunch (perflu).
Property Tax Rate	The Government Finance Database	The property tax rate is calculated as the total property tax revenue (Property_Tax_2017) divided by the aggregate home value from the 2017 ACS.
Total Expenditure Per Capita (2017)	The Government Finance Database	Total expenditures per capita is calculated as the total government expenditures (Total_Expenditure_2017) divided by the population (From 2017 ACS) of the municipality or township.
Building Permits All Units 2021	Building Permits Survey	The number of housing units permitted in 2021 divided by the population of the local government.
Year of Incorporation	Goodman, C. B. (2023). Municipal Incorporation Data, 1789-2020 (Version 1.0.0) [Computer software]. https://github.com/cbgoodman/municipal-incorporation/	The year a municipality was incorporated. Not available for townships.
Percent Democrat	"U.S. Voting by Census Block Groups", Bryan, Michael	The share of votes that are democrat in 2021.
New Housing Unit Elasticity	Nate Baum and Lu Han (2024)	The coefficient on the change in new housing units to the change in housing prices, specifically $\gamma_{01b_newunits_FMM}$.

Continued on next page

Table A3 continued from previous page

Variable	Source	Definition
Log Land Area	2022 Census Shapefiles	The area in acres of a local government with a log transform.
Log Neighbors within 25 Miles	2022 Census of Governments	The number of other local governments within 25 miles of a local government.
Miles to Metro Center	2022 Census Shapefiles	The number of miles from the centroid of a local government's shape file to the center of a metropolitan area, defined as the city hall of the center city.

Appendix: Question Details

This appendix provides detailed information about each question used in the study. Each question is presented with its original phrasing by the Pioneer Institute, the text that we embed for the question, background information and assumptions, question type, and the rephrased question that the language model sees. For some questions, we also include a value that triggers double-checking if the model's answer does not match it, along with the rephrased question used for double-checking and the keywords used to build context during the double-checking process. Additionally, certain questions involve subtasks, which are described in detail.

Question 4

Question Phrased by Pioneer: Is multi-family housing allowed, either by right or special permit (including through overlays or cluster zoning)?

Question Text That We Embed: Is multi-family housing allowed, either by right or special permit (including through overlays or cluster zoning)?

Question Background and Assumptions: Multi-family housing comes in a wide variety of forms and sizes. The ways municipalities define and categorize "multi-family" housing varies widely, as do the use-regulations that govern multi-family housing development. This study includes as "multi-family" any building with three or more dwelling units. Multi-family dwelling units can be rental or condominium. They can be in a freestanding residential building or part of a mixed-use building, new construction or conversion of a preexisting building. Zoning documents usually specify what kinds of buildings qualify for conversion to multi-family housing: single family houses, two family houses, mills, schools, churches, municipal buildings or other types of facilities. Freestanding new "Multi-family" housing is defined as any building with three or more dwelling units, excluding townhouses, unless a municipality includes townhouses in its broader definition of multi-family housing and effectively permits only townhouses as such. Assisted living facilities, congregate care homes, dormitories, and lodging houses are not considered multi-family housing. If the zoning laws allow for conversion to multi-family housing, but do not comment on whether new multi-family housing is allowed, then the answer is 'YES'. Most towns allow a form of multi-family housing.

Question Type: Binary

Rephrased Question the LLM Sees: Is multi-family housing allowed at all in any district or overlay? If multi-family housing is allowed by special permission in any district or overlay then that counts allowed.

Question 5

Question Phrased by Pioneer: Are apartments above commercial (mixed use) allowed in any district?

Question Text That We Embed: Are apartments above commercial (mixed use) allowed in any district?

Question Background and Assumptions: Zoning bylaws and ordinances in various municipalities often contain provisions for combining residential dwellings with commercial uses such as retail or office spaces, creating mixed-use developments. While some zoning regulations explicitly allow multi-family housing and retail to coexist within the same district, they may not clarify whether these uses can share the same building, leaving this to be determined in practice. Certain municipalities explicitly permit "combined dwelling/retail" configurations in their use regulation tables, sometimes noting that any uses allowed within the same district can occupy the same building. Additionally, detailed provisions for mixed-use are facilitated through special zoning arrangements like overlay districts (e.g., mixed use district, downtown overlay, or planned unit development) or conversion projects, such as transforming former mills to accommodate both retail and housing. However, it's important to note that some references to "mixed use" may actually pertain to commercial and industrial combinations, excluding residential components. If you cannot find any reference to residential and commercial uses in the same building within the context then you assume that the answer is 'NO'.

Question Type: Binary

Rephrased Question the LLM Sees: Is a combination of commercial and residential uses in the same building or structure allowed in any zoning district?

Question 6

Question Phrased by Pioneer: Is multi-family housing listed as allowed through conversion (of either single family homes or non residential buildings)?

Question Text That We Embed: Is multi-family housing listed as allowed through conversion (of either single family homes or non residential buildings)?

Question Background and Assumptions: The development of multifamily housing through the conversion of existing buildings encompasses two primary approaches: transforming single-family or two-family houses into structures with at least three units, and repurposing non-residential buildings, such as mills, other industrial buildings, schools, and municipal buildings, for multi-family residential use. This is different from the ability to construct new multi-family housing. The conversion of

non-residential structures often occurs through designated overlay districts, like Mill Conversion Overlay Districts, or within industrial zones, whereas the conversion of houses to accommodate more units typically takes place in residential or business districts. The question does not count the conversion of single-family homes into two-family dwellings as allowing conversion to multi-family dwellings because multi-family is defined as having at least three units. If the conversion requires a special permit then we consider that as allowing conversion. Assisted living facilities, congregate care homes, dormitories, and lodging houses are not considered multi-family housing. The allowance of multi-family housing does not imply the allowance of the conversion to multi-family housing. You must search for an explicit statement allowing the conversion to multi-family housing from another type of structure. If you do not find any mention of conversions in the context then you assume the answer is 'NO'.

Question Type: Binary

Rephrased Question the LLM Sees: In any district, is the conversion to multi-family explicitly allowed under any scope?

If The Answer Is Not This Value Then We Double Check: Yes

Rephrased Question the LLM Sees When Double Checking: In any district, is the conversion to multi-family explicitly allowed under any scope?

Keywords We Use to Build Context When Double Checking in Order of Importance: 'conver'

Question 8

Question Phrased by Pioneer: Are attached single family houses (townhouses, 3+ units) listed as an allowed use (by right or special permit)?

Question Text That We Embed: Are attached single family houses (townhouses, 3+ units) listed as an allowed use (by right or special permit)?

Question Background and Assumptions: The question asks whether some form of attached housing is allowed in the municipality. Common forms of attached housing are single-family attached homes, townhouses, rowhouses, and zero lot line dwelling units. Attached housing is often allowed through special zoning provisions, such as overlay districts or use provisions tailored for cluster developments, Planned Unit Developments (PUD), or communities for active adults aged 55 and over. Remember that accessory apartments to a single-family home or the ability to attach one unit to a single-family home do not count as attached housing. Duplexes also do not count as attached housing. A form of attached housing may be listed as a type of single-family or multi-family housing. However, the allowance of

single-family or multi-family housing does not imply the allowance of attached housing. This context raises the question of whether any type of attached housing are allowed either as their own category of housing or explicitly as a type of single family or multi-family housing. If you do not find any mention of a type of attached housing in the context then you assume that the answer is 'NO'.

Question Type: Binary

Rephrased Question the LLM Sees: Is some form of attached housing allowed in any district of the town?

If The Answer Is Not This Value Then We Double Check: Yes

Rephrased Question the LLM Sees When Double Checking: Is some form of attached housing allowed in any district of the town?

Keywords We Use to Build Context When Double Checking in Order of Importance: 'town house', 'town houses', 'townhouse', 'townhouses', 'attached dwelling', 'attached dwellings', 'row house', 'row houses', 'rowhouse', 'rowhouses', 'attached single family', 'attached unit', 'attached units', and 'attached'

Question 9

Question Phrased by Pioneer: Does zoning include any provisions for housing that is restricted by age?

Question Text That We Embed: Does zoning include any provisions for housing that is restricted by age?

Question Background and Assumptions: Many zoning bylaws/ordinances include provisions for housing that is deed restricted to occupants 55 (or another age) and older. Some of the provisions are for developments that are entirely age-restricted, while other provisions are incentives, often density bonuses, to include age-restricted units within an unrestricted development, such as cluster or multi-family. The restricted developments are called active adult housing, adult retirement village, senior village, planned retirement community, or something similar.

The answer should be Yes if any provisions exist for age-restricted single-family, townhouse, duplex, multi-family or accessory apartments. Provisions can be in the form of an age-restricted overlay, cluster development, density bonus for age-restricted units, or other zoning requirements or incentives for age-restricted housing.

Question Type: Binary

Rephrased Question the LLM Sees: Does zoning include any provisions for housing that is restricted by age?

Question 11

Question Phrased by Pioneer: Are accessory or in-law apartments allowed (by right or special permit) in any district?

Question Text That We Embed: Are accessory or in-law apartments allowed (by right or special permit) in any district?

Question Background and Assumptions: Accessory dwellings are separate housing units typically created in surplus or specially added space in owner-occupied single-family homes. Accessory dwellings can also be attached to the primary dwelling or be situated on the same lot (for example in a carriage house or small cottage.) An accessory dwelling typically has its own kitchen and bathroom facilities, not shared with the principal residence. Many zoning bylaws/ordinances call the dwellings “in-law apartments” or “family apartments” and restrict their occupancy to relatives of the homeowner - “related by blood, marriage or adoption.” Some of these also allow domestic employees, caregivers, elderly people or people with low incomes to live in the units. Some municipalities allow the apartment by right if a family member will occupy the accessory apartment, but require a special permit otherwise. If you cannot find any reference to accessory apartments in the context then you assume that the answer is ‘NO’.

Question Type: Binary

Rephrased Question the LLM Sees: Are accessory or in-law apartments allowed in any district? If they are allowed by special permit in any district then we count that as allowed.

Question 13

Question Phrased by Pioneer: Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?

Question Text That We Embed: Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?

Question Background and Assumptions: Flexible zoning, encompassing terms like open space residential design, cluster, planned unit development, or conservation subdivision, provides municipalities with a more adaptable approach to zoning beyond the traditional “as-of-right” options. This methodology allows developers to bypass the stringent requirements of standard zoning, such as specific lot sizes and setback mandates, and enables the incorporation of various residential unit types like townhouses, duplexes, and multi-family homes that might not be allowed under conventional zoning

regulations. The question only considers provisions that are primarily for residential uses. Most municipalities require special permits for cluster/flexible development.

Question Type: Binary

Rephrased Question the LLM Sees: Is the answer yes to any of the following question? Question 1: Is cluster development allowed explicitly by right in any district? Question 2: Is open space residential design allowed explicitly by right in any district? Question 3: Is any type of flexible zoning other than cluster development and open space residential design allowed explicitly by right in any district?

Question 14

Question Phrased by Pioneer: Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by special permit?

Question Text That We Embed: Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by special permit?

Question Background and Assumptions: Flexible zoning, encompassing terms like open space residential design, cluster, planned unit development, or conservation subdivision, provides municipalities with a more adaptable approach to zoning beyond the traditional “as-of-right” options. This methodology allows developers to bypass the stringent requirements of standard zoning, such as specific lot sizes and setback mandates, and enables the incorporation of various residential unit types like townhouses, duplexes, and multi-family homes that might not be allowed under conventional zoning regulations. The question only considers provisions that are primarily for residential uses. Most municipalities require special permits for cluster/flexible development so if you find suggestive evidence that the municipality allows cluster/flexible development by special permit then you assume that the answer is ‘YES’.

Question Type: Binary

Rephrased Question the LLM Sees: Is the answer yes to any of the following question? Question 1: Is cluster development allowed in any district, including by special permit? Question 2: Is open space residential design allowed in any district, including by special permit? Question 3: Is any type of flexible zoning other than cluster development and open space residential design allowed in any district, including by special permit?

Question 17

Question Phrased by Pioneer: Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?

Question Text That We Embed: Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?

Question Background and Assumptions: Inclusionary zoning requires or encourages developers to include affordable dwelling units within new developments of market rate homes. Some municipalities call it “incentive zoning” - when provision of affordable units is voluntary. The affordable units are typically located on site, but some municipalities also allow off-site development under certain circumstances. Often, payments may be made to a trust fund in lieu of building housing. Housing designated as “affordable” must be restricted by deed or covenant, usually for a period of 30 or more years, to residents with low or moderate incomes. The deed restrictions also limit sales prices and rents as the units are vacated, sold or leased to new tenants.

Do not include provisions for entirely affordable, subsidized housing development by public or non-profit corporations. Also do not include provisions under “rate of development” headings that exempt affordable units from project phasing and growth caps.

Question Type: Binary

Rephrased Question the LLM Sees: Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?

Question 20

Question Phrased by Pioneer: Is there a town-wide annual or biannual cap on residential permits issued, and/or is project phasing required?

Question Text That We Embed: Is there a town-wide annual or biannual cap on residential permits issued, and/or is project phasing required?

Question Background and Assumptions: Some municipalities enact town-wide caps limiting the number of units that can come on line annually or biannually. The number of permits is often set at the average in the previous years. Note that this question asks only about town-wide caps and does not consider caps exclusive to a specific district in the town. Some municipalities require phased growth for individual developments (also known as development scheduling or buildout scheduling) - a technique that allows

for the gradual buildout of approved subdivisions over a number of years. Note that we only consider project phasing when it is required and not when it is optional. Project phasing is usually triggered by a minimum number of units in the project, so small subdivisions can be constructed in one year. Some phasing provisions are only triggered at the town-wide level once a threshold number of units have been permitted. Most of the “rate of development” provisions include an expiration or “sun set” date (some that have expired have been updated and re-adopted). Many include a “point system” where points are awarded for provision of community goods such as open space or affordable units, and projects with more points are given priority for permits. If you do not find any information in the context about a town-wide annual or biannual cap or about project phasing then you assume the answer is ‘NO’.

Question Type: Binary

Rephrased Question the LLM Sees: Is the answer yes to any of the following question? Question 1: Is there a town-wide annual or biannual cap on residential permits issued Question 2: Is project phasing required?

Question 21

Question Phrased by Pioneer: Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?

Question Text That We Embed: How is lot area defined and how is the lot size calculated?

Question Background and Assumptions: Remember to first review your research so far on how a lot size is calculated and defined. If you have already found a restriction on including wetlands, sloped land, or easements in your prior research then the answer is ‘YES’.

Some municipalities require that the minimum lot size requirement be met by a percentage of land that does not include wetland resource areas, steeply sloped land or easements. A subset of those municipalities requires that the buildable area be contiguous on the lot – called “contiguous buildable area” or “contiguous upland area.” Upland area is non-wetland area. It is much more common for municipalities to restrict the use of wetlands areas in meeting lot size requirements than sloped land or easements.

Note that this question only asks about whether there are restrictions on calculating the lot size. It does not ask about whether there are restrictions to buildable area or whether there are any restrictions in wetland areas.

If you do not find any restrictions for lot size calculations in the context then you assume that the answer

is 'NO'.

Question Type: Binary

Rephrased Question the LLM Sees: Detail how lot area is defined and how a lot size is calculated. Then, answer the question of are there restrictions on counting wetlands, uplands, or sloped land in lot area/lot size calculation?

If The Answer Is Not This Value Then We Double Check: Yes

Rephrased Question the LLM Sees When Double Checking: Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?

Keywords We Use to Build Context When Double Checking in Order of Importance: 'wetland', 'upland', 'sloped land', and 'easement'

Question 27

Question Phrased by Pioneer: What is the minimum lot size for each zoning district?

Question Text That We Embed: What is the minimum lot size for each zoning district?

Question Background and Assumptions: The question asks to provide a list of the minimum lot size in each district of the town. If a district has different minimum lot sizes depending on the type of building like for example a different minimum lot size for single-family homes than for multi-family homes, then you pick the smaller of the minimum lot sizes. If a district allows smaller minimum lot sizes for historic properties or by special permission then you pick the standard minimum lot size for current buildings. If a district only lists a minimum lot size for a specific type of housing like housing for the elderly, then you pick that minimum lot size. Your answer should be structured as a list with district name, minimum lot size, and units for the minimum lot size which are usually square feet or acres. If a minimum lot size for a district is reported in both acres and square feet then only report it in square feet. If a district does not have a minimum lot size then record the town wide minimum lot size for that district if a town wide minimum lot size exists. If a town wide minimum lot size does not exist and a district does not have a minimum lot size then exclude it from your answer.

Question Type: Lot Size

Rephrased Question the LLM Sees: What is the minimum lot size for each zoning district?

Subtask:

- Subtask Question That Gets Embedded: List out each district in the town

- Rephrased Subtask Question the LLM Sees: List out each district in the town
- Additional Subtask Instructions: Please list out the name of each district in the town. Do not include overlay districts.
- How The Subtask Results Are Described to the LLM Afterwards: List of all districts to find the minimum lot size for

Question 28

Question Phrased by Pioneer: What is the minimum lot size for single-family homes in each residential district?

Question Text That We Embed: What is the minimum lot size for single-family homes in each residential district?

Question Background and Assumptions: When compiling a list of minimum lot sizes for districts that permit single-family housing, prioritize clarity by selecting the specific minimum lot size for single-family homes within each district. If multiple options exist, choose the most common standard size, excluding sizes for historic properties or special cases. Report sizes in square feet over acres unless only acre measurements are available. Only include districts with a defined minimum lot size or those adhering to a town-wide minimum if no district-specific size is established. Finalize the data in a CSV format with columns for 'District Name', 'Min Lot Size', 'Unit', and 'Estate', ensuring a straightforward, single entry for each district that reflects the standard requirement for single-family homes.

Question Type: Lot Size

Rephrased Question the LLM Sees: What is the minimum lot size for single-family homes in each residential district?

Subtask:

- Subtask Question That Gets Embedded: Find the name of each district that allows single-family housing
- Rephrased Subtask Question the LLM Sees: Find the name of each district that allows single-family housing
- Additional Subtask Instructions: Please list out the name of each residential district in the town that primarily consist of detached single-family housing. If you cannot find any districts that explicitly

allow single-family detached housing then just assume that any residential districts allow single-family detached housing. Respond with a detailed answer followed by a CSV format with the name of the district in the first column and whether a district has the label 'Estate' in the second column as a True/False statement. Use the column headers of 'District Name' and 'Whether Estate District'.

- How The Subtask Results Are Described to the LLM Afterwards: Your previous work finding which districts to find minimum lot sizes for and whether they are estate districts

Question 2

Question Phrased by Pioneer: How many zoning districts, including overlays, are in the municipality?

Question Text That We Embed: How many zoning districts, including overlays, are in the municipality?

Question Type: Numerical

Rephrased Question the LLM Sees: How many zoning districts and overlays are in the municipality?

Question 22

Question Phrased by Pioneer: What is the longest frontage requirement for single family residential development in any district?

Question Text That We Embed: What is the longest frontage requirement for single family residential development in any district?

Question Type: Numerical

Rephrased Question the LLM Sees: What is the longest frontage requirement for single family residential development in any district?

Subtask:

- Subtask Question That Gets Embedded: Find the name of each single-family residential district
- Rephrased Subtask Question the LLM Sees: Find the name of each single-family residential district
- Additional Subtask Instructions: Please list the names of each single-family residential district. Only include districts that are primarily residential. Usually, this means districts that start with the

letter R like R1. If there is only one residential district that permits single-family zoning then just name that one district. If you are unsure whether a residential district permits single-family zoning then assume that it does, but ensure that the district is primarily residential. An agricultural (A) or industrial (I) district would not be included for example.

- How The Subtask Results Are Described to the LLM Afterwards: Only consider the frontage requirements in the following districts

Question 17w

Question Phrased by Pioneer: Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?

Question Text That We Embed: Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?

Question Background and Assumptions: Zoning codes may require developers to include affordable housing in market-rate residential projects, but the applicability of these requirements can vary. Some inclusionary policies apply broadly to all residential development, while others are tied to optional zoning designations, incentive programs, or specific areas.

To determine if a zoning code contains a mandatory inclusionary requirement, look for clear language stating that all or most market-rate residential projects must provide affordable units as a standard condition of approval under normal zoning rules. The requirement should not be limited to projects that opt into a special zoning designation, participate in an incentive program, or are located in a particular overlay zone.

Focus on whether the code unambiguously requires all or most market-rate residential development to include affordable housing under the generally applicable rules. Do not select "YES" if affordable housing is only mandatory in narrow, specialized situations. The mere presence of affordable housing provisions is not sufficient if they are elective or only apply in atypical circumstances. If the affordable housing requirements are not clearly universally applicable, the likely answer is "NO".

Question Type: Binary

Rephrased Question the LLM Sees: Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?

Question 30

Question Phrased by Pioneer: How many mandatory steps are involved in the approval process for a typical new multi-family building?

Question Text That We Embed: How many mandatory steps are involved in the approval process for a typical new multi-family building?

Question Background and Assumptions: The approval process for constructing a new multi-family building typically involves multiple mandatory steps, each representing a distinct interaction or requirement that a developer must fulfill before construction can begin. Focus on identifying only the core, pre-construction approval steps that are required for all multi-family building projects, from initial application submission to final permit issuance. Each required interaction with a distinct city department or agency should be counted as a separate step, but be careful not to artificially separate closely related actions within a single process. For example, applying for and obtaining a building permit should be considered one step, not two. Be cautious not to include optional or discretionary steps, post-approval activities such as inspections during construction or certificate of occupancy issuance, steps that are only required in specific circumstances or for certain types of properties, or internal processes within departments that don't require direct developer interaction. When analyzing the ordinances, pay close attention to language indicating whether a step is mandatory (e.g., "shall", "must", "is required") versus optional or conditional (e.g., "may", "at the discretion of", "if applicable"). The goal is to identify the minimum number of distinct, mandatory steps that every multi-family building project must go through in the approval process, avoiding redundancy and over-segmentation of closely related actions.

Question Type: Numerical

Rephrased Question the LLM Sees: How many mandatory steps are involved in the approval process for a typical new multi-family building?

Question 31

Question Phrased by Pioneer: For a typical new multi-family building project in this jurisdiction, how many distinct governing bodies or agencies must give mandatory approval before construction can begin?

Question Text That We Embed: For a typical new multi-family building project in this jurisdiction, how many distinct governing bodies or agencies must give mandatory approval before construction can begin?

Question Background and Assumptions: When answering this question, focus on the approval process

for a typical new multi-family building project as described in the provided ordinance sections. Only count distinct governing bodies or agencies whose approval is explicitly required by the ordinances for all multi-family building projects, including those allowed "by right" under existing zoning. To be counted, an entity must have clear, independent approval authority that is mandatory for the project to proceed. This approval must be specifically for the multi-family project itself. Look for unambiguous language indicating required, independent approval steps. Distinguish between actual approval authority and advisory roles; entities that only review or provide input should not be counted. Consider roles like the Planning Board, Board of Health, Building Commissioner, and special permit granting authorities, but include them only if their approval is explicitly required and independent. For coordinated review processes, determine whether they represent multiple independent approvals or a single approval incorporating multiple inputs. Provide your answer as a number, followed by a brief explanation of which entities you counted and why. Cite relevant ordinance sections, explaining why each approval is considered independent and mandatory, and how it relates specifically to the multi-family project.

Question Type: Numerical

Rephrased Question the LLM Sees: For a typical new multi-family building project in this jurisdiction, how many distinct governing bodies or agencies must give mandatory approval before construction can begin?

Question 32

Question Phrased by Pioneer: Are there townwide requirements for public hearings on any type of multi-family residential projects?

Question Text That We Embed: Are there townwide requirements for public hearings on any type of multi-family residential projects?

Question Background and Assumptions: When answering this question, examine the zoning ordinances and bylaws for any townwide requirements that mandate public hearings or formal public input processes for multi-family residential developments. Focus on requirements that apply across all zones within the town. Answer YES if public hearings are required for any subset of multi-family projects, even if not all multi-family projects require hearings. For instance, if larger projects require public hearings while smaller ones don't, the answer should still be YES. Requirements specific to certain zones do not count towards a YES answer. Answer NO only if there are no townwide public hearing requirements for multi-family developments of any size or type, or if such requirements only apply in specific zones. Be sure to cite

relevant ordinance sections that support your conclusion. The goal is to determine whether there is any mandated opportunity for public input on new multi-family housing developments on a townwide basis, even if this only applies to certain categories of multi-family projects.

Question Type: Binary

Rephrased Question the LLM Sees: Are there townwide requirements for public hearings on any type of multi-family residential projects?

Subtask:

- Subtask Question That Gets Embedded: Do any types of multi-family housing projects require a special permit in this jurisdiction? If so, under what conditions?
- Rephrased Subtask Question the LLM Sees: What is the typical approval process for new multi-family building projects in this jurisdiction? Please describe any required permits, reviews, or other procedures that are standard for multi-family developments.
- Additional Subtask Instructions: Do any types of multi-family housing projects require a special permit in this jurisdiction? If so, under what conditions?
- How The Subtask Results Are Described to the LLM Afterwards: Special Permit Requirements for Multi-Family Housing Developments

Question 34

Question Phrased by Pioneer: What is the maximum potential waiting time (in days) for government review of a typical new multi-family building?

Question Text That We Embed: What is the maximum potential waiting time (in days) for government review of a typical new multi-family building?

Question Background and Assumptions: The review process for constructing a new multi-family building involves several stages, each of which may have a specific waiting period. The total waiting time includes the mandatory review periods as well as any discretionary days that can be added by the governing bodies or agencies. Each agency or department that a developer must interact with, such as city government departments like fire, police, sanitation, building, and planning, has its own review timeline. Additionally, discretionary days that may be required for public hearings, environmental reviews, or other procedural requirements must also be added to the total count of government review days.

Question Type: Numerical

Rephrased Question the LLM Sees: What is the maximum potential waiting time (in days) for government review of a typical new multi-family building?