The Costs of Housing Regulation: Evidence From Generative Regulatory Measurement

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

February 12, 2025

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 0.92 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) Zoning follows a monocentric pattern with regional variations, with suburban regulations particularly strict in the Northeast. (3) Housing regulations can be clustered into two main principal components. (4) The first of which corresponds to *value capture*, indicating how municipalities extract economic benefits in areas of high housing demand. (5) The second principal component associates with *exclusionary zoning*, resulting in higher housing costs and socioeconomic exclusion. We develop a model of non-cooperative muincipal government regulatory choice consistent with these facts.

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 Elliott Ash, Leah Brooks, Sophie Calder-Wang, Tom Cui, Nestor Davidson, Tatyana Deryugina, Giles Duranton, Ingrid Gould Ellen, Chris Elmendorf, Edward Glaeser, Joseph Gyourko (discussant), John Horton, Greg Howard, Ben Hyman, Jeffrey Lin, Stijn Van Nieuwerburgh, Theresa Kuchler, Adam Osman, Franklin Qian, Julian Reif, Daniel Rock, Albert Saiz, Gregor Schubert, Allison Shertzer, Jaehee Song (discussant), Lena Song, Johannes Stroebel, Bryan Stuart, and comments from conference and seminar participants at NBER Summer Institute (Real Estate/Urban), 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, CREED Conference, UIUC economics, the Metric of Regulation, the New York Fed, and the UEA. Preeyonuj Boruah, David Dai, Amber Gao, Grace Getman, Drew Harrington, Karin Hobelsberger, Nagharjun Mariappan, Scott Overbey, Alok Ranjan, Saul Richardson, Ritul Soni, Luofeng Zhou, and Kevin Zhou provided excellent research assistance. We gratefully acknowledge funding from Emergent Ventures and the NYU Stern Center for Global Economy and Business, and thank the Pioneer Institute for providing data. 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 regulations constitute the foundational legal environment governing the built environment of American cities, dictating the form of residential development and housing affordability (Glaeser and Gyourko, 2018; Gyourko et al., 2008). The influence of zoning laws and land use policies extends far beyond the real estate markets, with implications for social segregation, economic mobility, environmental sustainability, the growth of urban agglomerations, and construction sector productivity (Gyourko and Molloy, 2015; Kahn, 2000; Hsieh and Moretti, 2019; D'Amico et al., 2023). Despite their critical importance for such diverse outcomes, accurately measuring housing regulations remains a challenge due to the complexity and variation of municipal ordinances. These measurement gaps hinder our ability to understand the fundamental drivers of housing regulation and their broader impacts.

Our paper argues that advances in Large language Models (LLMs) enable scalable and accurate classification of regulatory documents, a task which we refer to as *generative regulatory measurement*. We obtain municipal codes for 63% of the population governed by local zoning ordinances, and use state-of-the-art Artificial Intelligence (AI) tools to estimate housing regulation on the full text of these documents. Our core methodological innovation lies in developing a comprehensive approach to regulatory analysis using LLMs which combines Retrieval Augmented Generation (RAG), semantic embeddings, multi-step processing, and carefully engineered prompts. We apply this approach to a set of regulatory questions initially developed by the Pioneer Institute for Massachusetts (Glaeser and Ward, 2009), and benchmark our LLM-generated regulatory categorizations against the human-coded measurements from this same Pioneer Institute study.

Our results indicate that LLMs deployed on frontier models (GPT-4 Turbo by OpenAI) have achieved near-human rates of precision in classifying regulation, with an accuracy rate of 96% for binary questions. LLMs also perform strongly on numerical questions with a correlation of 0.92 for residential lot size minimums, an important regulation governing density in single-family areas. We also manually verify a subset of housing regulations in California to ensure our results are not geographically biased.

We use the resulting LLM-produced dataset on national housing regulations, alongside other

housing data, to establish five key facts about land use policies and their impacts across United States. First, zoning ordinances appear to be broadly binding for the new development of housing across the United States. We highlight a key distinction between incorporated areas—areas with established local governments often formed to implement stricter housing regulations—and unincorporated areas. Over the period from 2000–2020, we find that 42% of the growth in new housing units in the U.S. was concentrated in unincorporated areas, including a majority of new housing units in the South. This is despite the fact that unincorporated areas are home to only 23% of the overall population, are typically low-demand regions with low prices and rents, face high vacancy rates, are far from municipal job centers and amenities, and expose residents to greater physical climate risks. We verify the role of regulations in driving this pattern through a discontinuity approach, suggesting higher housing supply just outside incorporated areas. The heavy reliance on greenfield development in far exurban locations for a substantial portion of housing production indicates significant barriers to infill development within incorporated areas. This pattern suggests that current zoning regulations in established municipalities limit housing supply in high-demand areas.

Our second key finding is that housing regulation varies within metropolitan areas in ways that are broadly consistent with a monocentric city model (Alonso, 1964; Muth, 1971), while also highlighting significant deviations. As standard monocentric city models predict, denser building is generally allowed in city centers, with stricter bulk regulations and lower density requirements found in inner-ring suburbs. However, we also observe substantial sorting of high-income households into expensive suburbs with higher minimum size requirements. This is in contrast to the most basic models in which all households, regardless of income, are assumed to prefer central locations to minimize commuting costs. Our findings suggest a more complicated spatial pattern, with some peripheral suburban locations able to maintain high prices and sorting by affluent households in areas with stricter zoning. We find that this pattern is particularly pronounced in the Northeast, which has substantially more onerous bulk regulation requirements than other regions of the country. This regional variation suggests that historical development patterns, local political economies, and urban path dependencies play a significant role in shaping contemporary zoning practices, alongside standard economic forces of agglomeration and transportation costs.

Third, we isolate two key dimensions of housing regulation that have distinct impacts on housing markets. The first principal component of our regulatory questions is associated with high prices and high construction, suggesting regulations that allow development but exact value in high-demand environments. In contrast, the second principal component associates with high prices but low construction, indicating regulations that restrict housing supply. Our ability to measure a more detailed, granular, and comprehensive set of regulations identifies distinct regulatory regimes which are associated with divergent housing market outcomes, in contrast to traditional unidimensional models of regulation (Gyourko et al., 2021).

Fourth, we support our interpretation of the first principal component in its role of value capture in high-demand environments. This component loads heavily on regulations that allow local governments to extract and redistribute housing surplus, such as inclusionary zoning mandates that require developers to include affordable units. These regulations are typically found in densely populated, centrally located cities that tend to support Democratic political candidates. Rather than completely blocking all development, these regulations impose implicit taxes on new construction to fund local amenities and redistribution.

Fifth, we connect the second principal component to exclusionary zoning practices aimed at limiting density and affordable housing options. This dimension relies heavily on bulk regulations, such as minimum lot size requirements, as well as procedural barriers to development. Among single-family zoned areas in municipalities across the country, 66% have town-wide 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. These tools are particularly prevalent in affluent, predominantly white suburban areas that tend to lean Republican relative to other areas in their metropolitan area. We suggest a range of potential motivations for exclusionary zoning practices, including concerns about spillovers from low-income residents. In particular, we find a strong association between these regulatory measures and indicators of local school performance and social mobility, suggesting a role for educational sorting. Unlike value capture regulations that allow development with conditions, these exclusionary practices directly restrict housing supply, constraining affordable housing supply in high-demand areas. These regulations therefore intensify economic and racial segregation by effectively pricing out lower-income and minority households

from neighborhoods with desirable amenities, particularly high-performing schools.

We interpret these facts through a model of inter-municipal competition in which local governments strategically and non-cooperatively select between two sets of housing regulations. Our framework builds on literature examining Tiebout sorting and local goods provision (Tiebout, 1956; Epple and Zelenitz, 1981), as well as traditional theories of zoning focused on housing sorting across public goods regimes (Fischel, 1987), Hamilton (1975), but generates distinct predictions about regulatory behavior which line up with our empirical results.

The model predicts that high-demand areas, typically metropolitan cores with high amenities and elevated housing costs, will implement value capture regulations. These regulations act as implicit taxes on housing development (similar to rent seeking behavior in other municipal contexts as in Diamond (2017)), with the seized housing surplus redistributed to residents. In contrast, the model predicts that suburban areas will adopt exclusionary regulations to limit perceived spillovers from lower-income residents. This leads to an equilibrium in which suburbs are home to higher-income residents, and urban cores feature a mix of both poor households as well as wealthy residents who value urban amenities enough to pay these implicit taxes (Brueckner, 1995, 2011). The model thus generates clear spatial predictions about regulatory patterns: value capture in high-demand urban cores and exclusionary zoning in suburbs, matching the empirical patterns we document.

The key economic contributions of our paper stem directly from our novel, detailed measurement of specific regulatory provisions across thousands of municipalities. This granular data reveals that housing regulations are fundamentally multidimensional rather than one-dimensional as measured in prior work. While previous research relied on aggregate indices or indirect measures of regulation, our direct measurement of specific provisions allows us to isolate and distinguish between value capture tools in urban cores (like inclusionary zoning mandates) and exclusionary practices in suburbs (like minimum lot size requirements). This detailed measurement also enables us to document how these regulatory approaches interact within metropolitan areas: urban regulations that raise housing costs push some residents toward suburbs, while suburban regulations maintain exclusivity through specific density restrictions. Second, by measuring actual zoning provisions rather than inferring them from outcomes, we provide direct evidence for how regulation suppresses density and displaces development to unincorporated areas. Our granular measurement of bulk regulations

like minimum lot sizes and multifamily restrictions quantifies these constraints precisely. Third, our detailed spatial mapping of specific zoning provisions challenges standard urban models—while broad zoning patterns follow monocentric predictions, our measurement of actual suburban land use restrictions helps explain why high-income households sort into strictly-zoned suburbs despite higher commuting costs.

Another key contribution of our paper is the development of a general-purpose methodology to reliably and precisely measure the content of textual documents. Our LLM-based approach achieves high accuracy rates, comparable to those of skilled non-experts, and offers several advantages for researchers. First, it provides unprecedented scalability at low cost: we successfully apply our regulatory classification measure across thousands of municipalities, a task that would be expensive and time-consuming for human analysts. This scalability opens up possibilities for comprehensive regulatory analysis across multiple domains. Second, our approach ensures verifiability and auditability by prompting the LLM to provide specific supporting text from the regulatory documents, enabling independent verification of classifications. Third, our approach is highly adaptable, allowing researchers to easily incorporate changes in definitions or incorporate advancements in AI models, facilitating replication and refinement of measurements over time.

The broad applicability of our approach extends to various domains where textual analysis is crucial, including building codes, tax regulations, legal cases, financial reports, newspapers, and other uses. This versatility is particularly valuable as the volume and complexity of regulations continue to increase (Singla, 2023). Our methodology therefore represents a significant advancement in the field of regulatory analysis and text-based research more broadly.

There are several caveats to consider regarding the results generated using our approach. First, LLMs require the user to source documents about zoning, which may be stale or not encompass the broader regulatory environment (i.e., interactions between state and local law). Second, our dataset only spans a portion of the questions that make up the regulatory landscape. Third, given the ambiguity and substantial heterogeneity of regulation across municipalities, our approach requires the LLM to make judgment calls when answering questions (i.e., choosing the more common type of single-family home minimum lot size requirements within a district). Fourth, our dataset reflects what the law says, but not how it is administered in practice (i.e., whether variances are liberally

granted). Finally, replicability remains a challenge for LLM-generated data, which we examine systematically. Despite these limitatios, our approach makes substantial progress in our ability to use LLMs for the systematic analysis of regulatory and administrative documents, and we suggest some best practices for researchers in this growing area.

Contributions to Literature The central contribution of our project is the creation of a standardized, comprehensive dataset of zoning across the United States, which we make publicly available for use by other researchers. Our approach towards regulatory measurement improves on existing approaches in accuracy, granularity of regulatory measurement, and in scope. One existing strand of the literature draws on survey-based approaches such as the Wharton Regulatory Index (Gyourko et al., 2008, 2021; Huang and Tang, 2012), which offer broad coverage on housing regulations, but are limited by low and potentially biased response rates, fixed questionnaires, and inaccuracy by respondents (Lewis and Marantz, 2019).

The second approach relies on other imputations of national housing regulation. Wedge-based approaches estimate housing regulations by examining the expected spatial macroeconomic distortions resulting from zoning. Examples in this literature include Hsieh and Moretti (2019), Glaeser et al. (2005), Herkenhoff et al. (2018), Babalievsky et al. (2021), and Duranton and Puga (2019). Other national approaches have examined textual data; for instance Ganong and Shoag (2017) use a scaled count of judicial decisions on "land use." 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 or tightening zoning restrictions and then analyze the effects of these changes in regulation on housing supply and rents. Mleczko and Desmond (2023) use a pre-LLM natural language processing (NLP) approach to measure a set of zoning regulations for the same sample of municipalities covered by the Wharton Regulatory Index (Gyourko et al., 2008).

The third strand in the literature uses more detailed analysis of specific regulations at more local levels. These include Shanks (2021) which 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). Similarly, (Lutz, 2015) studies the effects of property tax reforms on growth limitations and impact fees in New Hampshire. 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.... However, the enormity of that effort prevents it from being replicated in other markets by a single research team."

Existing research on housing regulations therefore leaves significant gaps in our understanding of their measurement and impacts. National studies identify broad impacts of regulations on housing costs and construction but lack specificity on key drivers, while more detailed state-level analyses are geographically limited and may not be nationally representative. We argue that the practical difficulties behind the scaling up of this approach have now been addressed through the development of modern AI LLMs, which provide both the comprehensiveness and granularity of the state-based approaches along with the scale of the national regulatory studies. Reflecting this breakthrough, our approach provides detailed measured of zoning regulations for over 5800 municipalities, more than twice as many municipalities as previous national studies such as the Wharton Regulatory Index (Gyourko et al., 2021), while providing the as much granularity or more as state and local studies such as the Pioneer Institute study (Glaeser and Ward, 2009).

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). 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. A growing literature has 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), as well as for social science hypothesis generation (Horton, 2023). Hoffman and Arbel (2023) argues for the use of LLMs in "generative interpretation" in estimating the meaning of legal contracts. A growing literature also examines broader implications of Generative AI (Eisfeldt et al., 2023; Brynjolfsson et al., 2023), as well as the role of algorithms applied to real estate (Calder-Wang and Kim, 2023; Raymond, 2023).

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.[1] 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 primary motive for incorporation.

Zoning, broadly, consists of two key sets of regulations: land use regulations, which partition local land into distinct use classes, and bulk regulations, whcih regulate the physical dimensions and density of buildings in different land use classes. Examples of bulk regulations include: minimum lot sizes (specifying the smallest allowed area for a buildable lot), lot coverage requirements, front and side setbacks (mandating minimum distances buildings must be from property lines), height restrictions, and floor area ratio caps (regulating total floor area relative to lot size). These regulations combine to regulate how densely areas can be developed, the size and spacing of buildings, and the overall built form of neighborhoods, and are the key focus of our analysis. Other mandates and requirements, such as parking minimums, further constrain both commercial and residential development in different areas.[2]

[1] In several states the "township" form of government also has jurisdiction in zoning which aligns with the Census County Subdivision definition.

[2] States and municipalities also enact building codes, which govern the building and safety standards that new

Municipalities enforce laws by issuing municipal codes which outline local regulation in different domains. 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 often aggregated by different companies online. Table 1 illustrates the breadth of our sample coverage. In total, we cover 25% of all municipalities in the U.S. and 6% of all townships. This coverage is skewed to larger cities, and so of the 76% of of the population in the U.S. that live in either a municipality or a township, we have relevant municipal documents for 63% of the population. 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` provides substantial coverage of the West and Northeast.[3]

The primary dataset for our analysis consists of the full-text of municipal ordinances. At the municipality-level, we also draw on demographic and housing information such as building permits data from the Census Building Permits Survey, rent and home price data from the American Community Survey (ACS), and many other data sources in Table A10.

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 represent a significant advancement in terms of both accuracy and runtime over previous models like Recurrent Neural Networks, which processed sequences linearly. LLMs are trained with semi-supervised learn-

construction needs to adhere to.

[3]When a municipality hosts its ordinance on multiple aggregators we prioritize Ordinance.com, and then Municode over American Legal Publishing. We show the percent of our sample that is from each source and so the share from Municode and American Legal Publishing understate the sources true coverage.

ing, 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 that the model does not require 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 in our our setting. The central advantage is scalability at low cost: we are able to load large quantities of municipal code data for classification and analysis, which exceeds the capacity of a typical human team to analyze at reasonable cost. Other advantages include the prospect for additional training, allowing for increased accuracy over time as LLMs improve in quality and additional training data is incorporated into the analysis. Another key benefit in creating a comprehensive and nationwide dataset resulting from the application of a uniform and standard set of criteria for analysis, rather than relying on a group of human analysts who may employ idiosyncratic legal interpretation.

Drawbacks in using LLMs for this purpose include potentially inaccurate measurement and the need for manual sourcing of relevant documents. Inaccurate measurement stems from a failure to locate or interpret the relevant sections of legal code. LLMs can only process a limited amount of text at once, and so they require a reliable process to locate the most relevant parts of the ordinance for a given question. Additionally, 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, 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 found in other legal documents. We provide the LLM with the municipal zoning ordinance for each municipality but not other potentially useful documents (i.e., state or county laws). We address these drawbacks by focusing on questions that can be answered with only a municipal zoning ordinance and measuring

performance against human-defined categorizations of regulation.

2.3 Processing Municipal Codes Using LLMs

In this subsection, we outline the general process we use to generate our housing regulatory dataset, a process we refer to as generative regulatory measurement. Figure 1 also illustrates the overall process.

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[4]. Any images of tables are transcribed using Amazon Textract.

The length of many zoning documents exceed the context windows currently usable by current LLMs (see Appendix Figure A13 for a histogram of token length across our entire sample compared with the feasibility of entry into different available LLMs). As a result, it is either impossible or cost-prohibitive to simply upload the entirety of municipal code documents into standard LLM services and query our questions directly. Moreover, LLM performance significantly declines with context legnth, rendering the effective length of LLMs well below the length of municipal codes (Modarressi et al., 2025). This is due to limitations in the attention mechanism, which struggles to retrieve relevant information as context length increases.

To address this limitation, our second step is to use a standard framework in computer science known as "retrieval-augmented generation" (RAG) (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 describe in more detail, in subsection 2.6, the accuracy and cost implications of this choice.

In order to implement RAG, we first partition each ordinance into small chunks of text[5]. Next,

[4]Ordinances sourced from Municode and American Legal Publishing are general municipal ordinances covering topics ranging from permits to purchase a pet to local election processes. Some general municipal ordinances omit a zoning section or only make references to a seperate zoning ordinance. We filter out such ordinances which do not contain zoning information by searching for key phrases in zoning documents, like common table headers (i.e. "Table of Uses") or zoning district names (i.e., R-1 for the first residential zoning district)

[5]We scrape each section within an ordinance separately, and then further split or combine sections so that their

we map each chunk of text into a vector representation called an embedding. Embeddings are vector representations of a text's semantic meaning, and are commonly used in AI research to efficiently measure the similarity of two sets of text (Reimers and Gurevych, 2019). We vectorize each subsection of the ordinance document using OpenAI embeddings models (link). Specifically, we use the "text-embedding-3-large" algorithm.

The third step in our process is to prepare a set of questions that we want to ask of our sample. We begin with the question base already used by the Pioneer Institute (i.e., "Is multifamily zoning allowed in this area as-of-right?"). In the initial step of our process, we simply use these questions verbatim. The text of the questions is also run through an embedding process to generate its own vector representation.

The fourth step of our process is to identify the most likely relevant information from the ordinance to show the LLM. We use cosine similarity, a standard measure of distance between two vectors, between the question embedding and the embedding of each chunk of text to rank chunks of text by relevance.

The fifth step is to refine the initial ranking of the most relevant text produced by cosine similarity and double-check for accuracy. We do so by using a cross-encoder reranking model[6] on the top 50 chunks of text, which processes the question and section text pairs simultaneously to determine the most semantically similar sections.[7] Reranking has been shown to increase retrieval performance (Anthropic, 2024). We then select text to show the LLM in order of highest relevance until a threshold. We choose 4,000 tokens (3,000 words) as the threshold, since beyond that length LLM performance is too low (Modarressi et al., 2025). The final output of this step is the set of 4,000 tokens representing the text in the initial document which is relevant to answer each question.

The sixth step in the process is the LLM call itself. We provide two key pieces of information to the LLM through an API call. First, we include 4,000 tokens of relevant text. Second, we provide the zoning question. Both pieces of information are provided in a single call to the LLM, in order to produce model output which is our answer. Each answer consists of an open-ended argument

length ranges between 50 and one thousand tokens of text. One token is roughly 3/4 of a word

[6]We specifically use the Cohere reranking model for this step.

[7]For some questions, we double-check the answer with different context creating by using keywords to rerank. Please see Appendix section A for more details on which questions we do this for and which keywords we check for.

followed by a parsable answer (i.e. "Yes" or "No"). The open ended answer allows for humans to audit the reasoning path of the LLM and has been shown to increase performance by providing space for the LLM to think out loud (Zhang et al., 2022).

These six initial steps are sufficient to produce an answer to each regulatory question sourced from municipal documents. However, they are not necessarily very accurate. Subsection 2.6 describes in more detail the accuracy in this, and at every stage, of the process. As a result, we followed a data-driven process to iterate and improve the accuracy of our approach. We engaged in substantial additional prompt engineering and other refinements to improve the accuracy of our analysis. To avoid overfitting, we conducted all such training on a distinct subsample of the Pioneer data, and then performed our final validity checks only once on a leave-out sample of the Pioneer data. We used three distinct strategies which we describe in more detail below: multistep processing, prompt engineering, and providing additional question background.

Prompt chaining is an additional technique which we used to produce additional information necessary to help answer each question. Prompt chaining breaks down LLM queries into multiple steps (prompts) where the output from one prompt is used as input in the next (Wu et al., 2022). 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 with distant embedding vectors. Additionally, LLM performance is enhanced when it is only required to answer a direct single step question in each call (Khot et al., 2023).

Additionally, we provide a "system prompt" in which 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 (Zhang et al., 2022).

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 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).

An additional step is to generate additional background information to provide to the LLM. The background information and model assumptions were initialized based on the Pioneer study (their "Issue Overview" and "Research Coding" sections for each question) when possible and were LLM generated otherwise. We further refined this background information to address areas of misinterpretation. Appendix A contains full information on the original Pioneer questions, our rephrased questions, as well as the additional background information and assumptions provided.

Additionally, we rephrase the 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. We did so through a data-driven approach in which we categorized model errors through more simple questions, and iterated on simplifications of questions to produce more accurate results. As mentioned above, our final validation estimates are free of any overfitting bias because they were estimated on a different set of municipalities than the ones which we used for training in this purpose. Appendix A contains a complete list of the original and rephrased questions used in our analysis.

2.4 Simple Example to Illustrate Approach

To help illustrate our process more concretely, we go through our general procedure for one question in our sample in Arlington, Massachusets on the presence of inclusionary zoning mandates or incentives.

First, we partition the ordinance into chunks and embed each chunk as well as the question. Panel A of Figure 2 visualizes a two-dimensional approximation of this embedding. Each dot represents one chunk of text and is colored by its broader article in the ordinance. Reassuringly, we find that chunks within the same article cluster together suggesting that embeddings are picking up on the underlying meaning of the text. The black star in the upper right quadrant represents embedding for the question itself. In short, RAG selects the chunks of text closest the question in embedding space to show the LLM.

Panel B of 2 zooms in on the embedding space around the question and shades its dot by its cosine similarity with the question vector. The closest subsections to the question at hand are indeed the relevant sections of the ordinance on affordable housing requirements:

"Section 11.08 - Affordable Housing Requirements Part 1: ...1. Fifteen percent (15%) of the Residential Units in new Projects shall be Affordable Units...

"Section 11.08 - Affordable Housing Requirements Part 2: where more than 50% of the units are affordable to households ... the number of parking spaces may be reduced to 50% of the requirements."

Our approach first isolates the nearest chunks of text, and then reranks these results. We then provide these texts, along with the question, and additional pieces of background information and assumptions:

Background Information: Inclusionary zoning requires or encourages developers to include affordable dwelling units within new developments of market rate homes. Some municipalities call it "incentive zoning" ... (omitted for brevity)

Assumptions: 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.

By providing relevant text and guidance to the LLM, we receive a final classification along with references for the source of this information:

GPT-4 Turbo Output: The zoning bylaw/ordinance provided does indeed include mandates and incentives for the development of affordable units. Here's a breakdown of the evidence supporting this conclusion:

1. **Mandates for Affordable Housing Development**: The ordinance mandates that 15% of the residential units in new projects be affordable units (Section 11.08, Part_1, d. REQUIREMENTS, 1.). This requirement applies to all new residential projects, including phased or segmented developments, with six or more units subject to Environmental Design Review (Section 11.08, Part_1, b. APPLICATION). The affordable units must conform to specific affordability criteria, ensuring they are accessible to households earning no more than 60% of the median income for rental units and no more than 70% of the median income for homeownership units (Section 11.08, Part_1, c. DEFINITIONS)...*(rest of answer omitted for brevity)*

2.5 Model Validation with Pioneer Data

A critical step in assessing the performance of LLM-based approaches lies in comparing model-generated classifications against a ground truth benchmark. To do so requires a high-quality annotated reference dataset. The Pioneer dataset serves as an excellent starting point 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 municipal ordinance text.

To address these issues, we construct a testing dataset based on 30 randomly chosen municipalities from the Pioneer Institute dataset. Importantly, as mentioned earlier, we made sure to leave out these municipalities from any prior training exercises to avoid overfitting. We also recode these munciipalities to 1) exclude question responses which relied on outside context, and 2) hand-correct inaccuracies in the original classification.[8]

Table 3 shows the performance results of our baseline GPT-4 Turbo model against the testing sample in Massachusetts. Among continuous questions (Panel A), our generated answers have an average correlation of 0.67 with 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 lowest 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). As shown in Panel B of Table 3, we observe a model accuracy of 96% across all binary questions. Because the raw accuracy measure may be biased depending on the base rate of answers, we also provide a Relative Squared Error (RSE) that compares each model result compared to a naive model which guesses the sample mode. We observe quite small RSEs as well.

2.6 Sources of Model Improvements

Although our benchmark results appear quite accurate, we provide in this section a detailed decomposition of where the accuracy gains come from, both overall as well as question-by-question. We do so to provide additional intuition for which aspects of our procedure are most important in improving final accuracy. Additionally, we also demonstrate

[8]Due 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.

robustness with respect to important features such as the choice of LLM model. We run each model five times to calculate standard deviation of accuracy for each specification.

The first basic choice we face is whether to use RAG models at all or not. Over our entire national sample, many municipal codes are simply too long to be fed into even the LLM models with the largest context windows. For these municipalities, RAG models are necessary. Because the municipalities that make up our core validation sample in Massachusets are shorter on average, we are able to compare the accuracy of specifications with and without RAG in Table 2. Here, we use Gemini Flash 1.5, a long context window model[9]. We test three specifications: No RAG, Basic RAG, and Enhanced RAG. No RAG provides the entire ordinance to the LLM as context. Basic RAG uses a smaller, less powerful, embedding model[10] and does not rerank text chunks, while Enhanced RAG follows our main specification of reranking and uses a more powerful embedding model. For each specification we only provide the question and ordinnce text to the LLM, omitting any prompt engineering or background information.

RAG approaches demonstrate improved performance over no-RAG models for both continuous questions (Panel A) and binary questions (Panel B). For continuous questions, we observe a dramatic reduction in "I don't know" responses when implementing RAG (dropping from 38.7% with no-RAG to 12.9% with enhanced RAG, with particularly strong improvements from basic to enhanced RAG). Continuous questions typically involve information that is localized to specific sections of the ordinance, such as in tables. The results confirm two findings from the literature for our setting: first, that refinements to RAG, such as re-ranking and the size of the embedding model, meaningfully improve retrieval (Anthropic, 2024); and second, that performance of long context models degrades as context length grows (Modarressi et al., 2025). These findings, combined with RAG's feasibility benefits and cost advantages, justify our choice to use RAG for the remainder of the paper.

Next, we examine the impact of iteratively layering the prompting strategies discussed in subsection refsubsec:processing. For this analysis, we switch to our primary model, GPT-4 Turbo, and use enhanced RAG throughout. We start with a "no prompting" specification that simply feeds the questions as originally phrased by the Pioneer Institute along with the context, achieving a correlation of 0.72 for continuous questions and an accuracy rate of 83.5% for binary questions. We then layer in additional prompting strategies: first adding question rephrasing through prompt engineering, then incorporating background information, and finally implementing subtask decomposition with multi-step processing and double checking. For binary questions, incorporating background information yields the largest improvement (from 86.0% to 93.9%), while for continuous questions, subtask decomposition provides the biggest boost (increasing correlation from 0.73 to 0.87). Our main specification, which uses subtask decomposition, reaches 0.87 correlation for continuous questions and 96.7% accuracy for binary questions, and we run it only once without standard deviations as we manually verify and adjust all errors.

In Table A5, we disaggregate these results to examine how different prompting strategies affect each question. The

[9]Gemini Flash 1.5 has a context winow of one million tokens, or roughly 750k words.

[10]Specifically we use the `multi-qa-mpnet-base-dot-v1` algorithm for basic RAG. See https://huggingface.co/sentence-transformers/multi-qa-mpnet-base-dot-v1.

impact varies substantially across questions. Some questions, like those about flexible zoning, achieve high accuracy even without prompting. Others show marked improvements from specific strategies - wetland restrictions in lot size calculations sees substantial gains from question rephrasing alone, while questions about affordable housing incentives benefit from both rephrasing and background information. For questions like longest frontage requirements, where relevant information often appears in semantically distant parts of the ordinance, subtask decomposition yields the largest improvements (correlation increasing from 0.46 to 0.70).

We also contrast the baseline approach with estimates drawn from other models in the model selection panel of Table 2. It appears that GPT 4-Turbo, using RAG and our full set of prompting, has the highest correlation for continuous questions (0.87) as well as the highest accuracy for binary questions (96.7%). GPT 3.5 Turbo is similar in accuracy for continuous questions, though considerably worse for binary questions, while Gemini Flash 1.5 is worse for both binary questions and continuous ones.

The key takeaway from our approach towards generative regulatory parsing is that, at least with models available at the 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.

2.7 Understanding Model Errors

To better diagnose reasons for model error in our baseline approach, 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 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 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. While ideally municipal regulations would identify a clear and unambiguous answer, we observe differences even among legal experts hired for the task of hand-classifying regulations. In principle, the ambiguous or unclear aspects of regulations can also be systematically classified through LLM-based approaches.

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 relevant text chunks 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. Appendix Table A2 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.

2.8 Additional Validation Checks

Additional Hand Validation: To assess accuracy outside of the Pioneer study Massachusetts sample, we conducted additional manual validation. First, we reviewed four process questions on a nationwide random sample of 30 municipalities (Panel A of Table A6). After removing ambiguous cases, the model achieved accuracy rates ranging from 89% to 100% across these questions. Second, we assessed accuracy on bulk zoning regulations by randomly selecting 30 municipalities in California (Panel B of Table A6). After dropping ambiguous cases—such as zoning rules that varied between interior and corner lots—the model achieved 89% accuracy for both measures. These results suggest that the model generalizes well beyond the Massachusetts-based training data, performing reliably across diverse regulatory contexts.

Comparison Against Wharton 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), which we describe in more detail in Appendix A. The key takeaways are that our approach improves substantially on Wharton in both coverage and accuracy.

A positive, though not perfect, correspondence between generated data and surveys is consistent with prior literature (Lewis and Marantz, 2019) investigating the reliability of survey-based responses on land use regulation. This paper finds planners have an incomplete understanding of their own municipality's regulation, and survey responses across years are inconsistent. Additionally, our data rely on the most recent zoning codes we were able to find whereas (Gyourko et al., 2021) uses survey questions from 2018. Consequently, some of the divergence likely also reflects changes in the underlying zoning codes over time.

2.9 Scalability and Replicability of LLM Regulatory Analysis

While LLMs hold promise for regulatory analysis based on accuracy, their key advantage lies in scalability and cost-effectiveness, which enables substantially larger systematic analysis of unstructured textual datasets.

Appendix Figure A1 illustrates the cost comparison between human analysis and two LLM-based approaches: one using retrieval-augmented generation (RAG) and another without RAG. Our analysis assumes several key parameters: lawyers require an average of five minutes per question-municipality pair at a rate of $50 per hour. For the RAG

approach, each LLM API call involves 4,000 tokens of text at $0.03 per query using GPT-4 Turbo pricing. For the No-RAG approach, each API call requires processing 330,000 tokens, reflecting the average length of ordinances in our sample. We analyze costs across 20 questions per municipality. Initial setup costs include $390 per question for preparation and $2.80 per municipality for data processing.

The key takeaway from the figure is that the RAG-based approach becomes substantially more cost-effective at larger scales, while the No-RAG approach exhibits a cost structure more similar to human analysis. Human-based analysis shows a linear cost increase as the number of municipalities grows. The RAG approach has a higher initial cost due to setup, but demonstrates significantly better scalability. The cost curves for RAG and human analysis intersect at approximately 300 municipalities, beyond which the RAG method becomes increasingly more cost-effective. While both LLM-based approaches require similar setup costs, the No-RAG method's heavy token usage results in steeper cost scaling, making it less advantageous for large-scale analysis compared to the RAG approach.

Moreover, the LLM approach offers additional benefits not captured in pure cost comparisons. These include faster processing times (the entire dataset for this study can be generated in under three days), consistent application of criteria across all documents (while human research assistants might vary in their interpretation of law), and the ability to easily update analyses as regulations change.

Another important consideration in adopting LLM-based approaches for regulatory analysis is the replicability of results. Here we summarize our key findings on LLM replicability, with complete details provided in Appendix Section A. LLMs' responses are not fully deterministic. Using a sample of 30 municipalities and two types of questions (a binary question about permit caps and a continuous question about the number of districts), we experiment with various approaches to mitigate this non-determinism, including ensemble methods (requesting multiple responses and aggregating by majority rule), temperature settings, and the use of random seeds.

We measure replicability using pairwise matching rates across multiple iterations of the same model specification, something we call a consistency score. There are two ways to request multiple iterations of output: requesting multiple completions in a single API call, or making separate API calls. This leads to two consistency metrics: internal consistency (agreement within a single API call requesting multiple completions) and external consistency (agreement between majority-rule answers across separate API calls). External consistency is our primary measure of interest as it captures replicability across independent runs of the model.

A key choice under the control of the researcher is the temperature, which is a hyperparameter controlling the randomness of the model's output. A lower temperature (closer to 0) makes the model's responses more deterministic by picking the most probable next token, while a higher temperature (closer to 1) increases randomness. We find that lower temperature settings generally lead to more consistent responses. With an ensemble size of 10, at temperature 0, we observe an internal consistency score of 0.9 for the continuous question about number of districts, compared to 0.61 at temperature 1, with similar patterns for external consistency. LLMs also allow users to set random seeds

for reproducibility, though we find these do not significantly affect consistency scores.

External consistency grows with ensemble size, especially for the continuous question, though at least 5% of pairwise comparisons do not match even with temperature 0 and ensemble size of 10. Interestingly, internal consistency scores are more predictive of external consistency for higher temperature models, suggesting that within-API call variation at low temperatures may not reflect the same sources of randomness as across-API call runs.

To mitigate non-determinism, we find a tradeoff between a low-temperature single-shot approach with its high external consistency and cost-effectiveness, or a high-temperature ensemble approach with at least five iterations for improved reliability. These strategies balance the trade-offs between consistency, accuracy, and computational cost. Our main specification implements a low temperature single-shot approach, prioritizing external consistency and cost-effectiveness. This approach also reflects our interest in creating a single national dataset, for which consistent interpretation of zoning regulations across municipalities is important. However, for other applications requiring more creative or diverse output, the high-temperature ensemble approach may be more appropriate.

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 corporation or a township, 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 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: Atlanta, Chicago, Philadelphia, and the San Francisco Bay Area. We pick these regions to illustrate patterns across a range of urban forms and geographical areas. Philadelphia represents an older, densely populated Eastern seaboard city with a long history of municipal incorporations. Chicago is an example of a Midwestern metropolis with a mix of urban density and suburban sprawl. Atlanta highlights a Southern city with extensive suburban development and growth in unincorporated areas. The San Francisco Bay Area highlights the West Coast region with complex topography and a patchwork of dense urban centers, suburbs, and unincorporated zones.

In these maps, we highlight incorporated areas in green and unincorporated areas in red, shading the areas to reflect housing unit density. Philadelphia and Chicago show a more continuous pattern of incorporation, with fewer and smaller unincorporated areas, reflecting their longer urban histories and denser development patterns. In contrast, Atlanta exhibits a more fragmented landscape with larger swaths of unincorporated land, particularly

in its outer suburbs, indicative of its more recent, automobile-centric development. Perhaps surprisingly, some of this unincorporated land is relatively dense, reflecting built up activity, which is concentrated in the fringes of this unincorporated zone closest to the city itself. The San Francisco Bay Area is a more complex mosaic of incorporated cities interspersed with significant unincorporated areas, influenced by its varied topography, state parkland, and growth patterns.

We plot in Figure 5 the concentration of housing growth in unincorporated areas. Appendix Section A describes our construction of incorporated and unincorporated areas in more detail, which derive from shapefiles matched to the Government Units Survey of the Census of Governments. 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 MSA level. Fully incorporated metros had substantially less growth in housing units over this period than partially incorporated metros. We show in Panel B the amount of new housing units added specifically in incorporated and unincorporated areas, generating two categories based on census block level housing data and the percent area of each block that is incorporated. We find 42% of new housing production takes 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, local ameneties, and public goods provisioning (Schoenholzer, 2018). We show in Table 4 associations between incorporated and unincorporated areas in the U.S. Unincorporated areas have substantially cheaper homes as measured 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. Unincorporated areas may also have a larger number of large plots of land for sale, reducing challenges associated with land assembly (Brooks and Lutz, 2016). Additionally, many incorporated municipalities have adopted growth control measures that limit new construction, pushing development to the fringes.

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. Higher commute lengths contribute to limited employment opportunities and more environmental pollutants and congestion. 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 Monocentric City Model and Zoning Gradients

We next interpret municipal regulations in the context of the monocentric city model (Alonso, 1964; Mills, 1967; Muth, 1971). In these models, there is a central location in each city where production is concentrated and rents decay as one moves away from this productive center, with the rate of decay governed by transportation costs. These dynamics may also affect the benefits and costs of zoning regulations at different distances from the city center. Housing regulations, in turn, may then affect the rent gradient as one moves away from the city center.

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. We also consider on minimum lot size requirements, which dictate the smallest allowable area for a residential plot. This regulation sets a lower bound on how small a piece of land can be for a single housing unit, impacting the potential density of an area. Larger minimum lot sizes result in fewer, more spread-out homes, while smaller minimums allow for denser development. They have been frequently estimated in prior research through bunching methods (Cui, 2024; Song, 2021) as important drivers of housing regulations; the contribution of our approach is to measure these regulations directly from municipal documents, rather than indirectly through their effects on observed housing development.

Figure 11 shows that these bulk regulations vary substantially across the U.S. In Appendix Figure A3, Panel A we show that among single-family zoned areas in municipalities, 66% have lowest 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. Panel B of this figure highlights the *highest* residential minimum lot size requirement in the municipality: prevailing requirements for individual properties will therefore be between these two values. 89% of municipalities have a maximum requirement above 5,000 square feet; 64% have a requirement above 10,000 square feet; and 42% of municipalities have a maximum requirement above half an acre. 12% of municipalities have a maximum requirement above 100,000 square feet. These findings underscore the prevalence and magnitude of minimum lot size requirements across the United States.

Minimum lot size requirements show a different pattern, and vary markedly across regions. In much of the country outside the Northeast and Midwest, 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 12 and 13 show maps of minimum lot sizes and affordable housing incentives, respectively, for jurisdictions within the metropolitan areas surrounding four select cities in the U.S., Atlanta, Chicago, Philadelphia, and San Francisco.

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 policies 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.

We also show these associations in Table A9. Across all regions, we observe that the number of zoning districts decreases robustly with distance from the city center, suggesting simpler zoning structures in more peripheral areas. This pattern is consistent across all regions but is particularly pronounced in the Midwest and South. Some components of allowable density decrease with distance from the center, especially the permission of townhouses and mixed-use development. Interestingly, the allowance of multifamily housing shows a positive correlation with distance in the West and South, contrary to the general expectation of decreasing density with distance. This might reflect the presence of suburban multifamily developments in these regions.

The Northeast stands out with several distinct patterns. Unlike other regions, it shows increasing restrictiveness with distance for several measures. For instance, the longest frontage requirement, highest residential minimum lot size, and mean residential minimum lot size all increase with distance from the city center in the Northeast, while these measures show no significant relationship in other regions. These results highlight the unique regulatory landscape of the Northeast, where bulk regulations and exclusionary zoning practices appear to intensify in suburban and exurban areas, contrary to patterns observed in other regions. These results are also especially surprising in the context of well-developed public transit and highway links in this area, which should, all else equal, facilitate greater development and density even outside of city centers in this region. The percentage of area unincorporated also increases robustly with distance from city center, as shown in Appendix Figure A2.

Several factors may help to explain this regional variation. The Northeast was the first region of the U.S. to urbanize and industrialize, leading to early adoption of zoning laws. Many of its suburbs were established earlier than in other regions, often as affluent enclaves seeking to preserve their character against urban expansion (Fischel, 2015). In addition, the region is characterized by a highly fragmented system of local governments, with many small, independent municipalities. This structure facilitates more localized and potentially more restrictive zoning policies. These suburbs also have a strong home role tradition of local control over land use decisions. The Northeast's early experience with industrial pollution and urban congestion may have fostered a culture of environmental protection that manifests itself in stricter land use controls, particularly for suburban lot size requirements, which were commonly justified on the basis of preserving natural land. Finally, the region has had particularly strong fights over access to local schools (i.e., school busing (Angrist et al., 2022)) that can increase the use of exclusionary zoning practices to maintain local school districts by limiting access to lower-income and minority households.

In contrast, California, which also has high house prices and where housing regulation is commonly thought to be tight, appears surprisingly to have more nationally typical bulk regulations. We show a map representation

of these associations in Figure 8, which highlights the nature of high minimum residential lot size requirements in the northeast relative to California. In contrast, California's housing market is heavily influenced by state-level regulations such as environmental reviews and an onerous permitting process. We find some evidence of this in Appendix Figure A4, which shows that the West has the highest potential waiting time for review of a typical new multi-family building (see Mayer and Somerville (2000) for a link between regulatory uncertainty and reduced construction), and also ranks highly on public hearing requirements for multi-family building. These results therefore suggest that while the Northeast and California have high housing costs, they are the result of very different factors that may call for different policy responses. Furthermore, California's housing market faces significant constraints due to its challenging topography (Saiz, 2010), which limits developable land, and the high demand for natural amenities and local jobs. These factors can create a situation in which even relatively less restrictive local bulk regulations can result in binding constraints on housing supply, especially in combination with a challenging permitting process.

Including this spatial variation in regulation has several implications for urban economic models. The broad patterns of zoning intensity by distance to city center is consistent with the predictions of standard monocentric city models. However, the sorting of high-income households into expensive, strictly zoned suburbs is inconsistent with the simplest models in which all households prefer central locations to minimize commuting costs, and the richest households locate centrally as they are able to pay the most for scarce urban land. The high market values of housing far from city centers and the choice of richer residents to locate there and commute long distances to urban jobs and amenities are a known challenge for this framework (Glaeser et al., 2008). This spatial pattern is more pronounced in the United States compared to other countries, where it is more typical for wealthy residents to reside in the city center. Our results suggest that this sorting is accompanied by regulatory restrictions on the minimum allowable housing size, which truncates the housing size distribution to the left, and results in sorting of higher-income residents to distant suburbs against the typical pattern expected in the monocentric city model.

The persistence of stringent zoning in affluent suburbs, often more restrictive than in central areas, suggests that regulatory requirements may lock in certain urban forms, contributing to the path dependency of metropolitan development patterns. The gradient in zoning complexity and intensity, particularly pronounced in regions like the Northeast, reflects differences in political processes, historical development, and local preferences between central cities and suburbs.

Fact 2. *Zoning regulations generally become less complex but more restrictive of density with increasing distance from city centers, with this pattern most pronounced in the Northeast where exclusionary measures intensify in suburban areas.*

4 Characterizing Housing Regulations

4.1 Principal Component Analysis of Housing Regulatory Dataset

To summarize our nationwide measure of housing regulations, we perform a PCA analysis. This technique reduces the dimensionality of our dataset by identifying key components that capture a large fraction of variation across our regulatory questions. Our analysis reveals two key principal components of housing regulation. Appendix Table A4 provides the loadings of each question on the first five principal components, and Figure 8 maps the first two principal components across the nation.[11]

In Panel B of Table A14, we find positive correlations of both PCs with the composite Wharton Index at the CBSA level. The first PC correlates at 0.33 against the Wharton Index, while the second PC correlates at 0.11. An index that sums the normalized values of each zoning regulation correlates at 0.22 with the Wharton index. 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.

To better understand the economic interpretations behind these two principal components, and 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 our two key principal components. We interpret the association of these two principal components with building permits and house prices in light of a simple model of supply and demand for housing. In places with rising demand for housing and inelastic supply, home prices will be high and building permits low (the lower, right quadrants of Panels A and B of Figure 9. In areas with rising demand for housing and elastic supply, there will be high house prices and high building permits (i.e. the upper right quadrant). In places with falling housing demand, there will be few building permits and low housing prices. In places with elastic housing supplies and constant or moderately rising demand, there will be low home prices and high permits per capita.

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 that the first principal component generally coincides with high housing demand environments. This interpretation is supported by the main regulatory loadings on this principal component (Appendix Table A4), which loads heavily on measures that are typically associated with more developed, high-demand housing markets. For instance, affordable housing and age restricted provisions are policy tools that are more likely to be implemented in areas with significant housing pressure and the administrative capacity to manage complex policies. Other associates of the first principal component relate to additional layers of local government: maximum review wait times and public hearing requirements. However, areas heavy in this component are much less likely to have bans or limits on multifamily and other dense housing typologies, though more likely to restrict permitting, require phasing, or restrict

[11] Appendix Figure A5 also shows a heatmap of correlations between regulations at the national level.

the inclusion of wetlands in lot size calculations.

We show associations of the second principal component in Table **??**. We focus on column 1, which controls for MSA fixed effects. The higher housing values and rents associated with areas scoring high on this component suggest that exclusionary zoning effectively creates pockets of affluence within metropolitan areas. This is further reinforced by the higher income levels and lower poverty rates in these localities. The lower share of affordable units in these areas is a direct consequence of these zoning practices, effectively pricing out lower-income households.

We summarize these associations as our third fact:

Fact 3. *The first two principal components of housing regulation capture distinct patterns in housing markets - the first component associates with high prices and high construction, while the second component associates with high prices but low construction.*

4.2 Value Capture and Housing Markets

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

The first Principal Component (PC) strongly correlates with various demand-side factors, as demonstrated by our analysis of county-level amenities. Using population-weighted county aggregates, Figure **??** displays normalized coefficients with 95% confidence intervals for the relationships between our PCs and three key measures. These measures include: a natural amenity index (capturing features like temperature and sunlight hours), patents per capita (measuring productivity), and an index of retail establishments per capita across industries (constructed by taking the first PC of industry-level establishment counts). We find that the first PC has statistically significant positive relationships with all three measures, while the second PC shows no statistically significant relationships with any of them. The relationship becomes even stronger when we construct a combined amenity index by taking the first principal component across these three measures.

The retail establishment patterns are particularly revealing: counties scoring high on the first PC have higher concentrations of retail outlets (like apparel stores and restaurants) and professional services (including educational institutions, healthcare facilities, and cleaning services) per capita, while having fewer establishments typically associated with negative externalities, such as gas stations, utility services, and truck transportation businesses per capita. This pattern suggests these areas are high-demand locations where businesses can command premium rents. For detailed correlations between the first PC and individual components of both the natural amenities index and retail establishments index, see Appendix Figure **??**. Further reinforcing the interpretation of the first PC as reflecting

demand-driven amenities, Appendix Figure A14 shows that the first PC exhibits a high degree of spatial autocorrelation, as captured by Moran's I. In contrast, the second PC has very low spatial autocorrelation, suggesting that it is less influenced by spatially clustered amenities and more likely driven by other economic forces. This supports the notion that the first PC captures factors that vary smoothly across space, such as quality-of-life amenities, which naturally exhibit spatial dependence.

To be sure, such regulations may also affect housing supply. We explore this variable in more detail in columns 1–2 of Appendix Table A8, 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. Perhaps surprisingly, the first principal component only moderately associates with lower housing elasticities once MSA fixed effects are included. This suggests that the bundle of regulations represented by this set of housing regulations—lower bulk regulations and limits on density, and positive association with more process, complexity, and good provision requirements—nearly balance each other out in housing supply.

Fact 4. *The first principal component corresponds to value capture in high-demand locations. This component loads heavily on regulations that allow municipalities to extract and redistribute housing value, such as inclusionary zoning mandates that require developers to include affordable units. These regulations enable local governments to convert market-rate housing demand into public benefits, and are most prevalent in densely populated, centrally located cities.*

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. The role of exclusionary zoning in shaping socioeconomic patterns has been a subject of significant research and debate in urban economics and policy circles. Our analysis provides new insights into this phenomenon, leveraging our comprehensive dataset to examine its prevalence, distribution, and impacts across the United States. Regulations commonly linked to exclusionary zoning practices 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 A4.

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 next show that municipalities with high second PCs have substantially fewer affordable housing units to the state median income household, as shown in Panel A of figure 10. We identify three key mechanisms driving this relationship. First, exclusionary zoning shifts the housing price distribution rightward, with significantly higher me-

dian home values in high second PC municipalities compared to low second PC areas (Panel B). Second, exclusionary zoning truncates the left tail of the housing price distribution. In Panel C, we compare the distribution of home values across granular home value bins for low and high second PC areas. High second PC areas have relatively more housing units at or right below the median home price and less well below the median, suggesting a leaner left tail. We further show in appendix figure A15 that this left truncation is more pronounced in higher median home values areas. Third, Panel D reveals that high second PC areas limit the degree of rental housing with a strong correlation of 0.40 between the second PC and the share of owner-occupied housing. The resulting lack of affordable housing options results in a strong negative correlation between the second PC and the share of low-income families (-0.26) and young households under 35 (-0.39).

We then validate that these density restrictions do have implications for observed density in Table A3. In this Table, we compute the association between our measured regulations and housing values, building permits, rents, observed density, and the affordability of local housing units (rental units cost $\leq$ 30% of state monthly median income, and owner-occupied units priced $\leq 3\times$ annual median household income). We present three sets of results on the relationship between regulations and housing outcomes. First, we use bivariate regressions to understand the unconditional associations between a given housing regulation and a particular housing outcome. Second, we estimate multivariate regression models using the full set of regulations using the Least Absolute Shrinkage and Selection Operator (LASSO) to investigate which regulations are most important and report conditional associations between these selected regulations and housing outcomes. Third, we use a random forest (XGBoost) specifications that allows for more complex interactions and potentially non-linear relationships between regulations and housing outcomes. For XGBoost there are multiple ways to measure feature importance but we use “gain" importance which is well suited to an environment that includes both binary and continuous variables. We then validate that these density restrictions do have implications for observed density in Table A3. In this Table, we compute the association between our measured regulations and housing values, building permits, rents, observed density, and the affordability of local housing units (rental units cost $\leq$ 30% of state monthly median income, and owner-occupied units priced $\leq 3\times$ annual median household income). We present three sets of results on the relationship between regulations and housing outcomes. First, we use bivariate regressions to understand the unconditional associations between a given housing regulation and a particular housing outcome. Second, we estimate multivariate regression models using the full set of regulations using the Least Absolute Shrinkage and Selection Operator (LASSO) to investigate which regulations are most important and report conditional associations between these selected regulations and housing outcomes. Third, we use a random forest (XGBoost) specifications that allows for more complex interactions and potentially non-linear relationships between regulations and housing outcomes. For XGBoost there are multiple ways to measure feature importance but we use “gain" importance which is well suited to an environment that includes both binary and continuous variables. We then validate that these density restrictions do have implications for observed density in Table A3. In this Table, we compute the association between our measured regulations

and housing values, building permits, rents, observed density, and the affordability of local housing units (rental units cost $\leq 30\%$ of state monthly median income, and owner-occupied units priced $\leq 3\times$ annual median household income). We present three sets of results on the relationship between regulations and housing outcomes. First, we use bivariate regressions to understand the unconditional associations between a given housing regulation and a particular housing outcome. Second, we estimate multivariate regression models using the full set of regulations using the Least Absolute Shrinkage and Selection Operator (LASSO) to investigate which regulations are most important and report conditional associations between these selected regulations and housing outcomes. Third, we use a random forest (XGBoost) specifications that allows for more complex interactions and potentially non-linear relationships between regulations and housing outcomes. For XGBoost there are multiple ways to measure feature importance but we use "gain" importance which is well suited to an environment that includes both binary and continuous variables.

We find that density restrictions (such as multifamily units not being allowed, minimum residential lot sizes, and frontage requirements) are strongly and robustly associated with lower observed density levels. The XGBoost model provides feature importance scores, normalized from 0 to 100, indicating the relative predictive power of each zoning regulation for various housing market outcomes, with higher scores suggesting greater influence on the predicted variable. The longest frontage requirement is identified as the most crucial factor in the random forest specification, with a maximum score of 100, closely followed by maximum (86) and minimum (84) residential lot size requirements.

While such a relationship is in principle consistent with lower demand in these regions, we also find that the same variables are strongly associated with higher house prices and rents, and a lower share of affordable housing units. In predicting median home values, the prohibition of multifamily housing is the most important predictor in the random forest specification, with a maximum score of 100. This is followed by affordable housing mandates (75) and prohibition of townhouses (56). We observe similar associations in these density regulations associating with higher median gross rents and a lower share of available affordable housing. Looking at affordability requirements has the added benefit of effectively normalizing for local income effects. These bulk regulations are also typically selected in the LASSO specifications on these outcomes, and have significant coefficients in bivariate regressions.

The consistent association between stringent zoning regulations, higher housing prices, and lower density points to an impact of these regulations on urban development patterns and housing markets. These regulations actively constrain housing supply in desirable locations. By restricting housing types, particularly in high-demand areas, these zoning practices appear to create artificial scarcity in the housing market

To further assess the predictive power of zoning regulations for housing market outcomes, we compare the performance of different modeling approaches in Table A7. This table reports the Root Mean Square Error (RMSE) for Lasso, OLS, and XGBoost models across various specifications predicting these same variables, contrasting the relative role of regulations and controls for land availability and income.

Zoning regulations demonstrate significant predictive power across all housing market outcomes. Including housing regulations improves model fit, even when other controls are already included. The XGBoost model consistently

outperforms Lasso and OLS, particularly when all variables are included (column 7), achieving the lowest RMSE of 0.66 for median home value, suggesting the importance of complex, non-linear relationships between zoning regulations and housing market characteristics.

A plausible rationale for such practices are the spillovers of lower-income residents, including perceived impacts on crime rates, traffic congestion, educational outcomes, and house prices more broadly. Consistent with this rationale, the demographic patterns associated with exclusionary zoning correspond to a higher proportion of white residents and lower foreign-born share. This suggests racial and ethnic segregation effects of these policies. This aligns with analyses of zoning as a tool for maintaining racial homogeneity in the absence of explicit racial covenants (Rothwell, 2011; Cui, 2024). Density patterns favor single-family residents, with owner-occupied fractions substantially higher and housing unit density substantially lower.

Areas characterized by more exclusionary zoning practices also show higher average math test scores as well as rates of economic opportunity (Chetty et al., 2014). By contrast, there are lower rates of free lunch eligibility among students. This pattern suggests that exclusionary zoning effectively creates enclaves of educational privilege, where resources and positive peer effects are concentrated. Consistent with a public goods motive, we also observe higher local revenue per student, property tax rates, and total expenditures per capita. These results are consistent with classic theories of fiscal zoning which emphasize the role of zoning and property taxes in conjunction to support local public goods (Fischel, 1987; Hamilton, 1975, 1976). The presence of such public goods helps to explain the sorting of higher-income residents in these areas, which are commonly not in the city center, helping to explain the spatial patterns discussed in the previous section. Ultimately, the key factor driving these associations is the lack of school finance equalization in the United States, unlike in other countries, which allows for more local differentiation in school quality reinforced through land use regulations which impact the composition and tax base of the local population.

Fact 5. *The second principal component captures exclusionary zoning practices aimed at restricting housing density and affordability through bulk regulations. This dimension relies heavily on minimum lot size requirements and procedural barriers to development, which are concentrated in affluent suburban areas. These regulations associate with higher housing prices, lower construction, and strong sorting on education and income, suggesting a role for public goods provision in motivating these restrictions.*

5 Discussion and Framework

We interpret our empirical findings through a model of inter-municipal competition in which local governments strategically select housing regulations. The model features two types of locations (city centers and suburbs) and two types of households (high and low wage). A more complete version of the model is discussed in Appendix E.

Suburbs adopt exclusionary policies and attract primarily high-income residents This pattern is reinforced by

complementarities: exclusion raises migration elasticities and makes value capture less attractive

The first principal component of our housing regulation dataset, corresponding to value capture, 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 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 children who pay particular attention to the level of school spending and composition of the local school district.

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.

Figure A7 provides a visual scatterplot representation of the relationship between these two principal components across U.S. municipalities. The plot also highlights the heterogeneity in zoning practices across municipalities. Suburban areas around high-demand metropolitan areas (such as Darien, CT or Billerica, MA) rank highly on both PCs. High-demand urban areas (like Santa Ana, CA) score low in PC2, but are relatively high in PC1. Low-demand

urban areas (like Cleveland, OH) rank low on both PCs. While exclusionary suburbs of low-demand areas (like Novi, MI) rank high in PC2 but low in PC1.

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 relative to existing hand-classification from the Pioneer Institute. Our approach also correlates with existing measures of regulation from the Wharton Index. This generative regulatory measurement approach therefore enables the creation of a comprehensive, nationwide dataset of municipal zoning regulations. By leveraging LLMs to extract structured information from unstructured textual data, our methodology opens up new avenues for analyzing vast amounts of previously untapped regulatory documents across multiple domains.

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 the importance of distinguishing 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.

Angrist, Joshua, Guthrie Gray-Lobe, Clemence M Idoux, and Parag A Pathak, "Still worth the trip? School busing effects in Boston and New York," Technical Report, National Bureau of Economic Research 2022.

Anthropic, "Introducing Contextual Retrieval," 2024.

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.

Brooks, Leah and Byron Lutz, "From Today's City to Tomorrow's City: An Empirical Investigation of Urban Land Assembly," *American Economic Journal: Economic Policy*, August 2016, *8* (3), 69–105.

Brueckner, Jan K, "Strategic control of growth in a system of cities," *Journal of Public Economics*, 1995, *57* (3), 393–416.

_ , *Lectures on Urban Economics*, MIT press, 2011.

Brynjolfsson, Erik, Danielle Li, and Lindsey R Raymond, "Generative AI at work," Technical Report, National Bureau of Economic Research 2023.

Bybee, J Leland, "The ghost in the machine: Generating beliefs with large language models," Technical Report, Technical report, Working Paper 2023.

Calder-Wang, Sophie and Gi Heung Kim, "Coordinated vs efficient prices: The impact of algorithmic pricing on multifamily rental markets," Technical Report, Working Paper 2023.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez, "Where is the land of opportunity? The geography of intergenerational mobility in the United States," *The Quarterly Journal of Economics*, 2014, *129* (4), 1553–1623.

Cui, Tianfang, "Did Race Fence Off The American City? The Great Migration and the Evolution of Exclusionary Zoning," Technical Report, Working Paper 2024.

D'Amico, Leonardo, Edward L Glaeser, Joseph Gyourko, William Kerr, and Giacomo AM Ponzetto, "Why Has Construction Productivity Stagnated? The Role of Land-Use Regulation," Technical Report, Working Paper 2023.

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.

Eisfeldt, Andrea L, Gregor Schubert, and Miao Ben Zhang, "Generative AI and firm values," Technical Report, National Bureau of Economic Research 2023.

Epple, Dennis and Allan Zelenitz, "The implications of competition among jurisdictions: Does Tiebout need politics?," *Journal of Political Economy*, 1981, *89* (6), 1197–1217.

Fischel, William, *Zoning Rules!: The Economics of Land Use Regulation*, Lincoln Institute of Land

Policy, 2015.

Fischel, William A, *The economics of zoning laws: A property rights approach to American land use controls*, JHU Press, 1987.

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.

_ , Matthew E Kahn, and Jordan Rappaport, "Why do the poor live in cities? The role of public transportation," *Journal of urban Economics*, 2008, *63* (1), 1–24.

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.

_ and Raven Molloy, "Regulation and housing supply," in "Handbook of regional and urban economics," Vol. 5, Elsevier, 2015, pp. 1289–1337.

_ , 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.

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," Technical Report, Working Paper 2023.

Horton, John J, "Large language models as simulated economic agents: What can we learn from homo silicus?," Technical Report, National Bureau of Economic Research 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," Technical Report, Chicago Booth Research Paper 2023.

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

Khot, Tushar, Harsh Trivedi, Matthew Finlayson, Yao Fu, Kyle Richardson, Peter Clark, and Ashish Sabharwal, "Decomposed Prompting: A Modular Approach for Solving Complex Tasks," 2023.

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.

Lewis, Paul G and Nicholas J Marantz, "What planners know: Using surveys about local land use regulation to understand housing development," *Journal of the American Planning Association*, 2019, *85* (4), 445–462.

Li, Junyou, Qin Zhang, Yangbin Yu, Qiang Fu, and Deheng Ye, "More agents is all you need," Technical Report, Working Paper 2024.

Lutz, Byron, "Quasi-Experimental Evidence on the Connection between Property Taxes and Residential Capital Investment," *American Economic Journal: Economic Policy*, February 2015, *7* (1), 300–330.

Mayer, Christopher J and C Tsuriel 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.

Modarressi, Ali, Hanieh Deilamsalehy, Franck Dernoncourt, Trung Bui, Ryan A. Rossi, Seunghyun Yoon, and Hinrich Schütze, "NoLiMa: Long-Context Evaluation Beyond Literal Matching," 2025.

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.

Raymond, Lindsey, "The Market Effects of Algorithms," Technical Report, Working Paper 2023.

Reimers, Nils and Iryna Gurevych, "Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks," 2019.

Rothstein, Richard, *The Color of Law: A Forgotten History of how our Government Segregated America*, Liveright Publishing, 2017.

Rothwell, Jonathan T., "Racial Enclaves and Density Zoning: The Institutionalized Segregation of Racial Minorities in the United States," *American Law and Economics Review*, 2011, *13* (1), 290–358.

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

Schoenholzer, David S, *Essays on State Capacity and Local Public Goods*, University of California, Berkeley, 2018.

Shanks, Brendan, "Land Use Regulations and Housing Development," 2021. manuscript.

Singla, Shikhar, "Regulatory costs and market power," Technical Report 2023.

Song, Jaehee, "The effects of residential zoning in US housing markets," Technical Report, Working Paper 2021.

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.

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*.

Wei, Jason, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus, "Measuring short-form factuality in large language models," 2024.

Wu, Tongshuang, Michael Terry, and Carrie Jun Cai, "AI Chains: Transparent and Controllable Human-AI Interaction by Chaining Large Language Model Prompts," in "Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems" CHI '22 Association for Computing Machinery New York, NY, USA 2022.

Yang, Stephen, "Predictive Patentomics: Forecasting Innovation Success and Valuation with ChatGPT," Technical Report, Working Paper 2023.

Zhang, Zhuosheng, Aston Zhang, Mu Li, and Alex Smola, "Automatic Chain of Thought Prompting in Large Language Models," 2022.

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

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

Notes: 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, and 2022 State-Level Census Population Data for census region and national population.
Links to data sources are American Legal Publishing, Municode, and Ordinance.com.

Table 2: Model Accuracy Contributions

Panel A: Continuous Questions

		Correlation		I Don't Know Rate	
Model	Specification	Mean	Std. Dev.	Mean	Std. Dev.
RAG Strategies					
Gemini Flash 1.5	No RAG	0.74	0.01	38.7%	2.3
Gemini Flash 1.5	+ Basic RAG	0.66	0.03	32.7%	1.7
Gemini Flash 1.5	+ Enhanced RAG	0.84	0.02	12.9%	0.6
Prompting Strategies					
GPT-4 Turbo	No Prompting	0.72	0.05	2.7%	0.6
GPT-4 Turbo	+ Question Rephrasing	0.78	0.05	2.9%	1.3
GPT-4 Turbo	+ Background Info	0.73	0.11	2.2%	1.4
GPT-4 Turbo	+ Subtask Decomposition	0.87	-	4.4%	-
Model Selection					
GPT-3.5 Turbo	Full RAG + Full Prompting	0.86	0.03	3.8%	1.0
Gemini Flash 1.5	Full RAG + Full Prompting	0.81	0.01	9.6%	1.7
GPT-4 Turbo	Full RAG + Full Prompting	0.87	-	4.4%	-

Panel B: Binary Questions

		Accuracy		I Don't Know Rate	
Model	Specification	Mean	Std. Dev.	Mean	Std. Dev.
RAG Strategies					
Gemini Flash 1.5	No RAG	81.0%	1.24	21.1%	1.0
Gemini Flash 1.5	+ Basic RAG	84.5%	1.18	22.0%	1.7
Gemini Flash 1.5	+ Enhanced RAG	86.2%	0.87	18.6%	0.8
Prompting Strategies					
GPT-4 Turbo	No Prompting	83.5%	1.4	14.2%	5.3
GPT-4 Turbo	+ Question Rephrasing	86.0%	0.8	10.7%	1.4
GPT-4 Turbo	+ Background Info	93.9%	1.0	1.1%	0.2
GPT-4 Turbo	+ Subtask Decomposition	96.7%	-	0.0%	-
Model Selection					
GPT-3.5 Turbo	Full RAG + Full Prompting	83.9%	0.6	3.1%	0.7
Gemini Flash 1.5	Full RAG + Full Prompting	89.1%	1.2	4.1%	0.4
GPT-4 Turbo	Full RAG + Full Prompting	96.7%	-	0.0%	-

We calculate the mean and standard deviation across five runs of each specification. See appendix section A for more details on sources of variation across runs. The GPT-4 Turbo Full RAG + Full Prompting, which includes the '+ Subtask Decomposition' prompting strategy, uses the main validation sample. We manually verified errors and adjusted the Pioneer dataset accordingly, reporting accuracy from a single run without standard deviation. See Figure 3 for more details on error adjustments.

Table 3: Accuracy Metrics of GPT-4 Turbo on Validation Sample

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
What is the lowest residential minimum lot size?	0.16	0.92
Cumulative Average	0.46	0.87
Cumulative Median	0.31	0.90

Notes: For Relative Squared Error (RSE) 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 calculate performance metrics and sample means (for RSE) only on the set of question municipality pairs that GPT-4 Turbo does not say "I don't know." 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 homes 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%

Notes: For Relative Squared Error (RSE) 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 where the model matches the adjusted Pioneer Institute answer for each question. We drop any errors where the answer is considered ambiguous. For details on adjustments to the Pioneer data, see Figure 3.

Table 4: Characteristics of Unincorporated Regions

Panel A: Housing					
Independent Variable:	Block Group is Unincorporated				
Specification	Average	Bivariate	Metro FE	Distance FE	Metro and Distance FE
Median Home Value		-6.0***	-2.6***	-1.9***	-0.3
		(0.1)	(0.8)	(0.1)	(0.7)
Median Year Built	1974	13.4***	7.2***	11.9***	5.6***
		(0.1)	(0.7)	(0.1)	(0.6)
Median Gross Rent		-98.2***	-81.1***	-8.8**	-37.4**
		(4.2)	(26.5)	(3.8)	(19.0)
Vacancy Rate		2.7***	2.2***	1.9***	1.6***
		(0.0)	(0.4)	(0.0)	(0.3)
Rental Rate		-12.6***	-15.1***	-9.9***	-12.5***
		(0.1)	(0.8)	(0.1)	(0.8)
Percent Commute Over 60		1.2***	2.9***	1.6***	2.8***
		(0.0)	(0.2)	(0.0)	(0.2)
Percent Over 65		2.0***	2.5***	1.3***	1.7***
		(0.1)	(0.3)	(0.1)	(0.3)
Panel B: Climate					
Independent Variable:	Tract is Unincorporated				
Expected Annual Loss	49.7	12.8***	4.0***	11.1***	4.9***
		(0.2)	(0.2)	(0.2)	(0.8)
Wildfire	41.7	25.5***	18.8***	18.7***	17.0***
		(0.2)	(0.3)	(0.2)	(1.6)
Hurricane	59.9	15.4***	0.1	15.9***	2.4***
		(0.3)	(0.1)	(0.3)	(0.6)
Coastal Flooding	21.1	13.0***	0.0	9.8***	1.8
		(0.5)	(0.5)	(0.5)	(2.4)
Riverine Flooding	46.9	15.2***	6.6***	10.0***	7.7***
		(0.2)	(0.3)	(0.2)	(1.2)

Notes: This specification regresses an indicator on whether a block group is incorporated against a variety of block group characteristics with different levels of fixed effects. 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. We use regression weights that are the product of the block group population and the percent unincorporated for the unincorporated observations and the percent incorporated for incorporated observations. For each block group, we identify the nearest metropolitan center city hall, 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. The coefficients displayed are for the indicator variable of whether a block group is unincorporated. The vacancy rate, rental rate, percent commute over 60 minutes, and percent aged 65 and over, are all expressed as percentages ranging from 0–100. For more information on how we measure whether an area is incorporated, see Appendix section A. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are shown in parentheses. We use robust standard errors and cluster standard errors at the metro level for specifications that include metro fixed effects.

	(1)	(2)	(3)	(4)
Dependent Variable:	Median Home Value	Median Year Built	Total Housing Units Units	Units Built Since 2000
Is Incorporated	151.123	-6.354***	-25.643***	-40.952***
	(4357.225)	(0.273)	(5.532)	(4.085)

Table 5: Average Difference between Incorporated and Unincorporated

Table 6: National Sample Question Means

Panel A: Continuous Questions

Question	National			Income Tercile			Urban/Rural		
	Mean	Weight	Count	Low	Mid	High	Rural	Mix	Urban
How many zoning districts, including overlays, are in the municipality?	14	19	5,471	13	14	14	10	16	13
What is the longest frontage requirement for single family residential development in any district?	92	69	5,213	74	83	113	93	97	79
What is the highest residential minimum lot size? (Thousand Square Feet)	52	45	5,424	36	46	70	60	57	31
What is the mean residential minimum lot size? (Thousand Square Feet)	25	18	5,424	17	21	33	31	26	16
What is the lowest residential minimum lot size? (Thousand Square Feet)	10	6	5,440	7	9	13	12	10	8
How many mandatory steps are involved in the approval process for a typical new multi-family building?	4.5	4.3	5,791	4.4	4.4	4.7	4.5	4.5	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	5,759	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?	218	211	5,109	195	222	233	200	222	226

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	99	5,703	99	98	90	95	96	92
Are apartments above commercial (mixed use) allowed in any district?	63	71	5,717	64	67	57	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	5,795	80	83	80	65	84	83
Does zoning include any provisions for housing that is restricted by age?	45	60	5,129	34	42	57	24	50	52
Are accessory or in-law apartments allowed (by right or special permit) in any district?	33	39	5,781	27	33	37	23	41	22
Is cluster development, planned unit development, open space residential design, or another type of flexible zoning allowed by right?	9	10	5,797	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?	80	80	5,679	79	81	80	69	86	73
Does the zoning bylaw/ordinance include any mandates or incentives for development of affordable units?	24	50	5,540	10	20	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	5,803	11	17	22	10	19	16
Are there restrictions on counting wetlands, sloped land or easements in lot size calculations?	10	7	4,617	4	8	16	7	12	7
Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?	7	10	5,784	1	4	16	2	7	10
Are there townwide requirements for public hearings on any type of multi-family residential projects?	30	32	5,709	23	30	37	27	30	32

Notes: This table reports the averages of sample questions from our generated national regulatory dataset across a range of demographic associates. 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. We use median income from the 2021 Five-Year American Community Survey (`B19013_001E`). For continuous questions we upper winsorize at the 1% level for frontage, minimum lot sizes, and maximum potential review waiting time.

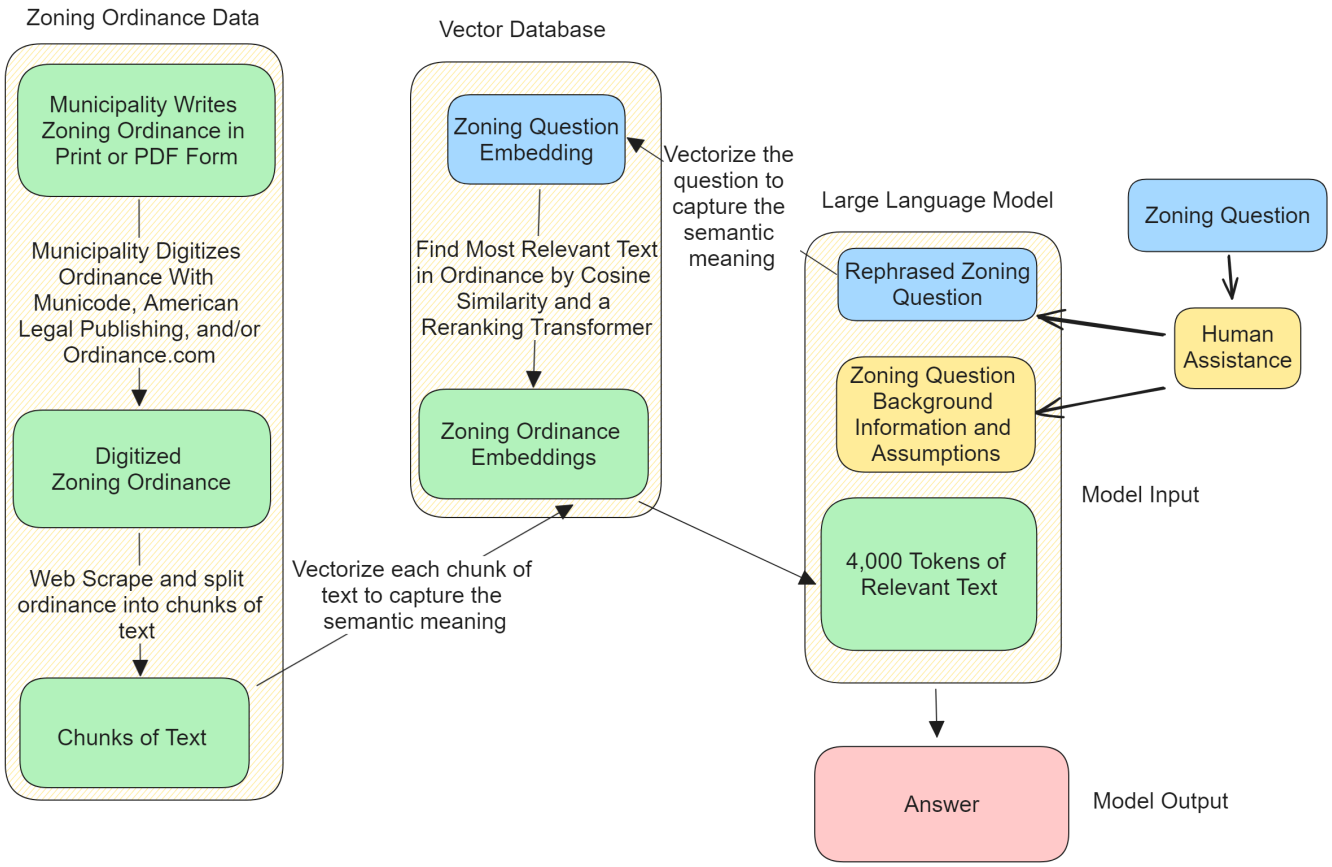
Table 7: Regression Results for Principal Components

	First PC		Second PC	
	No FE	MSA FE	No FE	MSA FE
Housing Density (Local Government Level)				
Housing Unit Density	0.01	-0.08**	-0.17***	-0.24***
	(0.01)	(0.03)	(0.03)	(0.05)
Share Structures with 2 or More Units	0.14***	0.09**	-0.30***	-0.34***
	(0.01)	(0.04)	(0.01)	(0.03)
Share Units Owner Occupied	0.01	-0.05	0.38***	0.32***
	(0.01)	(0.04)	(0.01)	(0.03)
Housing Density Index	0.06***	0.03	-0.34***	-0.34***
	(0.01)	(0.05)	(0.02)	(0.03)
Socioeconomic Characteristics (Local Government Level)				
White Share	-0.04***	-0.01	0.20***	0.22***
	(0.01)	(0.03)	(0.01)	(0.04)
Share Households Under 35	-0.05***	0.09**	-0.35***	-0.26***
	(0.01)	(0.04)	(0.01)	(0.04)
Share Low Income	-0.20***	-0.10***	-0.22***	-0.18***
	(0.01)	(0.03)	(0.01)	(0.05)
Socioeconomic Index	-0.11***	0.00	-0.35***	-0.28***
	(0.01)	(0.04)	(0.01)	(0.04)
Government Services (Local Government Level)				
Math Test Scores	0.23***	0.15***	0.29***	0.14***
	(0.01)	(0.04)	(0.01)	(0.03)
Local Revenue Per Student	0.12***	0.00	0.28***	0.17**
	(0.03)	(0.09)	(0.06)	(0.08)
Total Expenditure Per Capita	0.12***	0.04	0.12***	-0.06**
	(0.01)	(0.03)	(0.01)	(0.02)
Government Services Index	0.24***	0.12**	0.36***	0.13***
	(0.02)	(0.06)	(0.02)	(0.05)
Geographic Attributes (Local Government Level)				
Log Land Area	0.47***	0.47***	0.07***	-0.04
	(0.01)	(0.03)	(0.01)	(0.06)
Log Neighbors within 25 Miles	0.15***	0.21***	0.19***	-0.11***
	(0.01)	(0.06)	(0.01)	(0.04)
Log Miles to Metro Center	-0.14***	0.00	-0.03**	0.13***
	(0.01)	(0.04)	(0.01)	(0.04)
Geographic Attributes Index	-0.29***	-0.28***	-0.16***	0.16***
	(0.01)	(0.05)	(0.01)	(0.03)
Fundamental Amenities (County Level)				
Retail Establishments Index	0.46***		-0.03	
	(0.02)		(0.02)	
Natural Amenities Index	0.27***		0.01	
	(0.02)		(0.02)	
Patents per Capita	0.30***		0.01	
	(0.02)		(0.02)	
Combined Amenity Index	0.49***		-0.01	
	(0.02)		(0.02)	

For variable definitions, see Appendix Table A10. All right-hand side variables are measured as Z-scores. Fixed effects are for MSAs with State FE for municipalities not within an MSA. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are shown in parentheses. We cluster standard errors at the MSA level when including MSA fixed effects and use robust standard errors otherwise. County level regressions are based on population weighted county level regulations. The Natural Amenities Index comes from the USDA Natural Amenities Scale dataset and represents the first principal component of climate, topography, and water area features. The Retail Establishment Index combines per-capita counts of different types of retail businesses into a single measure using principal component analysis. Patents per Capita measures innovation activity using domestic patents granted between 2000-2015, normalized by each county's average population during that period. The Combined Amenity Index synthesizes all three measures (Natural Amenities, Patents, and Retail Establishments) into a single index using principal component analysis. Detailed components of the Natural Amenities and Retail Establishment indices can be found in Figure A11.

Figures

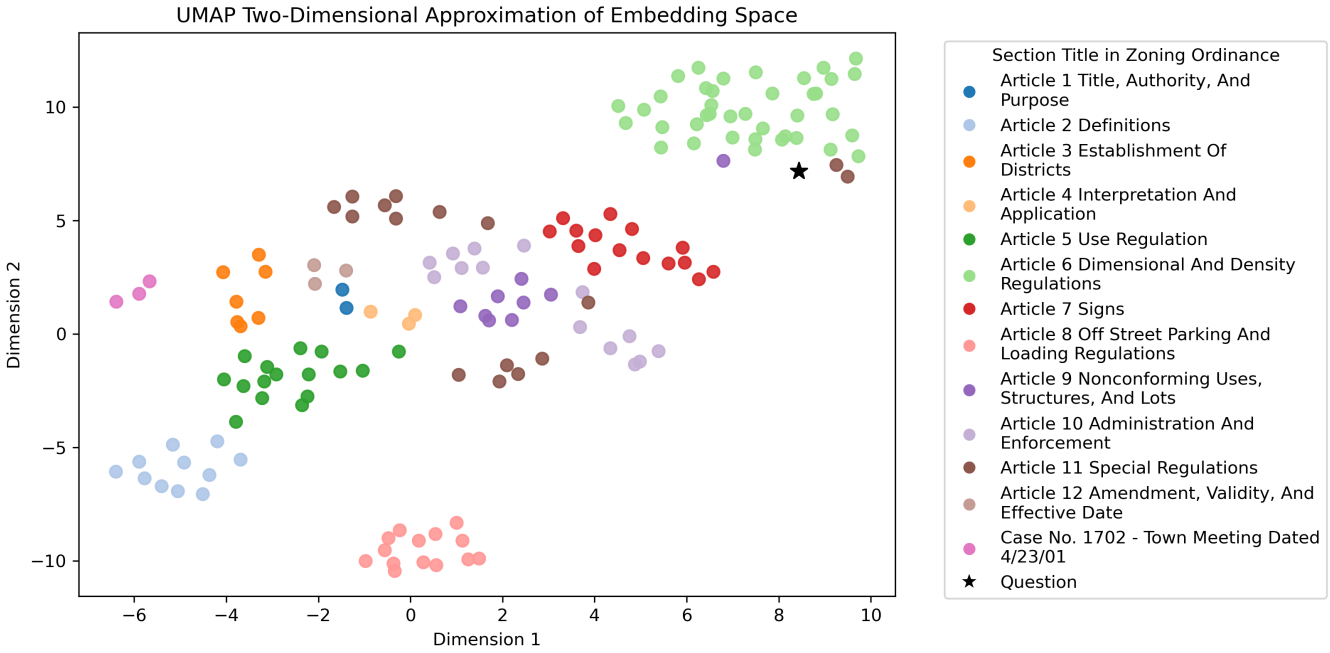
Figure 1: Model Overview



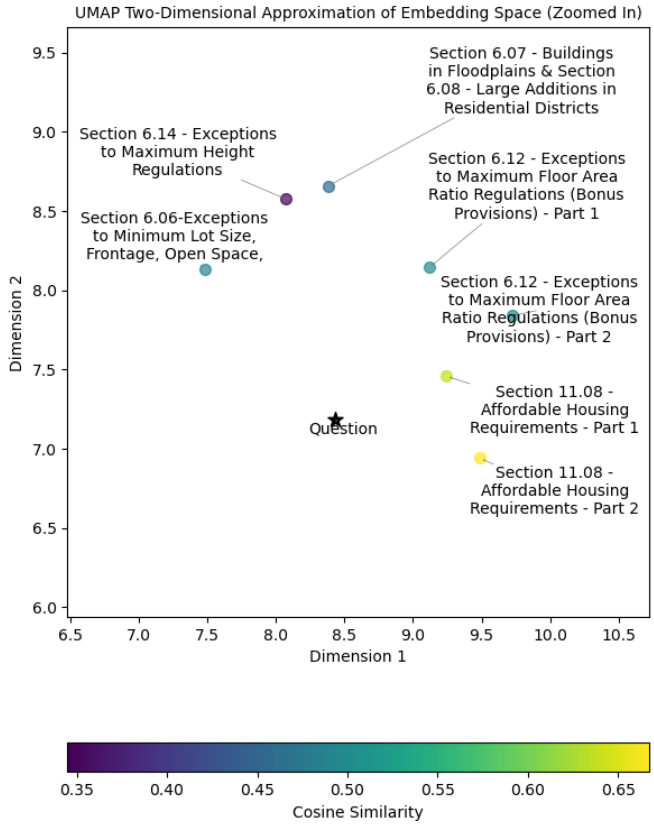
Notes: This Figure illustrates our overall process for Generative Regulatory Measurement. We scrape each section within a zoning ordinance separately. We split up sections that are longer than one thousand tokens into chunks of at most one thousand tokens. We also combine adjacent sections of less than 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 use the "text-embedding-3-large" algorithm. 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 for key phrases, table headers, and zoning district names (i.e., R-1 for the first residential zoning district). See Appendix Section A for further details on question background information and assumptions.

Figure 2: Embeddings of Municipal Code Sections

Panel A: Full Embedding Space

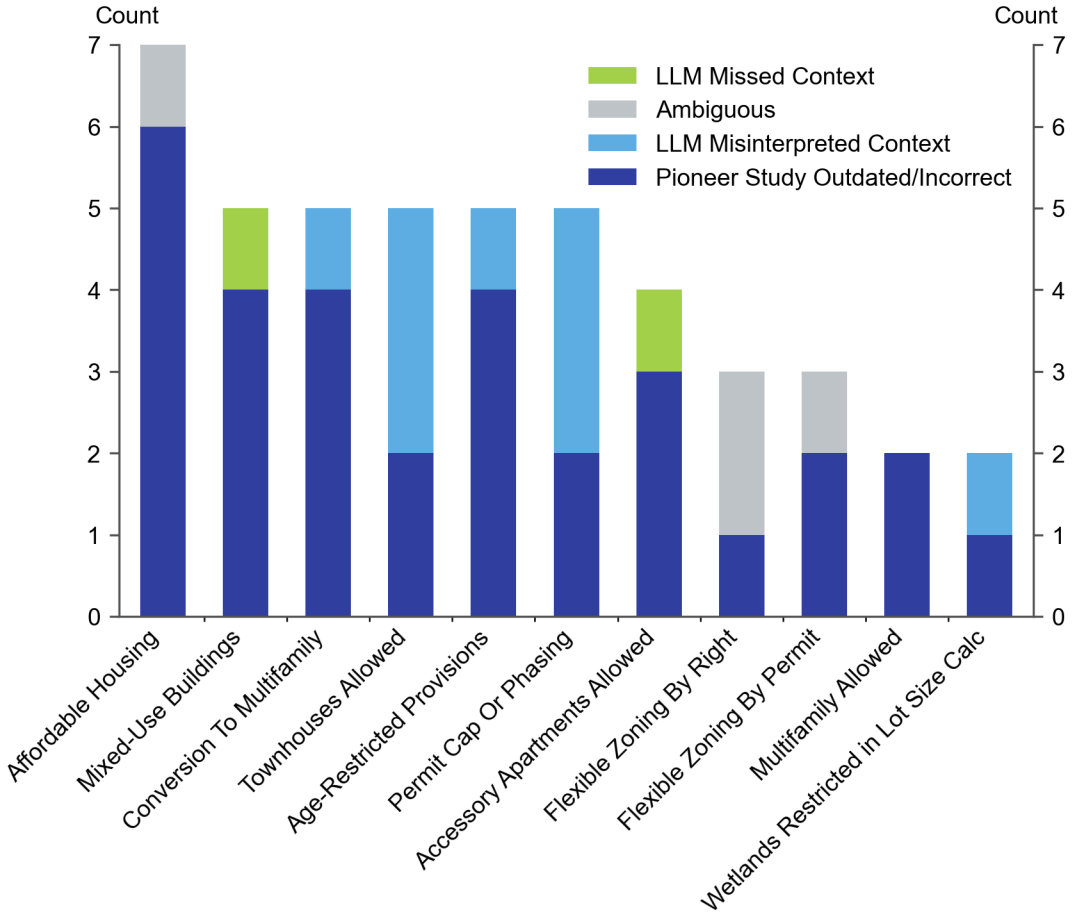


Panel B: Zoomed View of Similar Sections



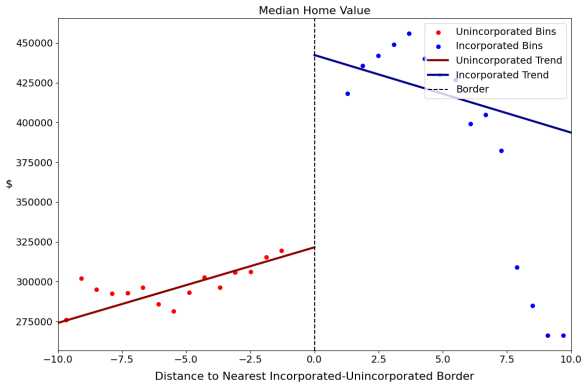
Notes: This Figure shows the two-dimensional projection of embeddings from sections of municipal codes for Arlington, Massachussets. Panel A shows the full embedding space, while Panel B provides a zoomed-in view of closely related sections. The embeddings are created using OpenAI's text-embedding-3-large model and visualized using UMAP dimensionality reduction.

Figure 3: Sources of Discrepancy Between LLM and Human Classifications of Zoning Regulations

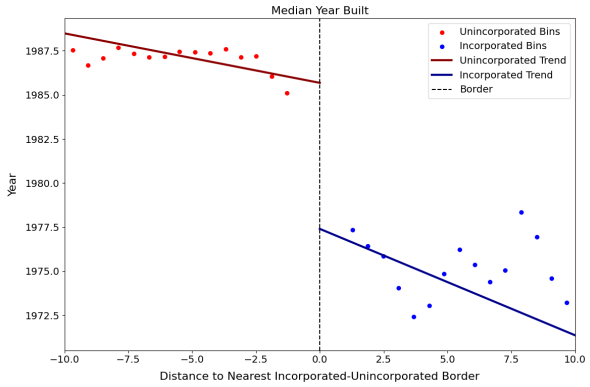


Notes: This Figure reports reasons for disagreement between our generated model data from GPT-4 Turbo and the original version of the Pioneer Institute data for binary questions. We first ran GPT-4 Turbo on the testing sample of 30 randomly selected municipalities that were included in the Pioneer Institute's 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. When measuring the performance of the model we adjust for disagreements where the Pioneer study was outdated/incorrect and also drop ambiguous cases.

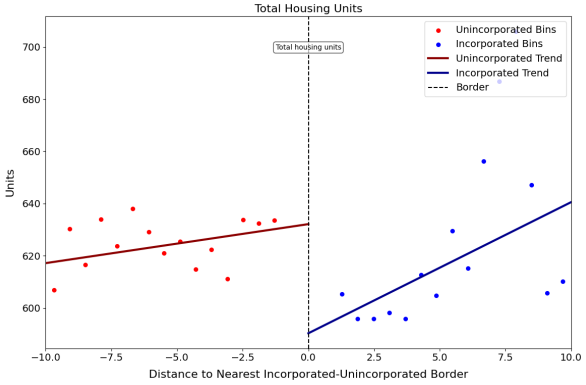
(a) Median Home Value



(b) Median Year Built



(c) Total Housing Units



(d) Units Built Since 2000

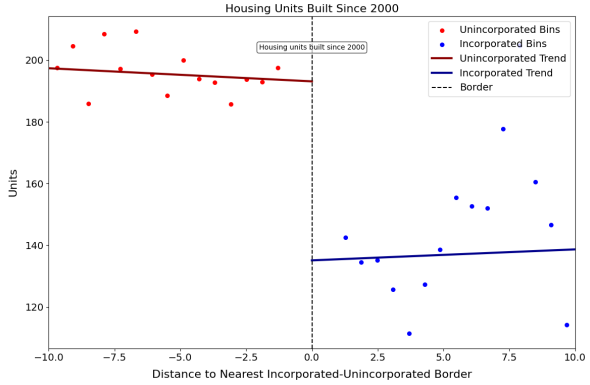
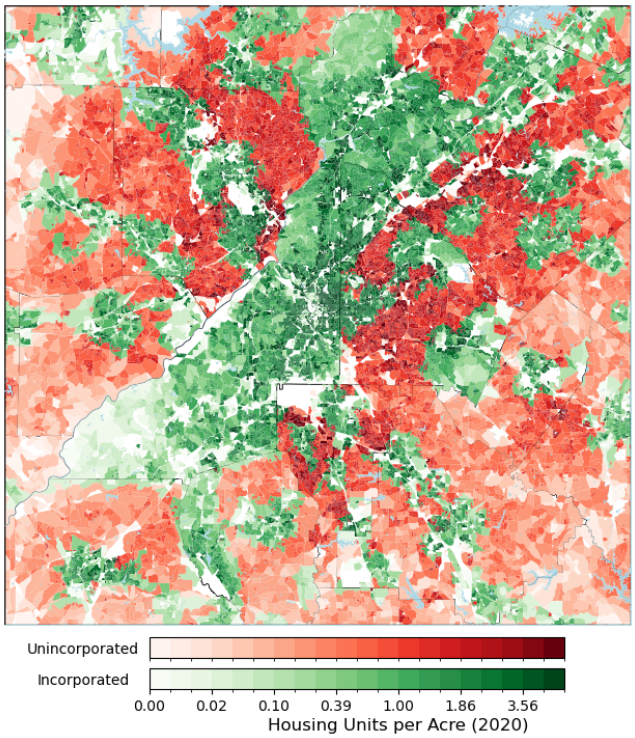


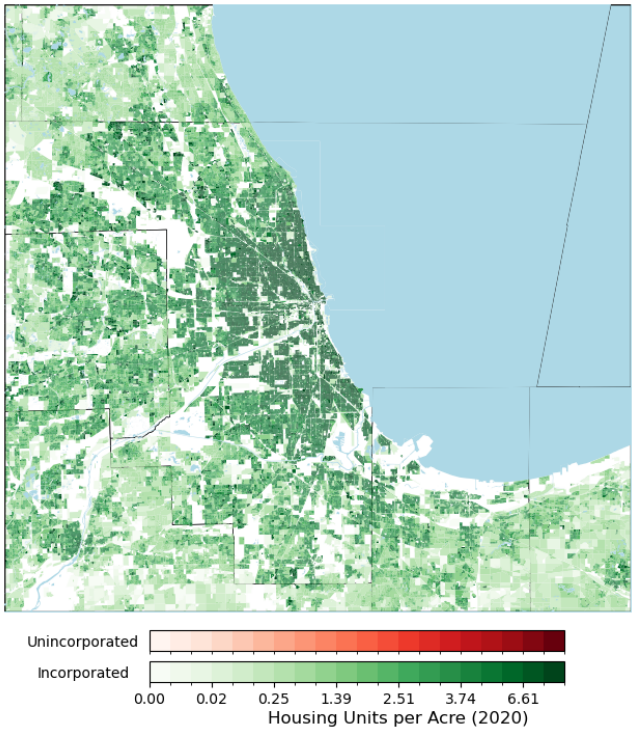
Figure 6: Binscatter Plots for Housing Outcomes

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

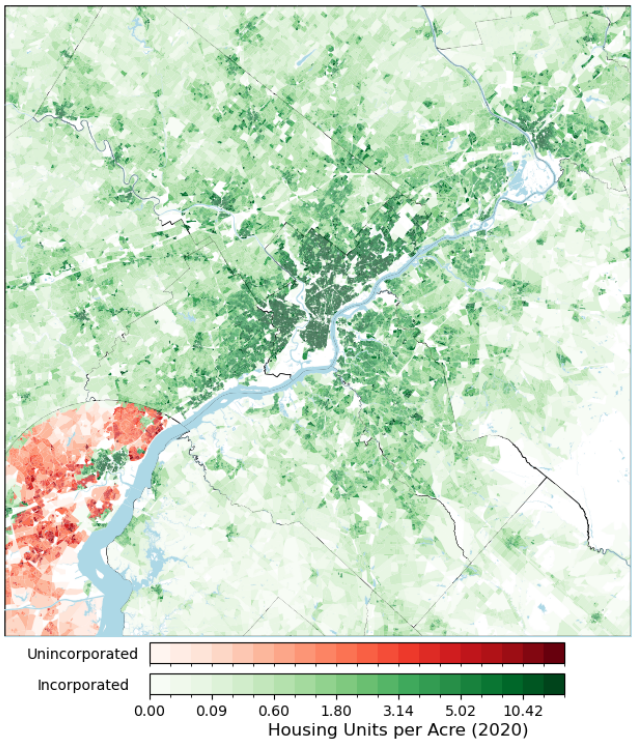
Panel A: Atlanta



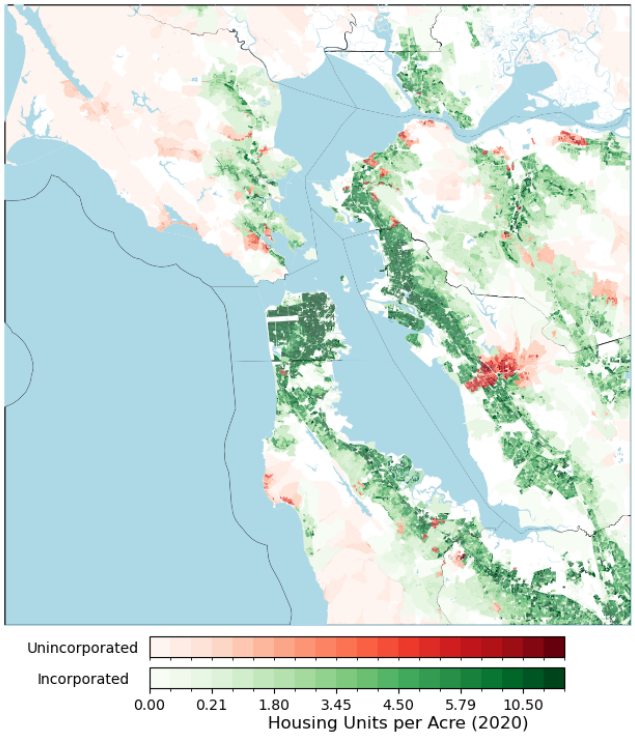
Panel B: Chicago



Panel C: Philadelphia



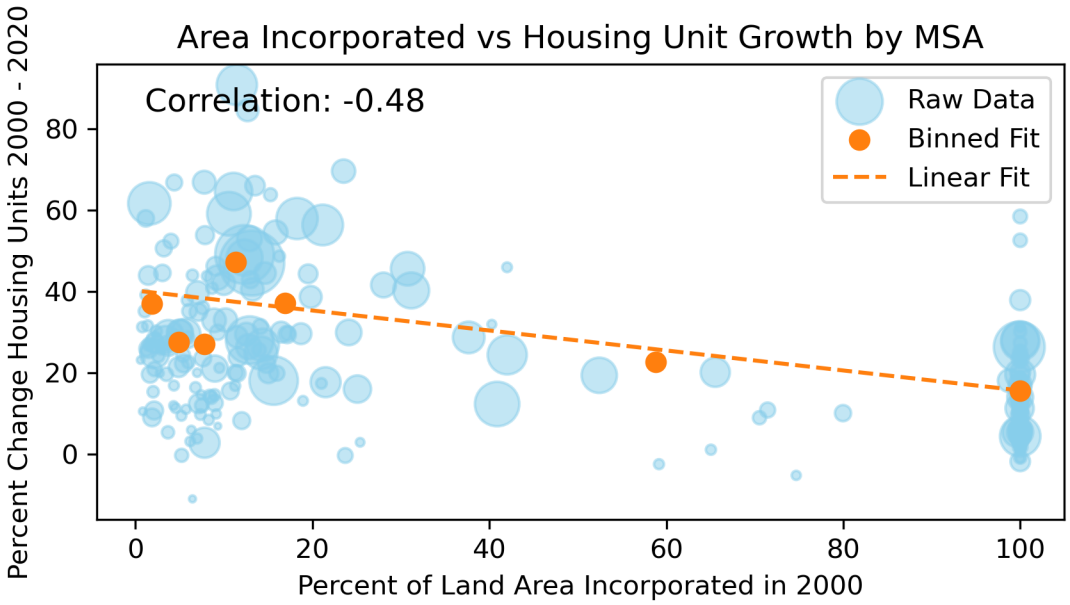
Panel D: San Francisco



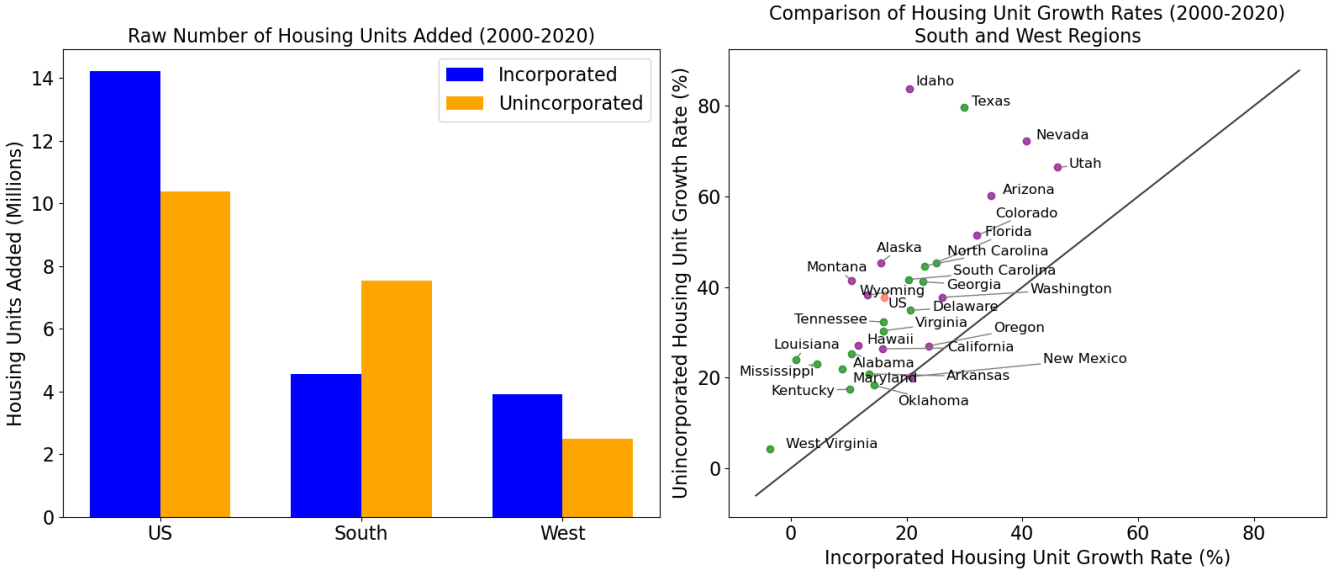
Notes: We plot housing unit density in four metropolitan regions: Atlanta (Panel A), Chicago (Panel B), Philadelphia (Panel C), and the San Francisco Bay Area (Panel D). Incorporated areas are colored in green, while unincorporated areas are colored in red. We plot housing unit density at the census block level drawn from 2022 ACS data and categorize blocks as incorporated if any of their area overlaps with an incorporated area and unincorporated otherwise. See Appendix Section A for further details about measuring incorporated areas.Each map shows roughly a 100km × 100km square area.

Figure 5: Housing Growth in Unincorporated Areas

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

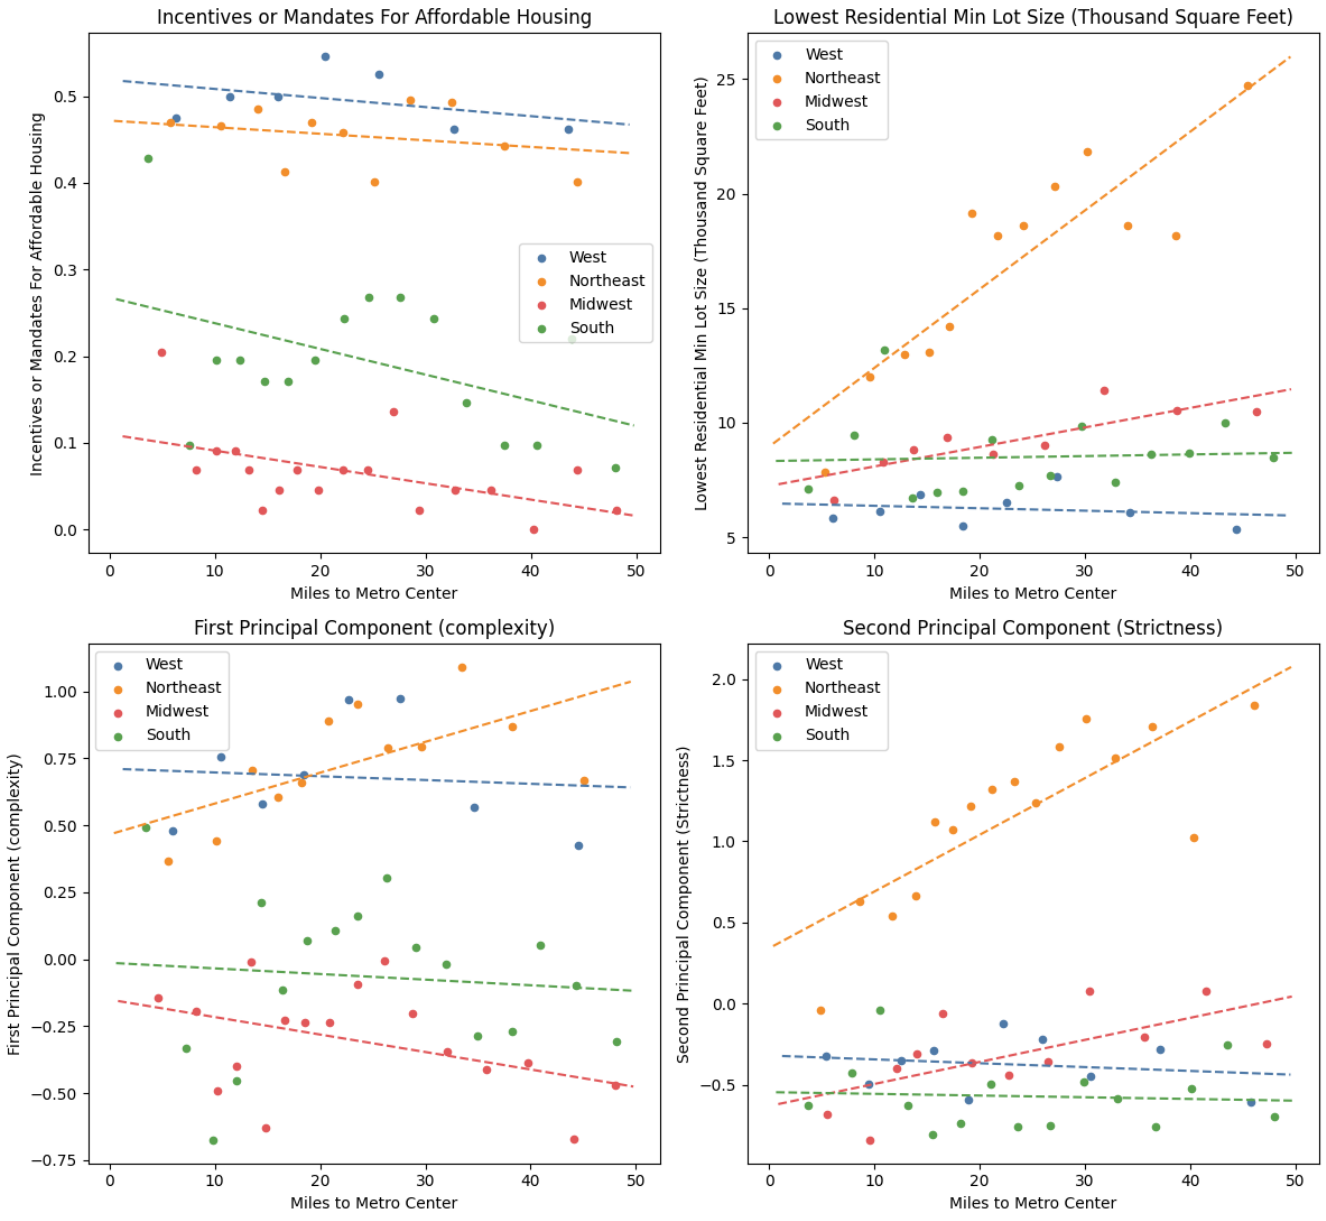


Panel B: Growth in Housing Units From 2000–2020, Incorporated vs Unincorporated



Notes: Panel A: For each MSA we calculate the percent of land area (county shape file ex water shape file) that intersects with the incorporated area in 2000 to calculate the percent of land area incorporated in 2000. Panel B: For each census block in both 2000 and 2020 we calculate the percent of its area that was incorporated in 2000. We then count the number of housing units that are incorporated by summing the product of housing units and percent incorporated for each block. We count the number of number of housing units that are unincorproated by summing the product of housing units and 1 - percent incorporated for each block. We focus on the South and West regions in this panel since practically all of the Northeast and Midwest are incorporated. See Appendix Section A for further details on how we measure incorporation.

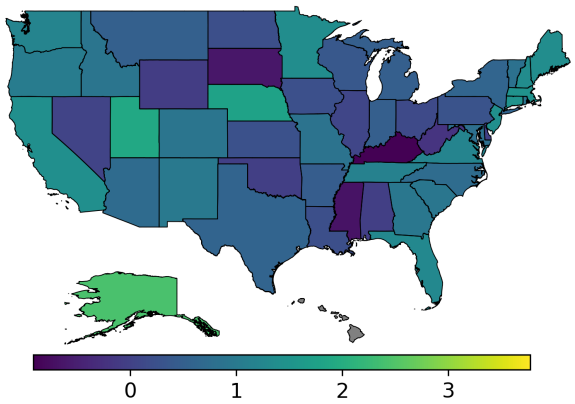
Figure 7: Spatial Variation of Zoning Regulations Relative to Metropolitan Centers



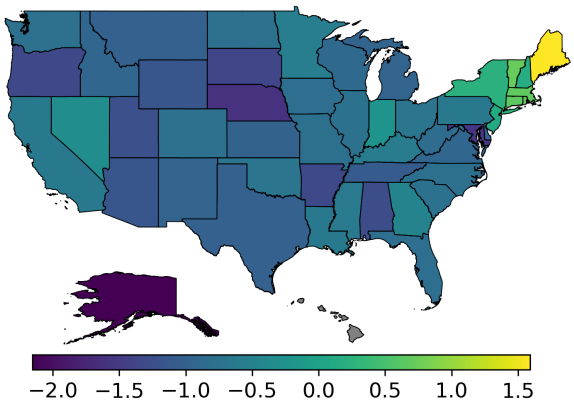
Notes: Note: We plot regulatory variables at the municipality level based on the distance from the center of the city for the 73 largest MSAs. We define city hall as the center of the city. We show whether a city has an affordable housing incentive or mandate; the minimum lot size; the first principal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

Figure 8: Nationwide Maps of Population-Weighted Averages

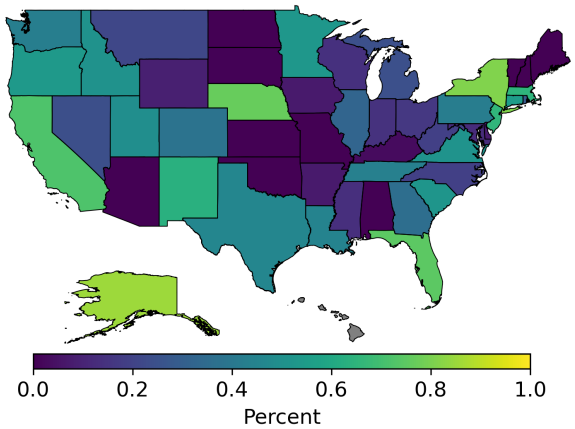
Panel A: First PC



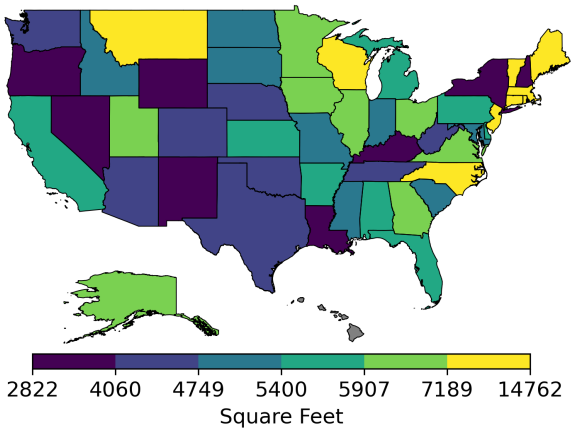
Panel B: Second PC



Panel C: Affordable Incentives/Mandates



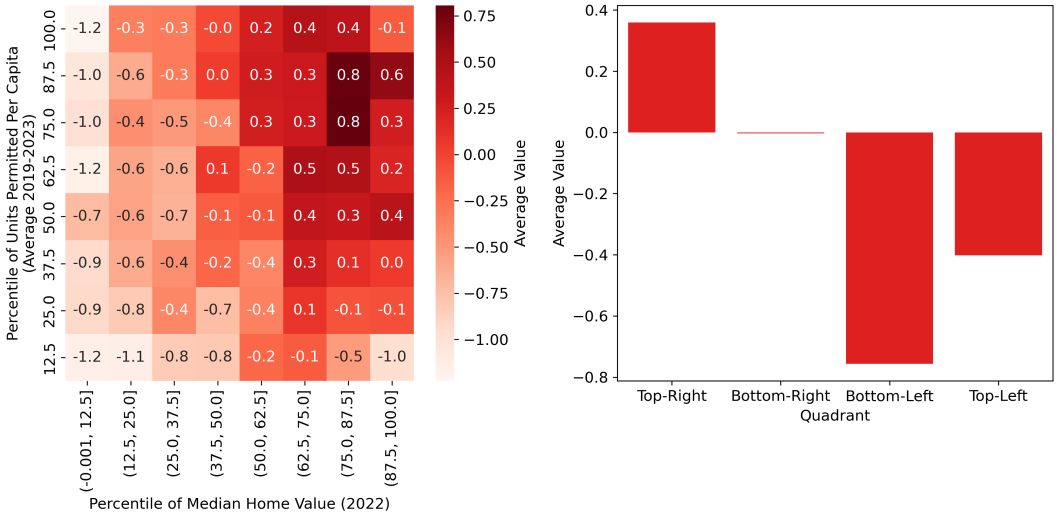
Panel D: Lowest Residential Minimum Lot Size



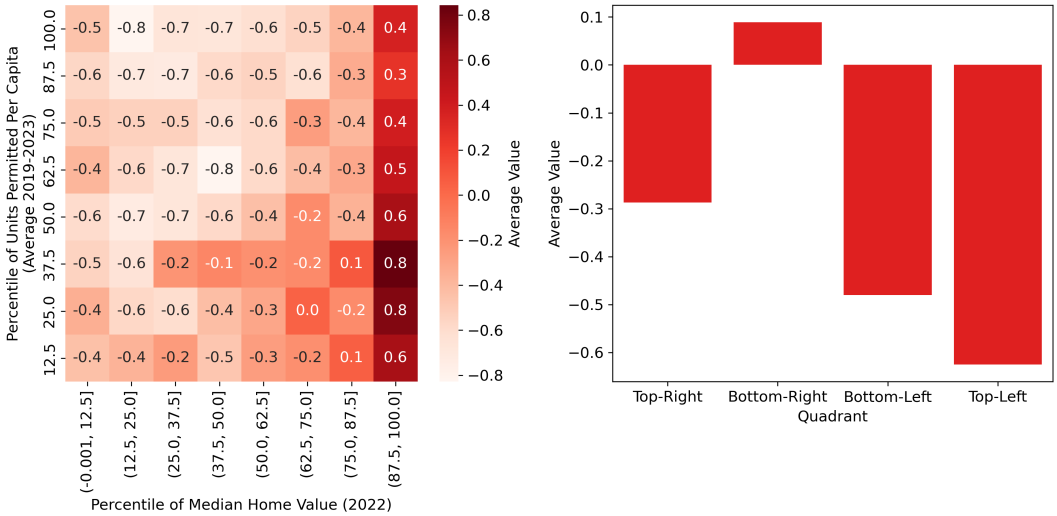
Notes: This Figure plots regulatory variables by state, weighted by local municipality population. We use the 2022 ACS Population as the population weight. Hawaii is grey because only one municipality (Honolulu) is in the dataset. For unweighted state maps please see Appendix Figure A9 and for county level maps please see Appendix Figure A10.

Figure 9: Regulatory Dimensions and Housing Market Dynamics

Panel A: First Principal Component (Value Capture)



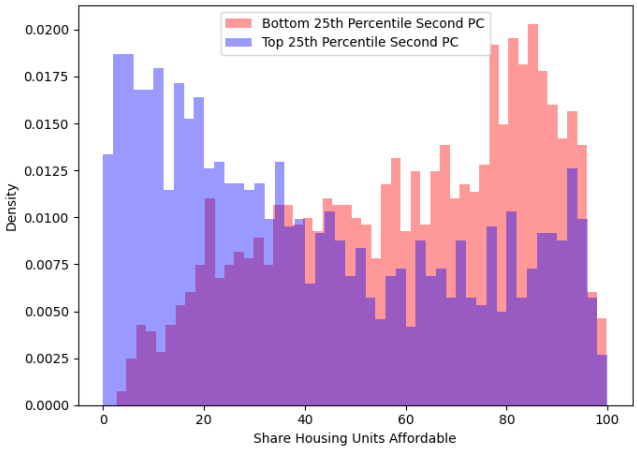
Panel B: Second Principal Component (Exclusionary Zoning)



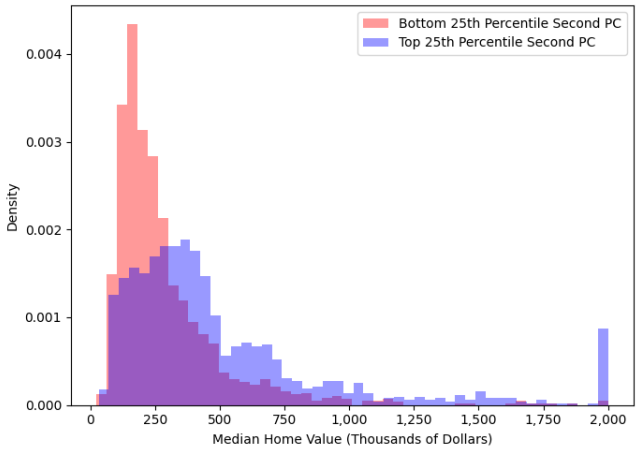
Notes: This Figure illustrates the relationship between two principal components of housing regulation and key housing market outcomes. Panel A shows the first principal component, associated with value capture. Panel B displays the second principal component, which corresponds to exclusionary zoning practices. The x-axis in both panels represents the percentile of median housing value, drawn from the 2022 ACS, while the y-axis measures the percentile of units permitted per capita (averaged from 2019–2023). Darker colors illustrate a larger correlation in the heatmap between each regulatory principal component and each coordinate of prices and building. The right figures show the overall correlation between each quadrant of the space of housing prices and quantities and the regulatory outcome. Please see Table A10 for details on variable definitions.

Figure 10: Second Principal Component and Housing Affordability

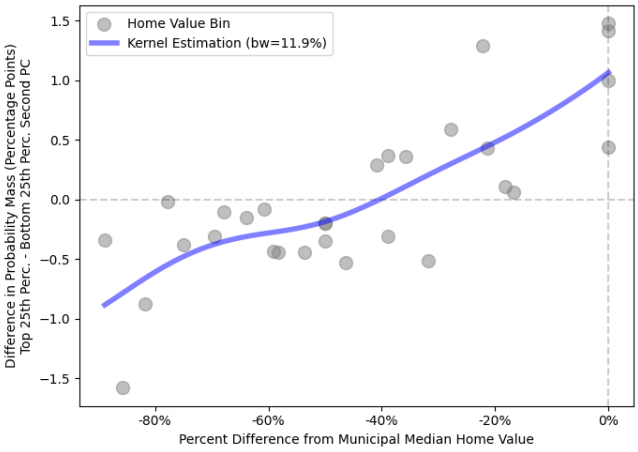
Panel A: Share Units Affordable



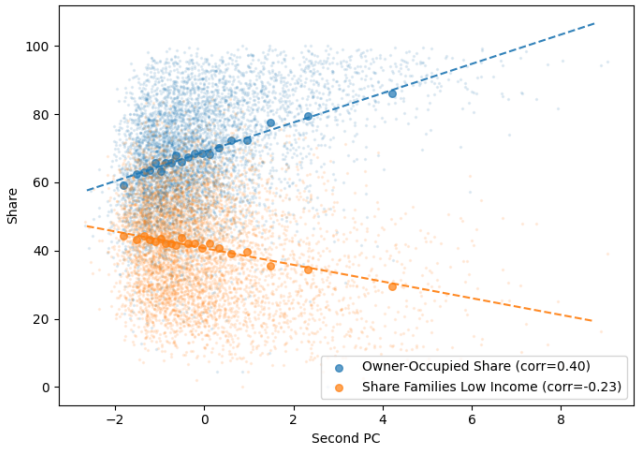
Panel B: Median Home Value



Panel C: Difference in Probability For Left Tail of Home Value

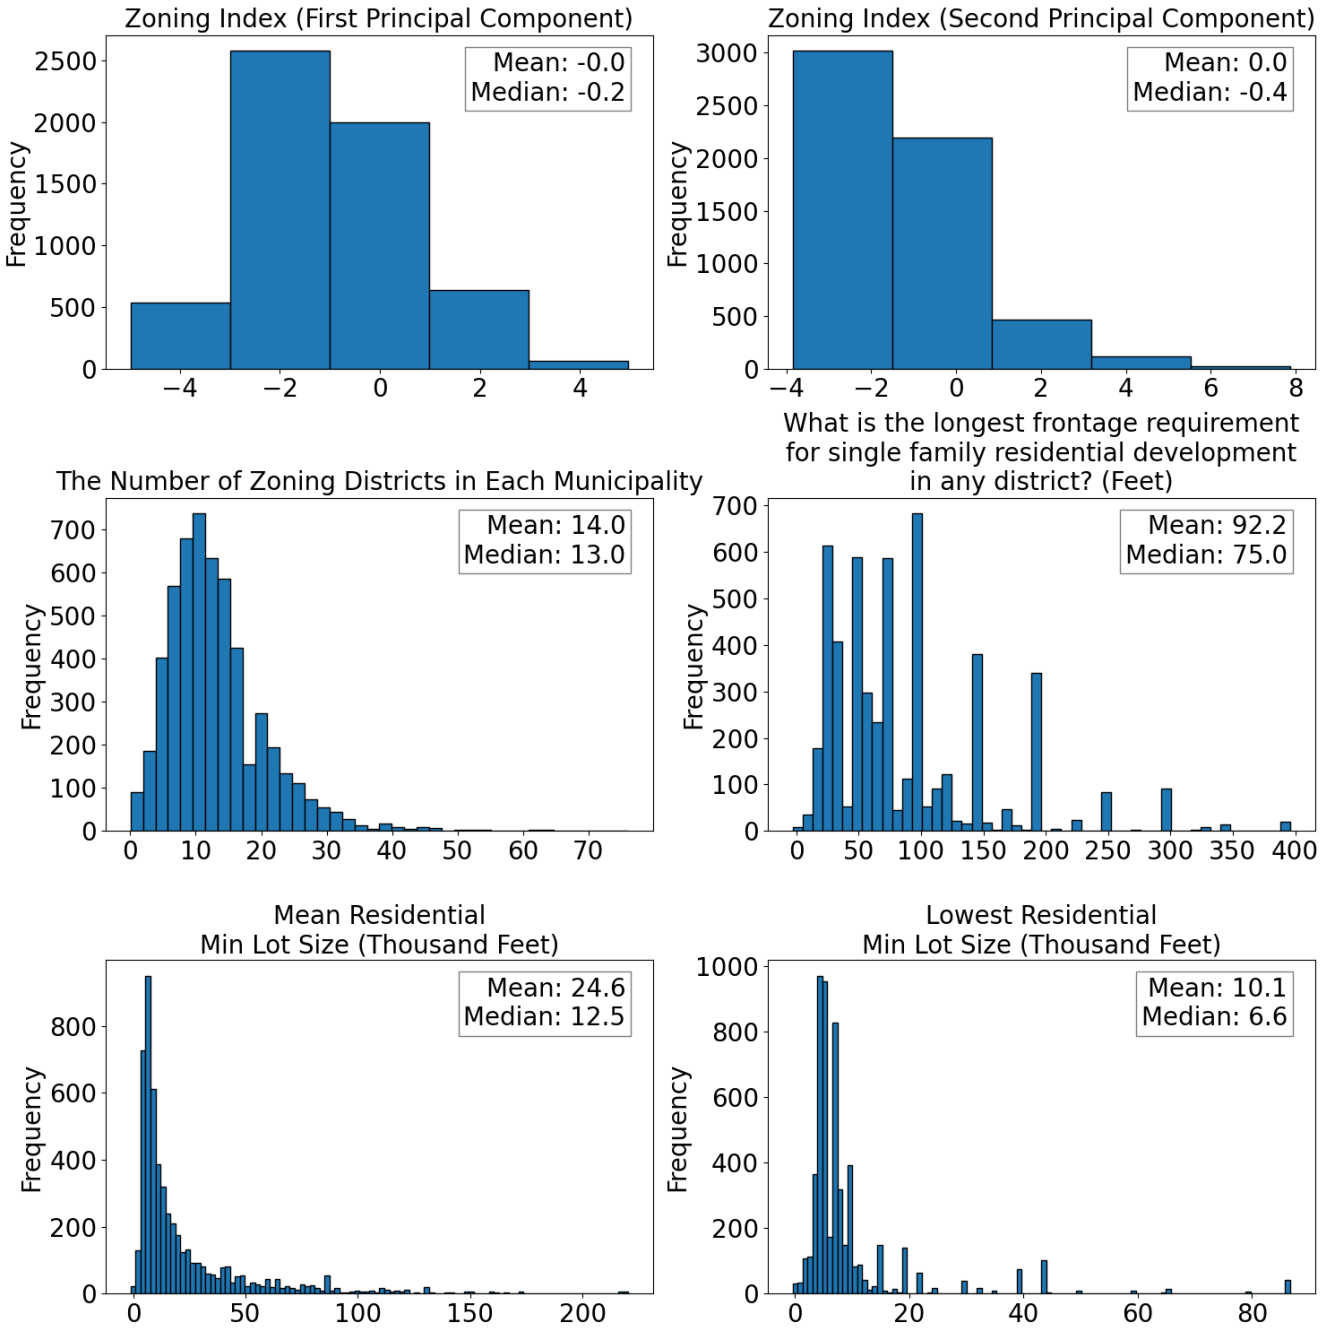


Panel D: Owner-Occupied Share and Low-Income Family Share by Second PC



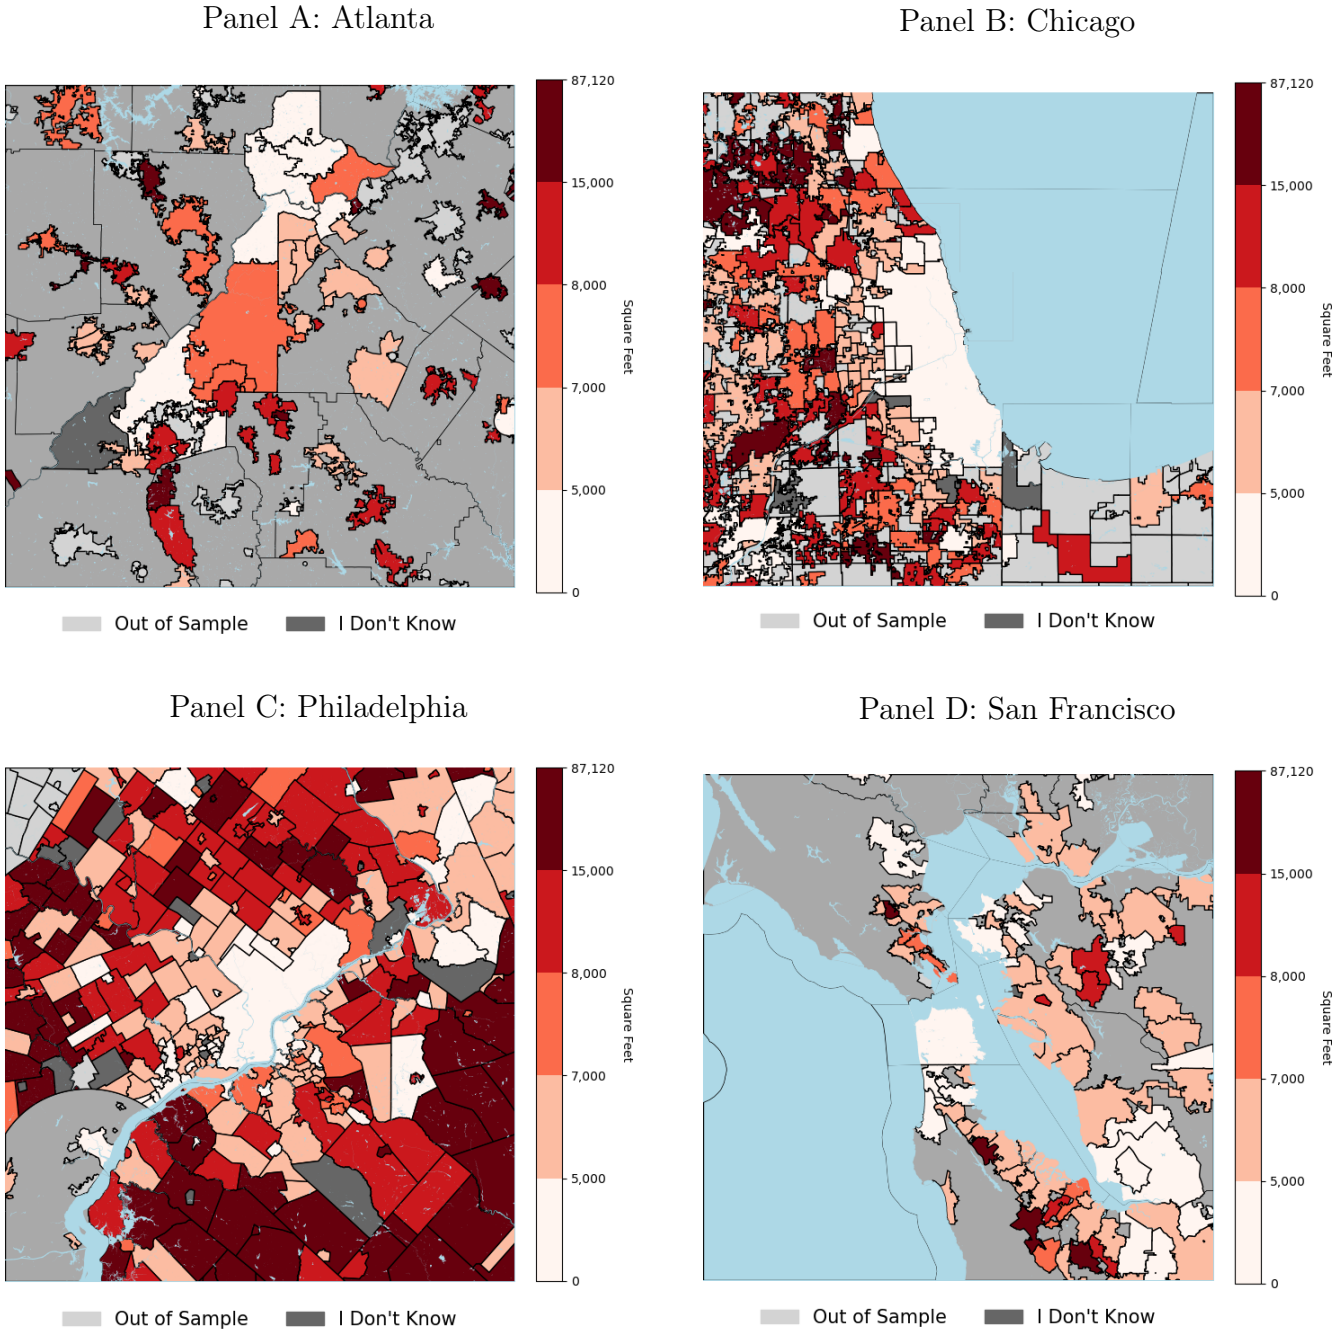
Notes: In Panel C, we compare the distribution of home values between municipalities with high and low second principal component (PC) scores using the distribution of owner-occupied housing units from the 2022 American Community Survey (ACS). We focus on municipalities with median home values between $200k and $500k, where the most overlap occurs in the second PC median home value histogram (Panel B). For each ACS home value range:$200k-–$249k, $250k-–$299k, $300k—$399k, and $400k—$499k—we filter the dataset to include only municipalities with median home values in that range. Within each filtered subset, we divide municipalities into the top and bottom 25th percentiles of the second PC. We then compute the difference in the distribution of home values between these two groups by calculating the percentage point difference in the probability mass of housing units for each ACS home value bin. To simplify the analysis, we collapse all home values below $100k into a single category with a midpoint of $50k. The x-axis measures the percentage difference between the bucket's median home value midpoint (e.g., $225k for the $200k—$249k bucket) and the midpoint of each ACS home value bin, while the y-axis represents the difference in probability mass between high and low second PC municipalities. Finally, a kernel regression, with bandwidth determined using Silverman's rule of thumb, is fitted to the plotted points to highlight the underlying trend.

Figure 11: Distribution of Zoning Indices and Housing Regulations



Notes: This Figure shows a histogram of the frequency of different regulatory measures at the municipality level. The first two panels show the distribution of the first two principal components. The third panel shows the distribution of the number of zoning districts. The fourth panel shows the longest frontage requirement. The fifth panel shows the average minimum lot size, and the sixth panel shows the minimum of all minimum lot size requirements. The last three variables are plotted in thousands of feet. See Table 6 footnote for details on the sample.

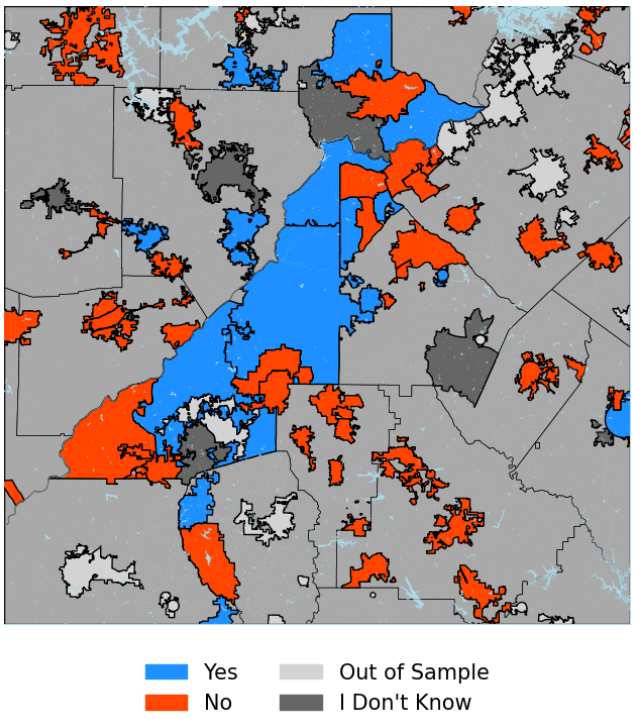
Figure 12: Lowest Residential Minimum Lot Sizes For Select Metro Areas



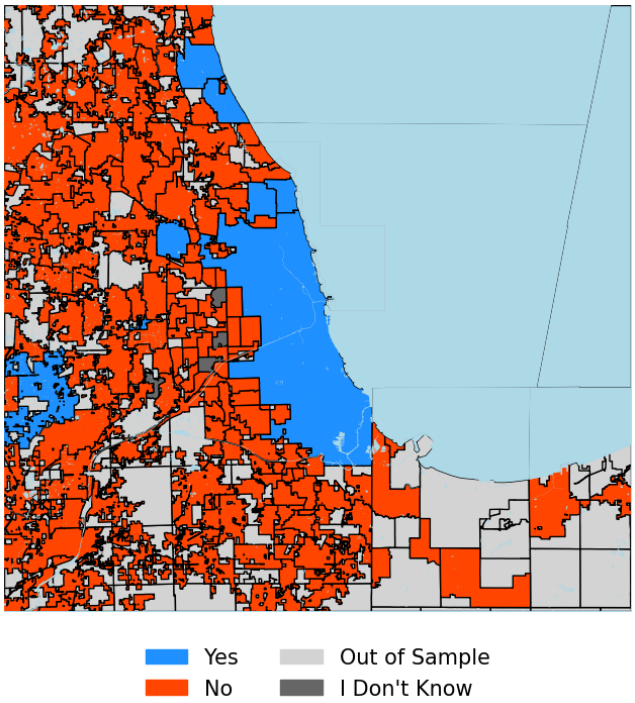
Notes: This Figure plots the lowest local government wide minimum lot size for each local government in the Atlanta (Panel A), Chicago (Panel B), Philadelphia (Panel C), and the San Francisco Bay Area (Panel D) metro areas. Within each map we plot all Census-designated places as well as county subdivisions that represent local governments. Both Census-designated place and Census county subdivisions data comes from the 2022 Census TIGER/Line shape files. Each map shows roughly a 100km × 100km square area. Non-Incorporated areas are shown in light grey, and areas for which the model reported "I don't know" are shown in dark gray.

Figure 13: Affordable Housing Incentives or Mandates For Select Metro Areas

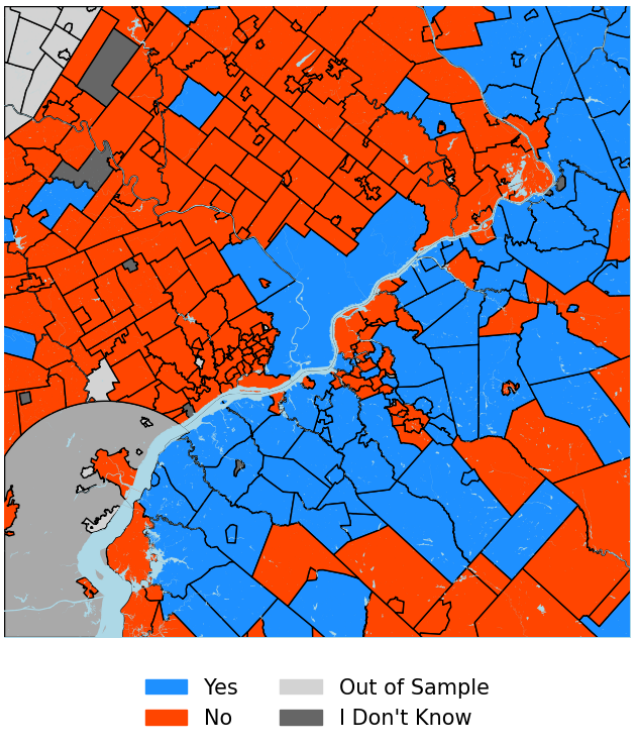
Panel A: Atlanta



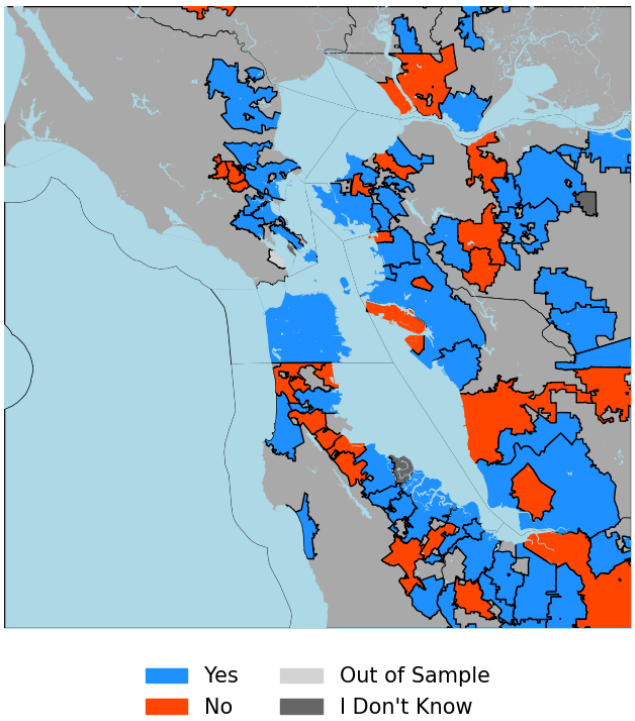
Panel B: Chicago



Panel C: Philadelphia



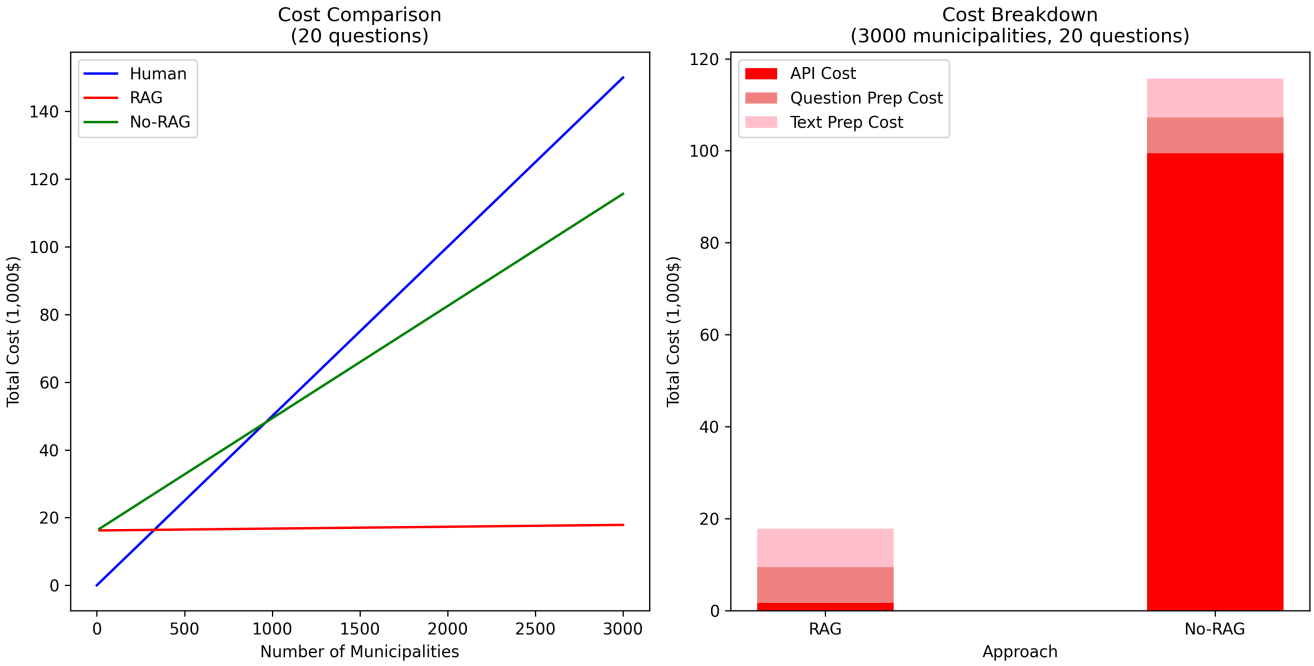
Panel D: San Francisco



Notes: This Figure plots the presence of affordable housing incentives for each local government in the Atlanta (Panel A), Chicago (Panel B), Philadelphia (Panel C), and the San Francisco Bay Area (Panel D) metro areas. Areas with incentives or mandates are plotted in blue; areas without such mandates are shown in red. Within each map we plot all Census-designated places as well as county subdivisions that represent local governments. Both Census-designated place and Census county subdivisions data comes from the 2022 Census TIGER/Line shape files. Each map shows roughly a 100km × 100km square area. Non-Incorporated areas are shown in light grey, and areas for which the model reported "I don't know" are shown in dark gray.

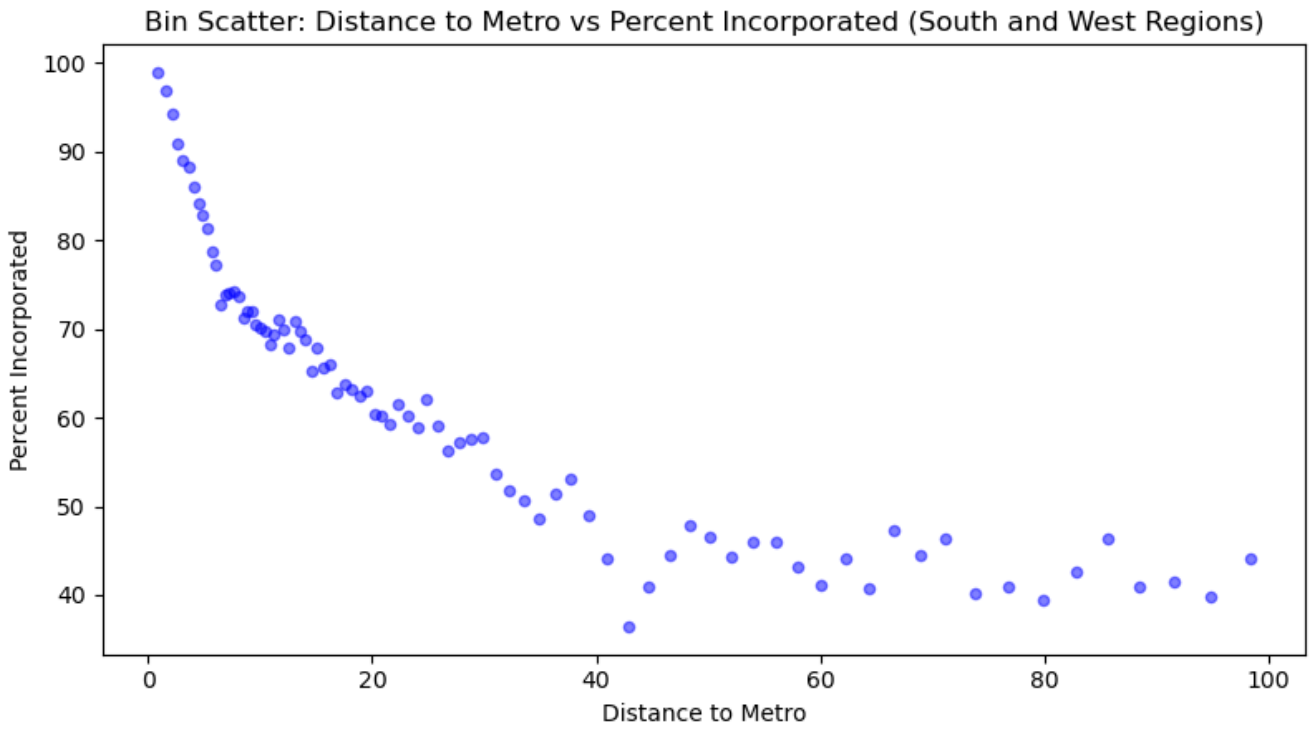
A Appendix A: Additional Graphs and Tables

Figure A1: Comparison of Human vs. LLM Cost



Notes: This figure compares the cost of human versus LLM-based analysis of zoning regulations across different numbers of municipalities. We assume: (1) Lawyers take an average of five minutes per question-municipality pair, compensated at $50/hour; (2) For the RAG approach, each LLM API call involves 4,000 input tokens and 500 output tokens, with costs of $5 and $15 per million tokens respectively; (3) Question preparation costs $390/question, including lawyer time for answering 60 municipalities (30 for training, 30 for testing) and 10 hours of human labor at $140 to train the model; (4) Scraping, cleaning, and embedding each municipality's ordinance costs $2.80. For the No-RAG approach, input tokens average 330,000 per question-municipality pair based on mean ordinance length. All costs are assumed to scale linearly with the number of municipalities and questions, shown here for 20 questions.

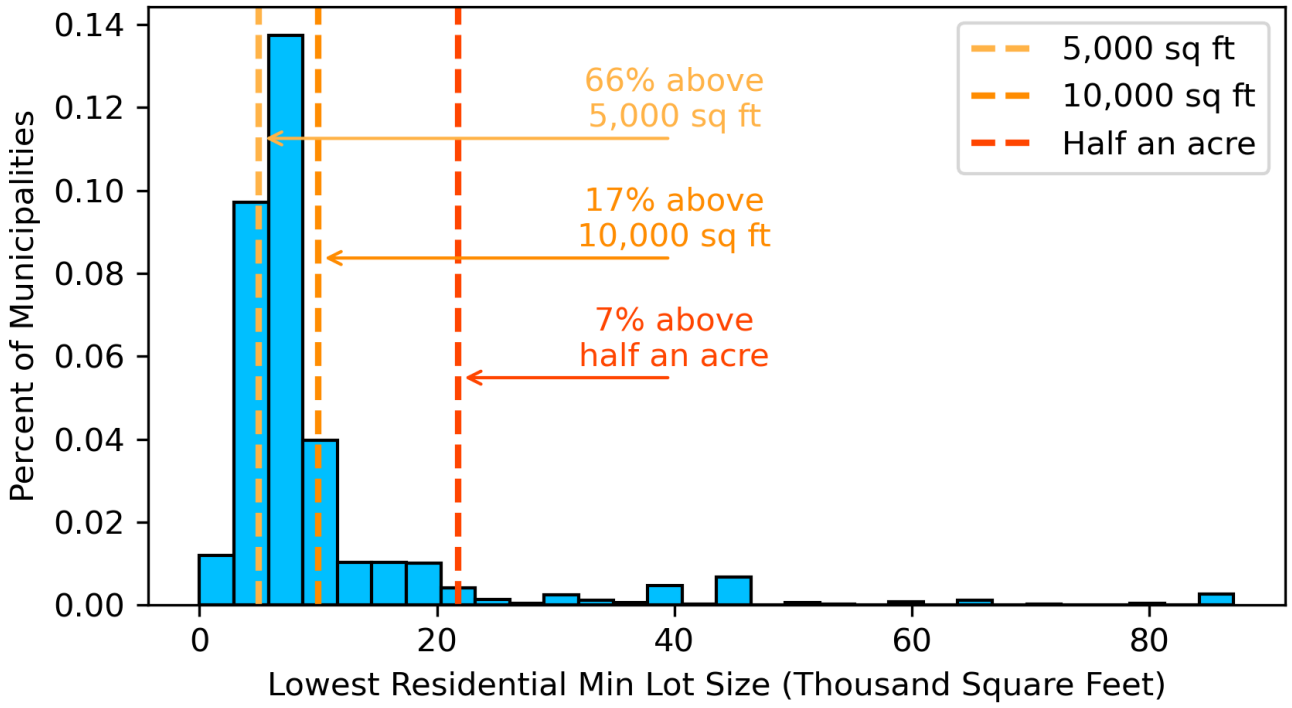
Figure A2: Municipal Incorporation Status Relative to Metropolitan Core Distance



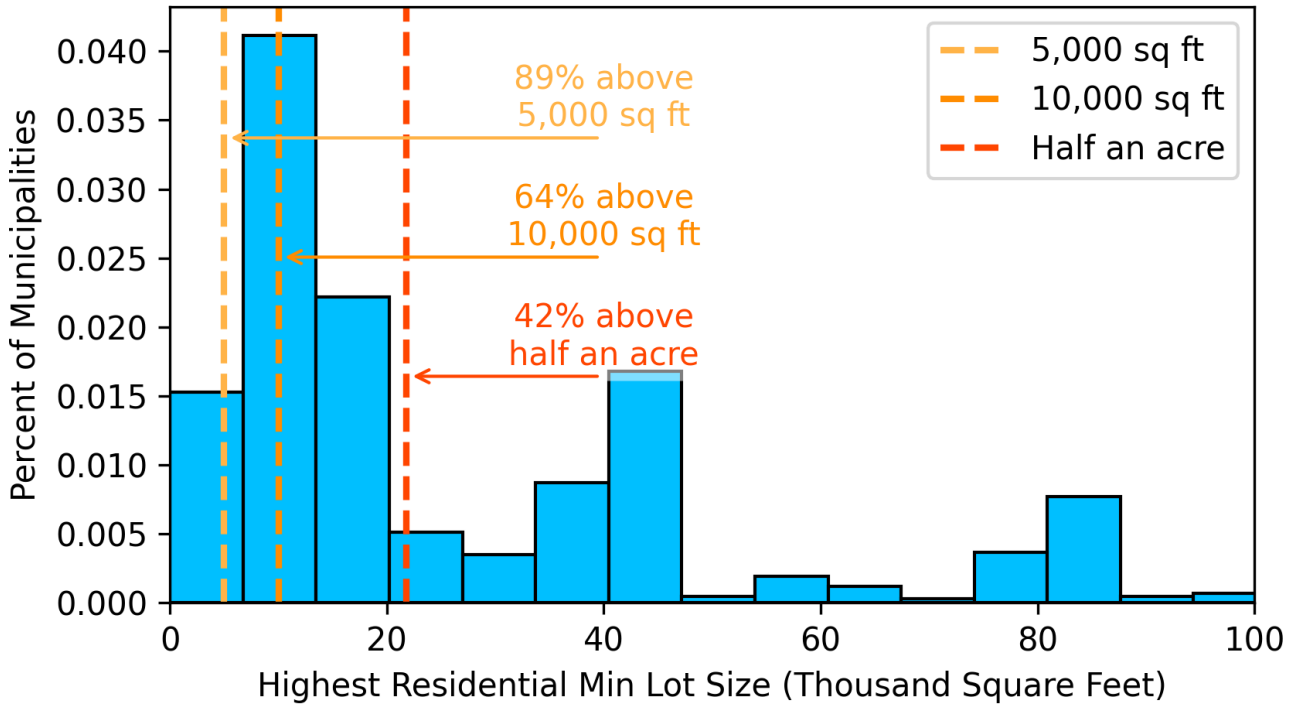
Notes: This Figure illustrates the relationship between municipal incorporation status and distance from metropolitan centers. The x-axis represents the log distance from the nearest metropolitan center, measured in miles, which is defined as City Hall. The y-axis shows the probability of a block group being incorporated, ranging from 0 to 1. Each point represents a binned average, with the blue line indicating the fitted relationship. For details on incorporation measurement methodology, see Appendix Section A.

Figure A3: Minimum Lot Size Distribution

Panel A: Lowest Residential Min Lot Size

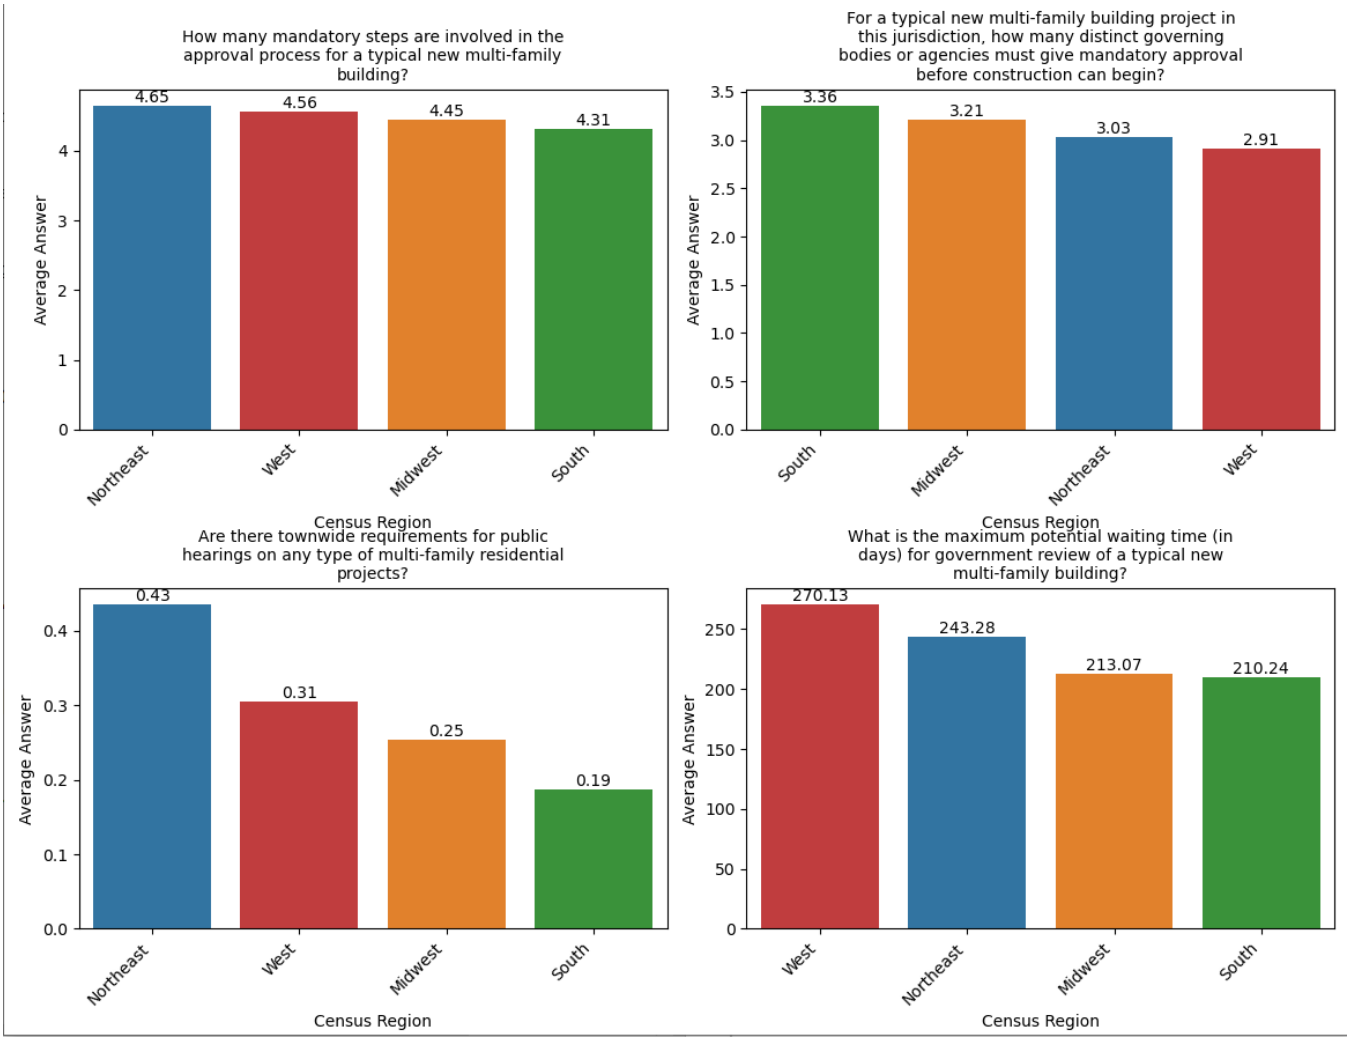


Panel B: Highest Residential Min Lot Size



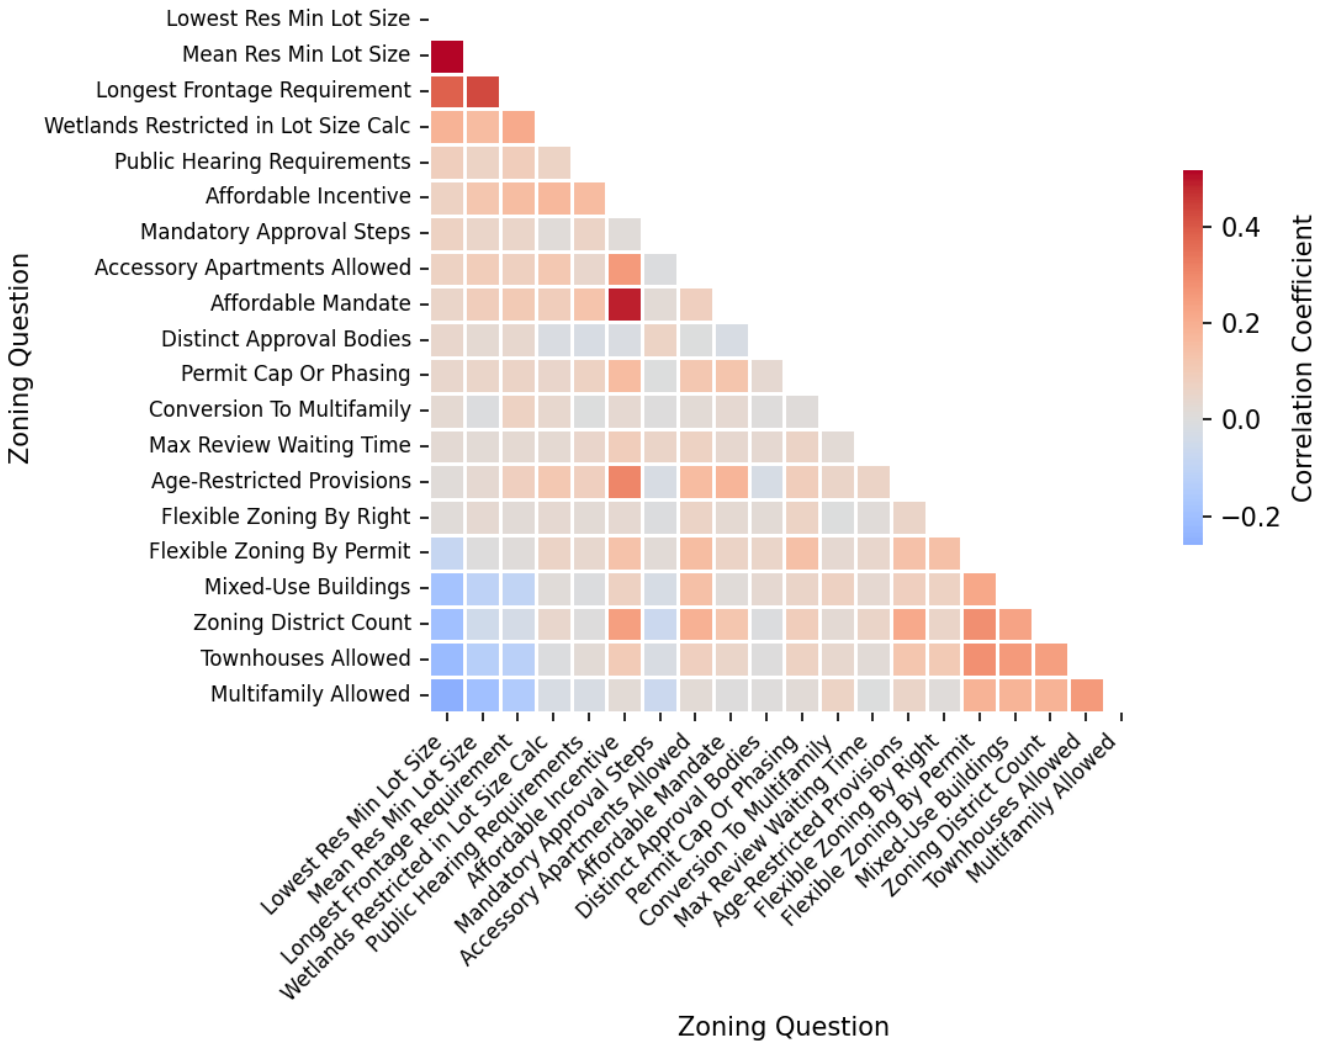
Notes: This histogram shows the distribution of minimum lot size requirements across municipalities. The x-axis represents lot size in square feet, with key thresholds labeled. The y-axis shows the percentage of municipalities falling into each lot size category. Vertical lines mark important thresholds: 5,000 sq ft (common suburban lot size), 10,000 sq ft (quarter-acre), and 21,780 sq ft (half-acre). The x-axis in Panel B stops at 100 thousand square feet, though 12% of municipalities have highest residential min lot sizes above this level.

Figure A4: Measuring Housing Process Variation



Notes:. This histogram shows the distribution of four housing process questions which augmented the initial Pioneer Institute set of questions. Survey questions are grouped into Census Region.

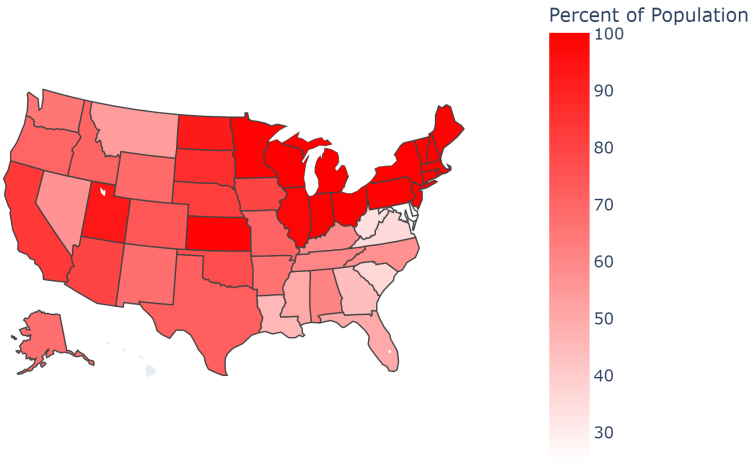
Figure A5: Pairwise Correlations Between Zoning Questions



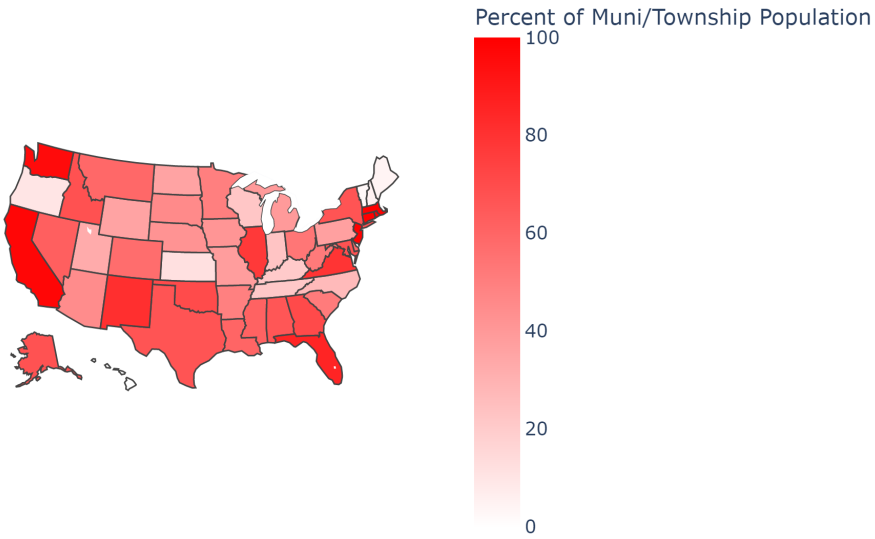
Notes: This heatmap illustrates the pairwise correlations between various zoning regulations across U.S. municipalities. Each cell represents the correlation coefficient between two zoning measures, with color intensity indicating the strength and direction of the relationship. Darker red indicates stronger positive correlations, while darker blue represents stronger negative correlations. White or light-colored cells suggest weak or no correlation. The diagonal represents perfect correlation of each measure with itself. See Appendix Table A1 for full question names. See table 6 footnote for details on sample construction.

Figure A6: Population Coverage

Panel A: Percent of the Population Living in Either a Municipality or Township Government By State

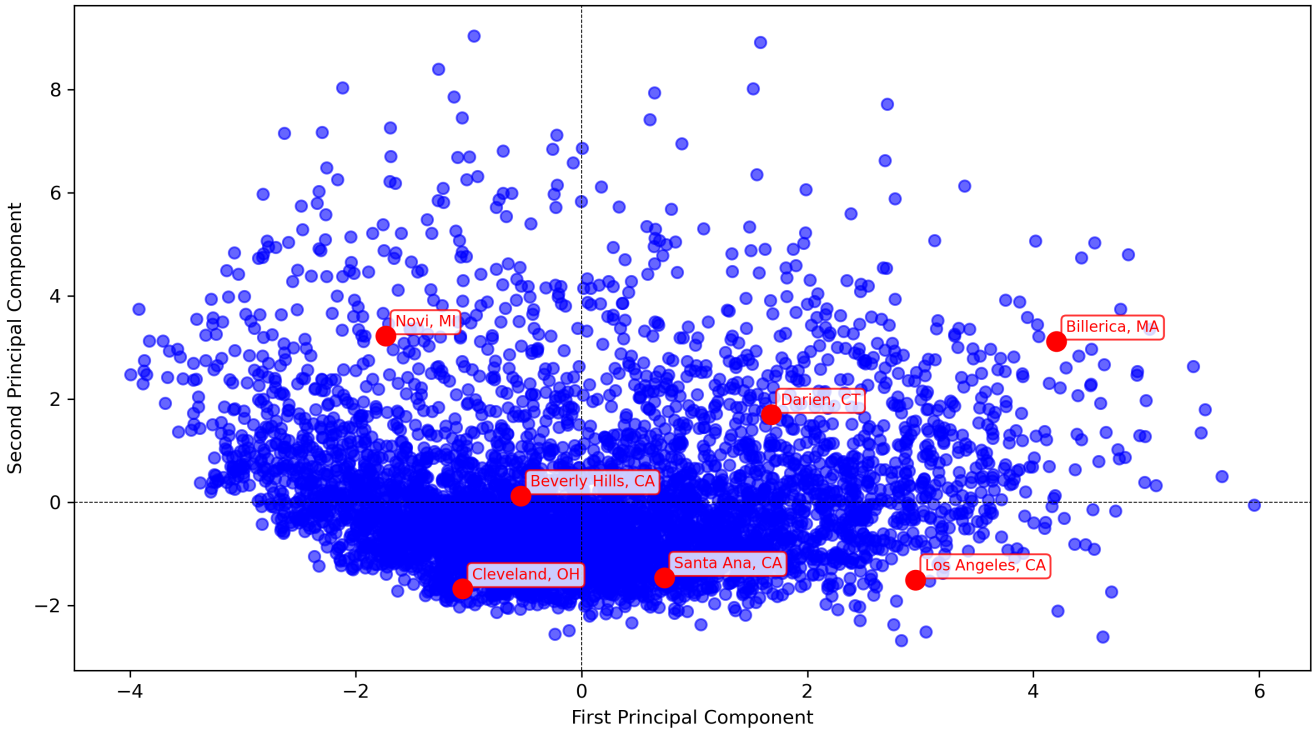


Panel B: State-Level Coverage: Percentage of Municipality and Township Residents in Our Sample



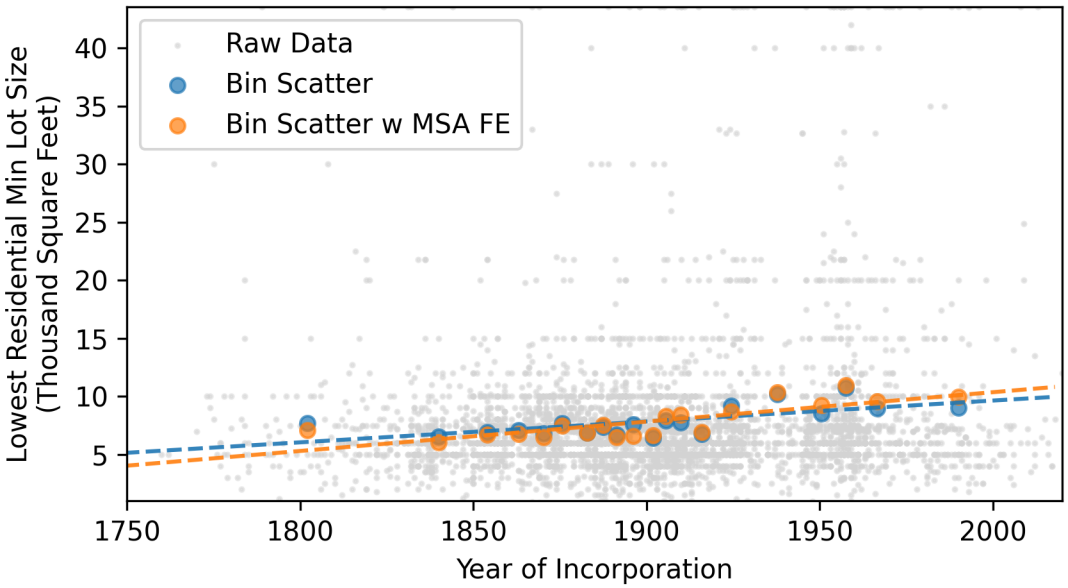
Notes: Panel A of this Figure shows the fraction of the population in each state who resides in a municipality or township government by state. The remaining fraction of the population live in an unincorporated area. Panel B of this Figure shows the fraction of the population in each state, who lives in a municipality or township, which we have municipal code coverage for. See Table 1 footnote for more details on sample coverage.

Figure A7: Relationship between First and Second Principal Components



Notes: This scatter plot illustrates the relationship between the first two principal components of our zoning regulation analysis across U.S. municipalities. The x-axis represents the first principal component (PC1), which we interpret as a measure of regulatory complexity and housing demand. The y-axis shows the second principal component (PC2), which corresponds to exclusionary zoning practices. Each point represents a municipality. Municipalities in the upper right quadrant tend to have both complex regulations and more exclusionary practices, while those in the lower left have simpler, less restrictive zoning.

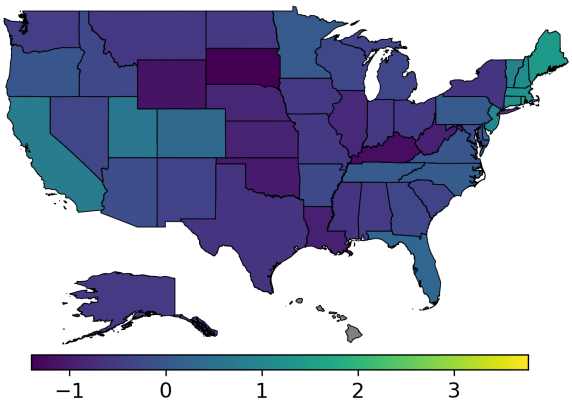
Figure A8: Historical Trends in Minimum Lot Size Requirements by Incorporation Date



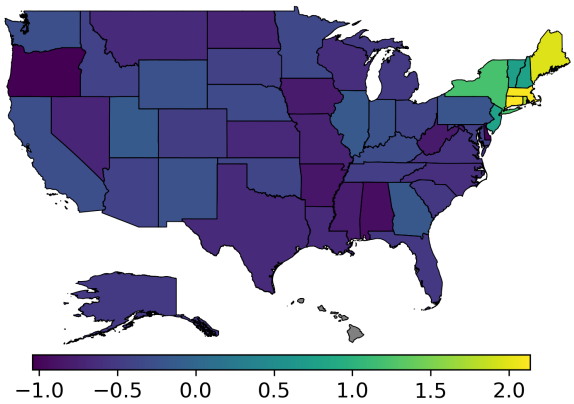
Notes: This scatter plot illustrates the relationship between a municipality's year of incorporation and its current minimum lot size requirement. The x-axis spans from 1750 to present, capturing the majority of U.S. municipal incorporations. The y-axis shows minimum lot sizes from 1,000 to 43,560 square feet (1 acre), though some municipalities have larger requirements not shown here. Each point represents a municipality. Blue dots indicate binscatter points and blue lines indicate a linear trend in raw data, while the orange line includes a MSA fixed effect. Note that while the visible plot is truncated, the regression lines are based on the full dataset, including earlier incorporations and larger lot sizes.

Figure A9: Nationwide Maps of Regulatory Variables

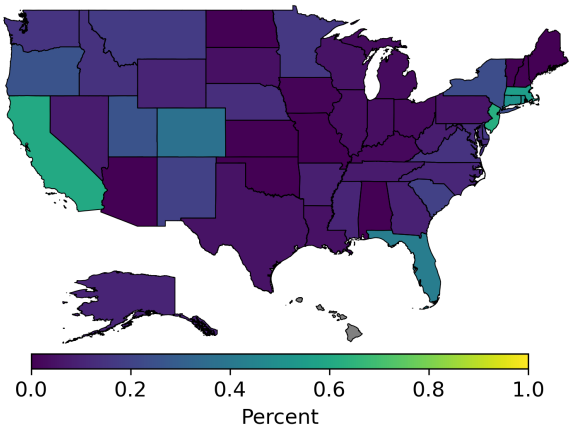
Panel A: First PC



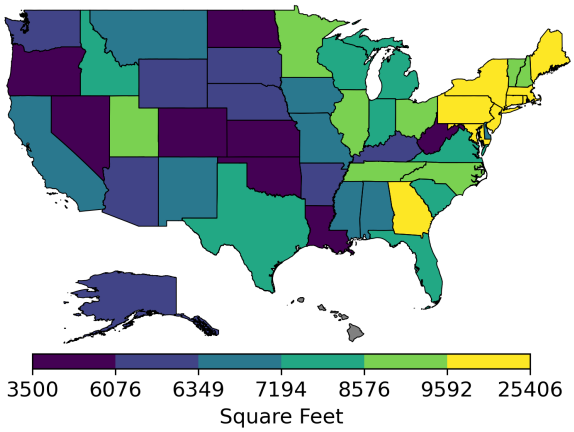
Panel B: Second PC



Panel C: Affordable Incentives/Mandates

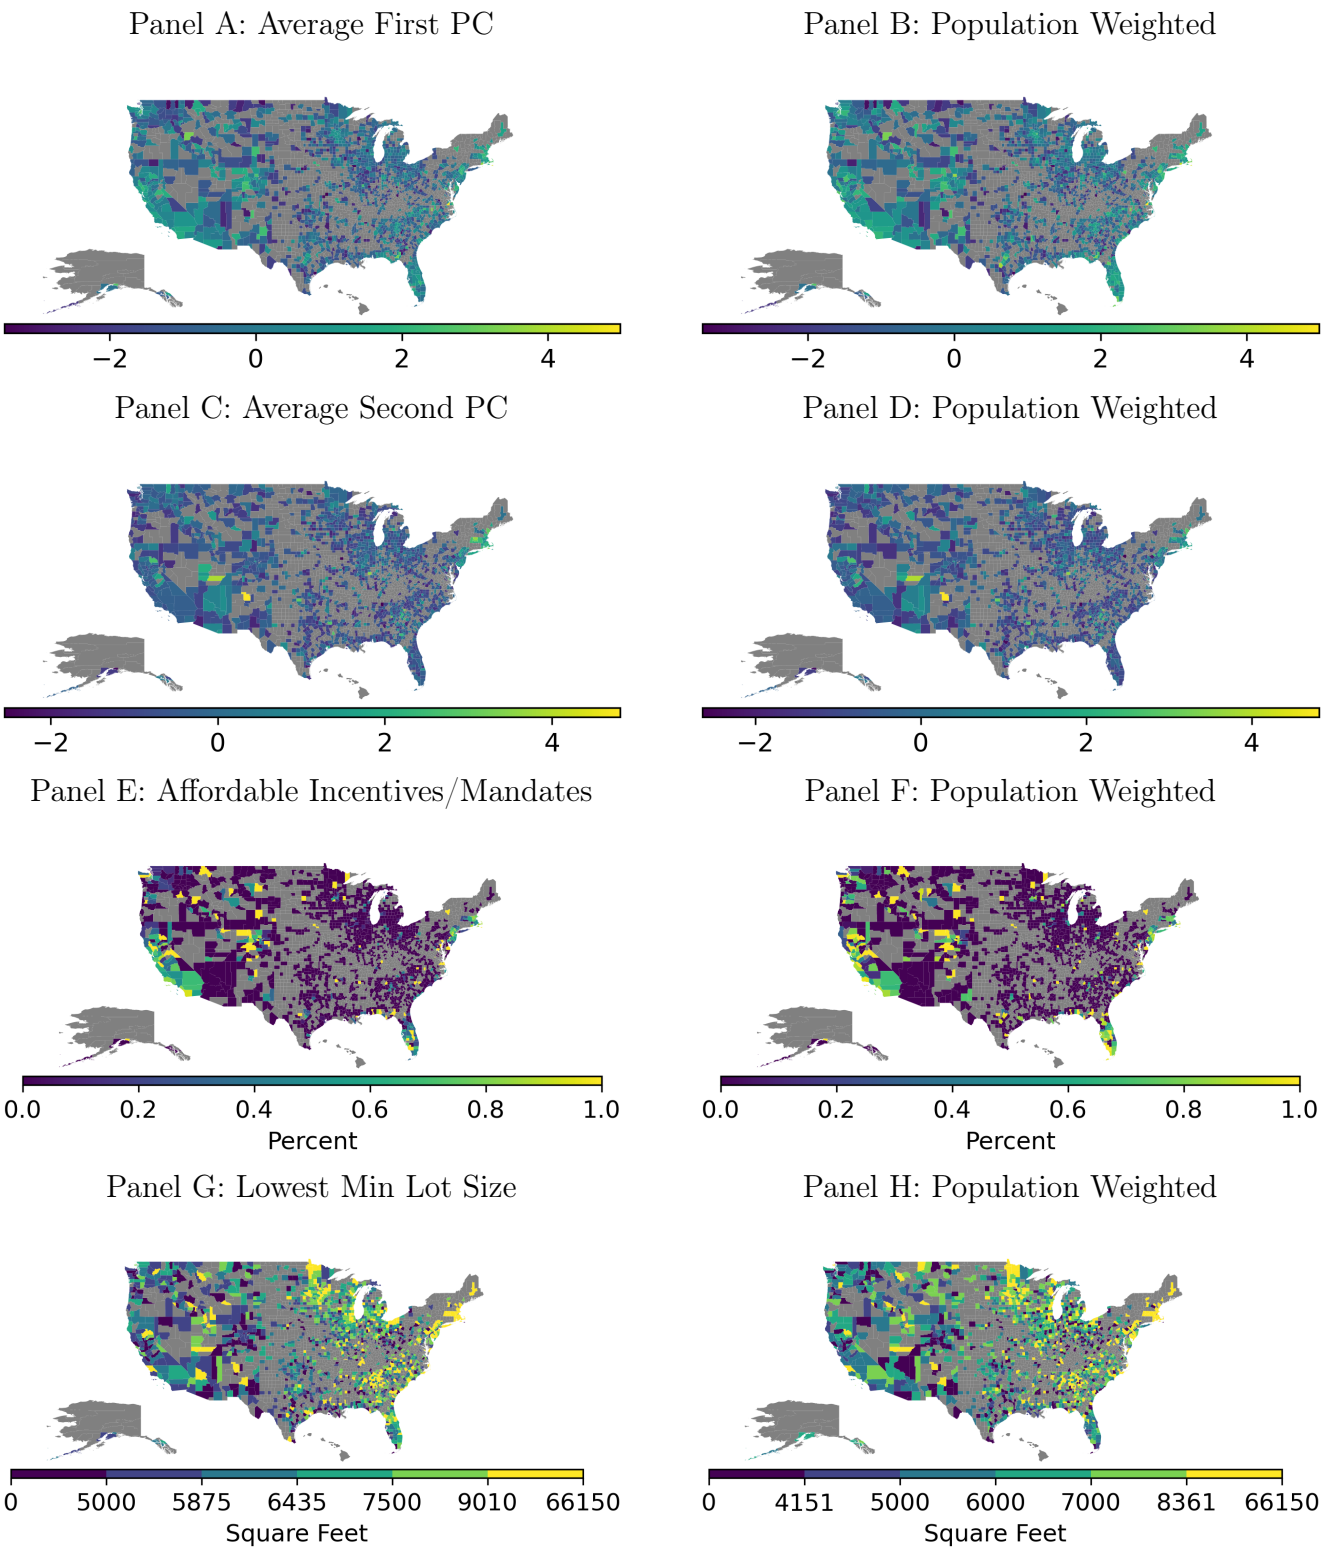


Panel D: Lowest Residential Min Lot Size



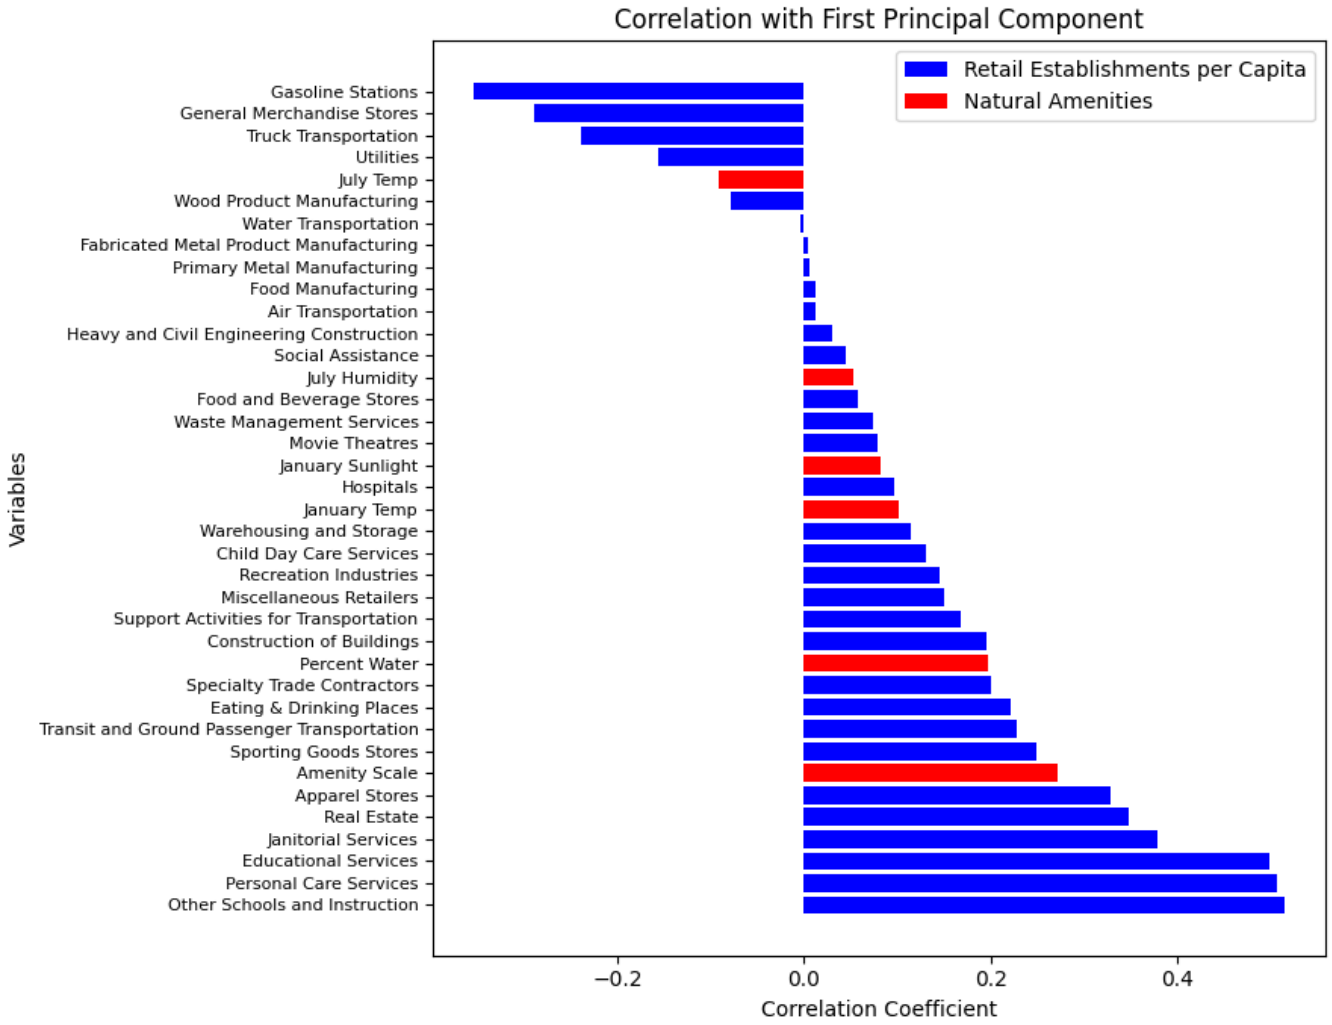
Notes: This Figure shows a comparable averages of regulatory variables across states as in Figure 8, but without population weights. This Figure therefore shows the raw average across all municipalities. Hawaii is grey because only one municipality (Honolulu) is in the dataset. For county maps please see Figure A10.

Figure A10: County Maps of Zoning Measures



Notes: This Figure shows county-level maps of four key zoning measures across the United States. Panels A–B show the first principal component (PC1) of zoning regulations, interpreted as regulatory complexity. Panels C–D display the second principal component (PC2), associated with exclusionary zoning practices. Panels E–F illustrate the prevalence of affordable housing incentives or mandates. Panels G–H depict the minimum residential lot size requirements. For each measure, we present both unweighted averages (left column) and population-weighted averages (right column). Darker colors indicate higher values or greater prevalence of each measure. We use the 2022 ACS Population as the population weight. Hawaii is grey because only one municipality (Honolulu) is in the dataset. For population weighted state maps please see Figure 8 and for unweighted state maps please see Figure A9.

Figure A11: Granular County Amenities Correlations with Principal Components



Notes: Retail establishment data are from the U.S. Census Bureau's County Business Patterns (CBP) 2022 dataset. Correlations are calculated at the county level, with the number of establishments for each industry normalized to per capita measures using county population estimates. Principal Component indices are population-weighted averages of municipality level data, aggregated to the county level. Industries are classified using 2017 NAICS codes. Natural amenity data comes from USDA county natural amenities dataset.

Figure A12: Second PC and Rent by Number of Bedrooms

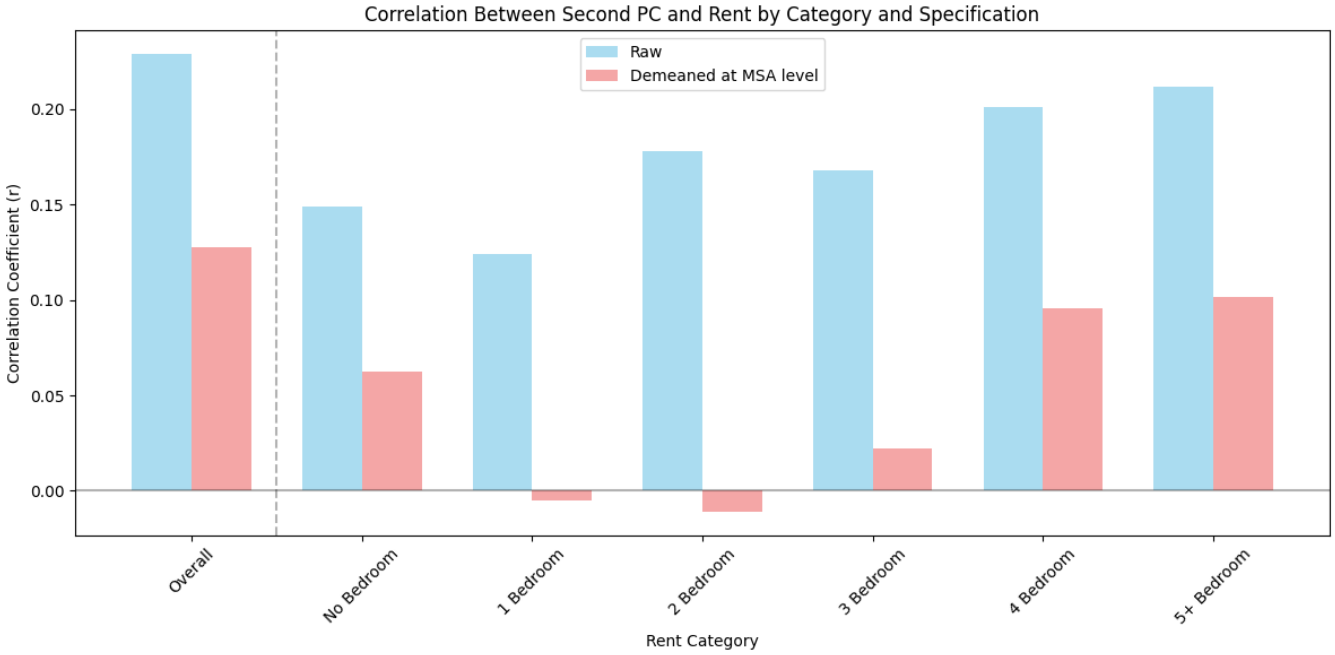
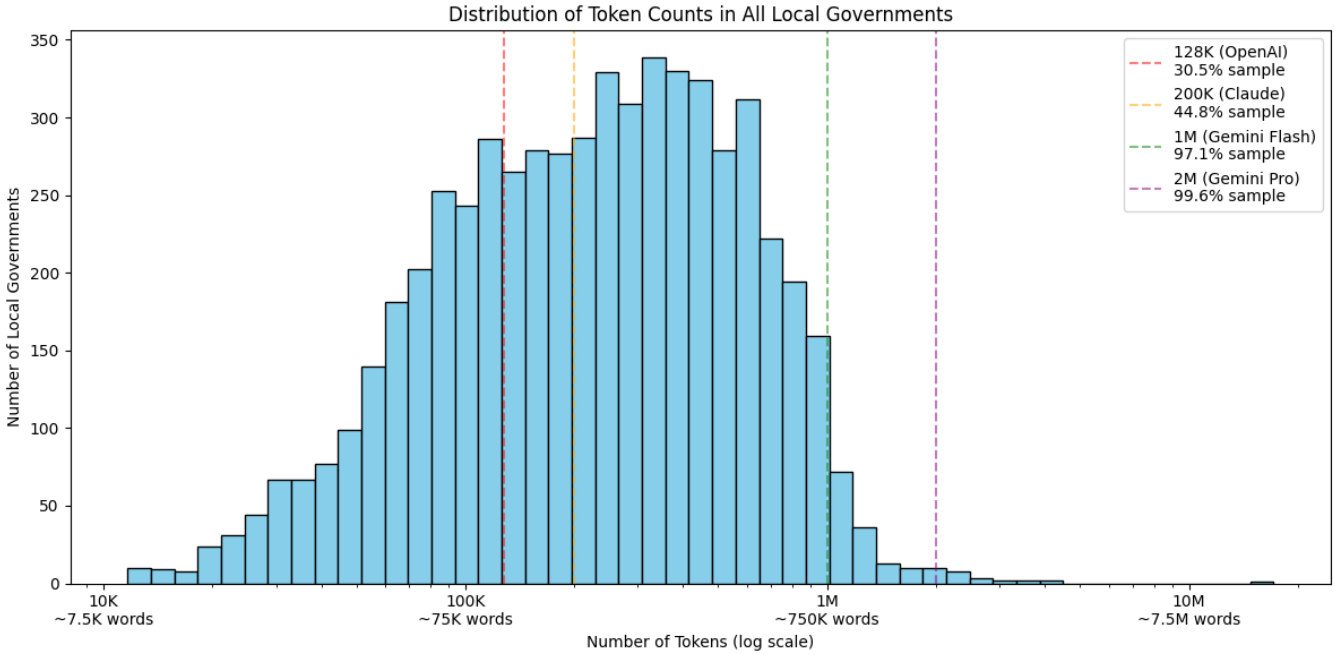


Figure A13: Distribution of Ordinance Lengths



Notes: Vertical dashed lines represent maximum context length for various popular Large Language Models in number of tokens. We use the OpenAI tokenizer used in GPT-4 Turbo (cl100k_base), though tokenizers for Claude and Gemini produce similar numbers.

Figure A14: Spatial Correlation of Zoning Regulations and Indices

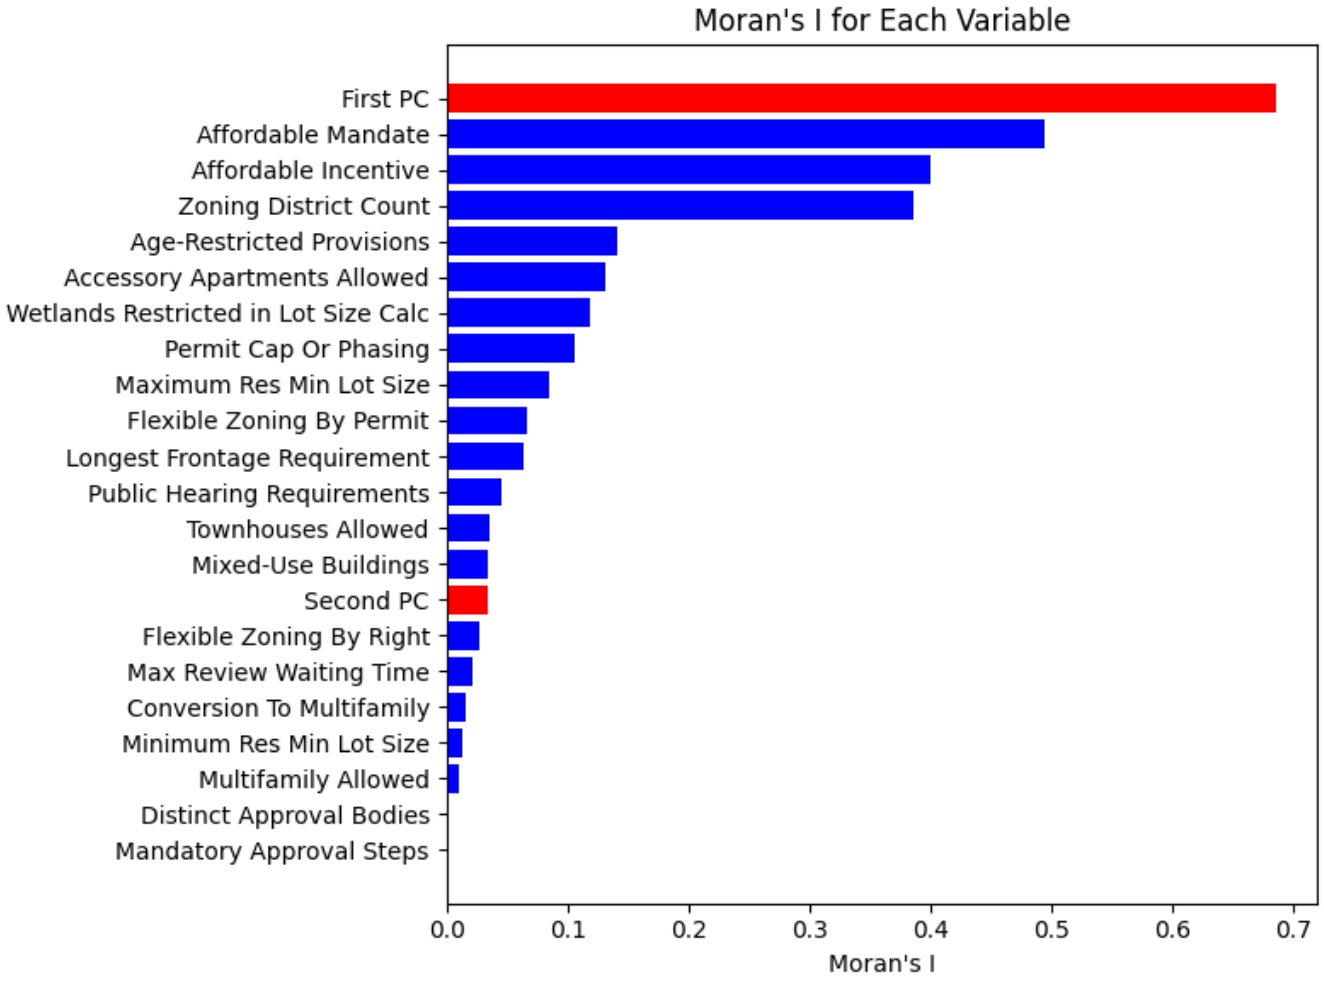
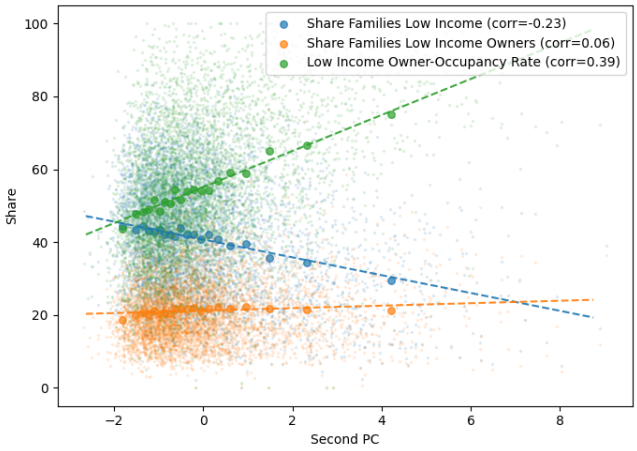
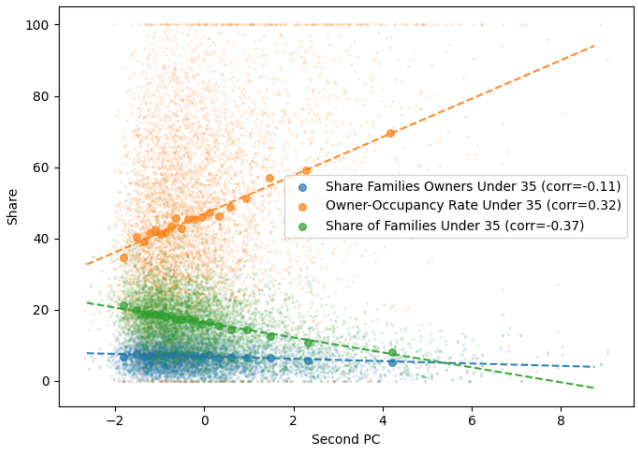


Figure A15: Correlations Between Second Principal Component and Housing Outcomes

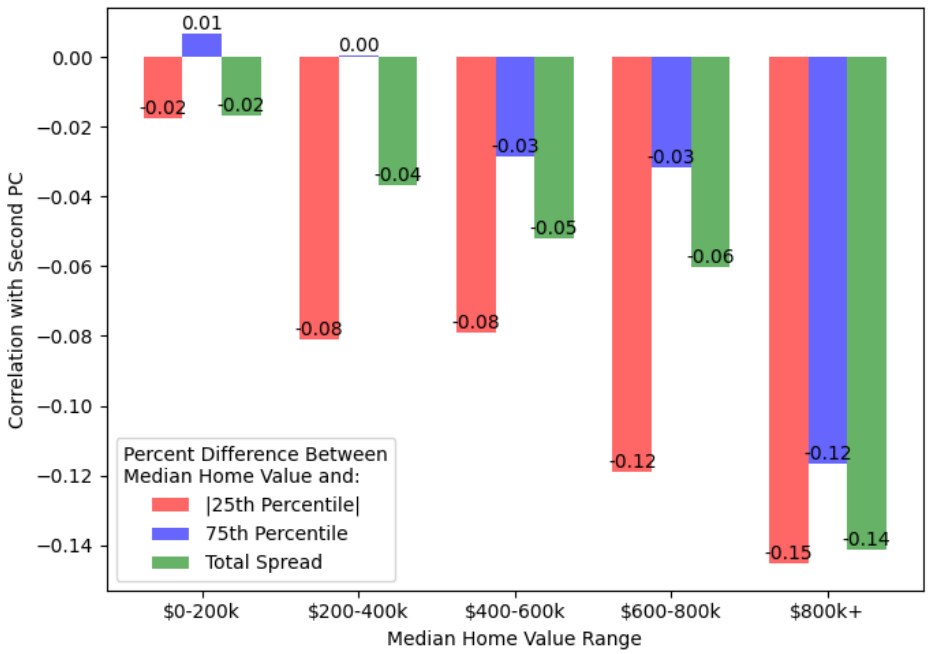
Panel A: Low-Income Housing and Ownership



Panel B: Housing and Ownership Among Younger Households



Panel C: Correlation of Second PC with Home Value Dispersion



Notes: This figure presents the relationship between the second principal component (PC) and key housing indicators. Panel A plots the relationship between the second PC and the share of low-income families, the share of low-income homeowners, and the low-income owner-occupancy rate. Panel B examines similar trends for younger households, plotting the share of family owners under 35, the owner-occupancy rate among those under 35, and the overall share of families under 35. Panel C presents correlations between the second PC and measures of home value dispersion across different median home value bins. The correlation with the 25th percentile home value reflects the absolute percent difference from the median, so a negative correlation indicates that higher-second PC municipalities have a tighter lower tail, with the 25th percentile closer to the median. Similarly, a negative correlation with the 75th percentile means that the upper tail is also compressed. The interquartile range correlation captures the combined effect of these shifts, with negative values indicating an overall tighter home value distribution in higher-second PC municipalities.

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

Full Question	Short Question
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 Incentive
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 lowest residential minimum lot size?	Lowest Res Min Lot Size
What is the mean residential minimum lot size?	Mean Res Min Lot Size
What is the highest residential minimum lot size?	Highest Res Min Lot Size
How many zoning districts, including overlays, are in the municipality?	Zoning District Count
What is the longest frontage requirement for single family residential development in any district?	Longest Frontage Requirement
Do developers have to comply with the requirement to include affordable housing, however defined, in their projects?	Affordable Mandate
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

Notes: This Table shows the mapping between the full regulatory questions and the short question name use throughout the paper. See Appendix Section A for further details on questions.

Table A2: Confusion Matrix For Binary Performance Results

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

Notes: 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. See Panel B footnote of Table 3 for details about the sample.

Table A3: Predicting Housing Market Outcomes With Zoning Regulation

Panel A: Individual Housing Regulations

Dependent Variable:	Median Home Value			Total Building Permits			Median Gross Rent			Housing Unit Density			Share Housing Units Affordable		
	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate
Lowest Res Min Lot Size	50	0.02	0.12***	32		0.01	66	0.03	0.08***	100	-0.13	-0.24***	100	-0.11	-0.17***
Multifamily Not Allowed	99	0.17	0.24***	18		0.00	100	0.14	0.16***	49	-0.04	-0.10***	73	-0.08	-0.12***
Affordable Mandate	62	0.01	0.05***	100		-0.00	79	0.06	0.09***	58	-0.01	-0.02	35	-0.01	-0.05***
Highest Res Min Lot Size	51	0.01	0.09***	11		0.02	49	0.02	0.07***	96	-0.11	-0.21***	74	-0.06	-0.15***
Affordable Incentive	62		0.01	35		0.02	61	0.03	0.08***	59		-0.02	48	-0.03	-0.07***
Townhouses Not Allowed	100	0.02	0.12***	9		-0.03**	56	-0.02	-0.01	28	-0.04	-0.07***	48		0.01
Longest Frontage Requirement	47		0.06***	37		0.01	52	0.01	0.06***	50	-0.12	-0.21***	43	-0.04	-0.12***
No Conversion to Multifamily	49		0.02*	23		0.02	64		0.02	36	-0.03	-0.06***	48		-0.03***
Age-Restricted Provisions	67		-0.07***	18		0.01	58	0.00	0.02	28	-0.00	0.00	48		-0.02*
Zoning District Count	56		-0.04***	16	0.00	0.06***	58	0.06	0.06***	40	0.03	0.03***	48	-0.04	-0.06***
Max Review Waiting Time	30		-0.00	47		0.02	42		0.01	48	-0.01	-0.02	42		-0.02*
No Mixed-Use Buildings	59		0.07***	29		-0.02	45	0.03	0.04***	33	-0.06	-0.09***	41		-0.02*
Public Hearing Requirements	57		-0.04***	12		0.04**	56	-0.00	-0.01	24	0.00	0.00	45		0.01
Wetlands Restricted in Lot Size Calc	32		0.04***	17		0.01	44	0.02	0.04***	33	-0.06	-0.11***	53	-0.05	-0.10***
Mandatory Approval Steps	43		0.03**	19		0.03*	44	0.01	0.03**	29		-0.01	41	-0.00	-0.04***
Flexible Zoning By Permit	33		0.08***	12		-0.05***	37		-0.01	49	0.09	0.08***	44	0.03	0.07***
Permit Cap Or Phasing	19		-0.04***	26		0.03**	37		0.01	33	-0.03	-0.05***	49		-0.04***
Distinct Approval Bodies	37		0.02	17		0.01	29	0.01	0.03**	36	-0.01	-0.04***	43		-0.02
Accessory Apartments Banned	22		-0.03***	21		-0.05***	31	-0.02	-0.05***	27	0.06	0.08***	23	0.08	0.12***
No Flexible Zoning By RIght	15		-0.00	13		0.00	35		-0.01	19	0.01	0.04***	38	0.02	0.06***

Panel B: Housing Regulations Indices

Dependent Variable:	Median Home Value			Total Building Permits			Median Gross Rent			Housing Unit Density			Share Housing Units Affordable		
	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate	XGBoost	LASSO	Bivariate
Overall Index	82	-0.05	-0.01	73	0.02	0.03**	91	-0.06	0.00	100		0.01	91	0.08	0.00
First PC	91		0.02	92		-0.02	96	-0.02	0.00	80		0.02	91	0.02	-0.00
Second PC	100	0.05	0.01	100		0.02	100	0.07	0.03*	76		0.01	100	-0.10	-0.03**

Notes: This specification reports the results of three sets of regressions: a random forest specification, LASSO, and bivariate regressions of our generated regulatory variables against four dependent variables. Median house value and median gross rent are drawn from the 2022 ACS. Total building permits are defined as the number of housing units permitted divided by the population of the local government, averaged over 2019–2023, from the Census Building Permits Survey. Housing unit density is the number of housing units in a local government in the 2022 ACS divided by the area from its shape file. The share of affordable housing units is defined as the percentage of housing units affordable to someone earning the state median income. Rental units are considered 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. For bivariate regressions, stars indicate statistical significance: *** p<0.01, ** p<0.05, * p<0.1. LASSO coefficients are shown where selected, with blank cells indicating variables not retained in the model. XGBoost scores represent "gain" importance from a random forests specification, measuring the average gain of splits using each feature, normalized so the most important feature for each dependent variable has a score of 100, with others scaled relatively. Higher scores indicate greater importance in the model's predictions. All variables are first demeaned at the MSA level, or for municipalities not within an MSA at the state level, then transformed into z-scores (mean=0, std=1). We use imputed regulations from our PCA analysis for LASSO and Bivariate regressions when the LLM reports "I don't know," see footnote of Appendix Table A4 for further details. We express variables so that a more positive value is associated with stricter zoning regulations, i.e. we transform the question of whether multi-family housing is allowed to whether it is not allowed. We allow missing data for XGBoost letting the algorithm both impute and predict. The Overall index in Panel B is a sum of normalized individual housing regulations for a municipality.

Table A4: Loadings on Principal Components

	First	Second	Third	Fourth	Fifth
Affordable Incentive	0.42	0.10	-0.37	-0.10	0.04
Affordable Mandate	0.32	0.11	-0.46	-0.17	0.06
Age-Restricted Provisions	0.31	0.00	-0.18	0.15	-0.02
Zoning District Count	0.30	-0.20	-0.00	0.14	-0.07
Wetlands Restricted in Lot Size Calc	0.23	0.20	0.06	0.21	0.11
Permit Cap Or Phasing	0.22	0.03	0.05	-0.22	-0.24
Highest Res Min Lot Size	0.19	0.37	0.22	0.15	-0.02
Longest Frontage Requirement	0.17	0.40	0.21	0.15	0.03
Public Hearing Requirements	0.15	0.11	-0.13	-0.29	-0.09
Variance Explained	0.13	0.11	0.06	0.06	0.05
Max Review Waiting Time	0.12	0.03	0.06	-0.39	0.41
Lowest Res Min Lot Size	0.05	0.47	0.19	0.05	-0.06
Distinct Approval Bodies	0.02	0.02	0.43	-0.31	-0.08
Mandatory Approval Steps	0.01	0.10	0.17	-0.52	0.39
No Conversion to Multifamily	-0.09	0.00	-0.19	-0.38	-0.48
No Flexible Zoning By Right	-0.13	0.04	-0.23	0.11	0.53
Multifamily Not Allowed	-0.14	0.35	-0.08	-0.10	-0.13
No Mixed-Use Buildings	-0.21	0.29	-0.20	-0.04	-0.21
Townhouses Not Allowed	-0.23	0.33	-0.14	0.04	0.03
Accessory Apartments Banned	-0.30	-0.02	-0.06	-0.10	0.06
No Flexible Zoning By Permit	-0.31	0.20	-0.31	0.05	0.10

Notes: This Table reports loadings between the first five principal components of our regulatory dataset and specific regulatory questions. We upper winsorize at the 1% level the values for highest residential min lot size, lowest residential min lot size, longest frontage requirement, and max review waiting time. We transform the highest residential min lot size variable into a dummy for whether it is above one acre. Missing data, where the model output "I don't know," were imputed with k-nearest neighbors. Prior to performing principal component analysis, all variables were normalized into z-scores. 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 A5: Question Level Accuracy Contributions From Prompting

Panel A: Binary Questions

Question	Accuracy (%)			
	No Prompting	Question Rephrasing	Background Info	Subtask Decomposition
Accessory Apartments Allowed	86.1%	99.0%	96.7%	96.7%
	(3.9)	(2.1)	(0.0)	(-)
Flexible Zoning By Right	100.0%	100.0%	100.0%	100.0%
	(0.0)	(0.0)	(0.0)	(-)
Flexible Zoning By Permit	96.0%	94.8%	95.8%	100.0%
	(0.3)	(3.4)	(2.9)	(-)
Affordable Incentive	81.4%	86.6%	99.3%	100.0%
	(2.5)	(1.9)	(1.5)	(-)
Permit Cap Or Phasing	83.3%	67.1%	91.3%	90.0%
	(0.0)	(5.2)	(1.8)	(-)
Wetlands Restricted in Lot Size Calc	76.0%	90.6%	92.7%	96.7%
	(1.5)	(5.6)	(1.5)	(-)
Multifamily Allowed	93.0%	94.7%	95.3%	100.0%
	(2.3)	(1.8)	(3.0)	(-)
Mixed-Use Buildings	79.7%	91.4%	92.6%	96.7%
	(7.5)	(2.9)	(1.5)	(-)
Conversion To Multifamily	59.5%	60.8%	88.0%	96.7%
	(1.3)	(4.1)	(3.8)	(-)
Townhouses Allowed	73.6%	68.1%	82.6%	90.0%
	(6.2)	(1.5)	(2.8)	(-)
Age-Restricted Provisions	89.5%	93.4%	98.6%	96.7%
	(5.1)	(2.9)	(1.9)	(-)

Panel B: Continuous Questions

Question	Correlation			
	No Prompting	Question Rephrasing	Background Info	Subtask Decomposition
Zoning District Count	0.88	0.95	0.94	0.98
	(0.03)	(0.03)	(0.02)	(-)
Longest Frontage Requirement	0.50	0.53	0.46	0.70
	(0.10)	(0.12)	(0.07)	(-)
Lowest Res Min Lot Size	0.80	0.87	0.79	0.92
	(0.11)	(0.04)	(0.29)	(-)

Table A6: Further Manual Validation

(a) Process Questions on Nationwide Random Sample of 30 Municipalities

Question Description	Correct	Ambiguous	Incorrect	Correct (Non-Ambiguous %)
Public Hearing Requirements	80%	10%	10%	89%
Max Review Waiting Time	87%	13%	0%	100%
Mandatory Approval Steps	93%	7%	0%	100%
Distinct Approval Bodies	80%	13%	7%	92%

(b) Bulk Questions on Random Sample of 30 California Municipalities

Question Description	Correct	Ambiguous	Incorrect	Correct (Non-Ambiguous %)
Lowest of Residential Min Lot Sizes	83%	6%	10%	89%
Longest Frontage Requirement	80%	10%	10%	89%

In Panel A we randomly draw 30 municipalities from our national dataset and have a law student review each answer. For Panel B, we draw a random sample of 30 municipalities from California and have a law student review each answer. Answers that are ambiguous include situations where the bulk regulations depend on whether a lot is interior or corner and the model chose one of the cases that could be correct. The column "Correct (Non-Ambiguous %)" calculates the percentage of correct responses among non-ambiguous cases.

Table A7: Comparative Model Performance in Predicting Housing Outcomes

		(1)	(2)	(3)	(4)	(5)	(6)	(7)
LHS Variable	Model							
Median Home Value	Lasso	1.07	1.11	0.73	1.05	0.71	0.73	0.71
	OLS	1.06	1.11	0.72	1.05	0.70	0.72	0.70
	XGBoost	1.00	1.12	0.74	0.96	0.66	0.77	0.67
Total Building Permits	Lasso	1.07	1.07	1.06	1.07	1.07	1.07	1.07
	OLS	1.06	1.06	1.06	1.06	1.06	1.06	1.06
	XGBoost	1.32	1.31	1.31	1.32	1.31	1.31	1.31
Median Gross Rent	Lasso	1.08	1.10	0.82	1.07	0.81	0.82	0.81
	OLS	1.08	1.10	0.81	1.07	0.81	0.81	0.81
	XGBoost	1.06	1.08	0.80	1.04	0.76	0.80	0.77
Housing Unit Density	Lasso	1.08	1.16	1.13	1.08	1.07	1.13	1.07
	OLS	1.08	1.16	1.13	1.08	1.07	1.13	1.07
	XGBoost	0.98	1.13	1.10	0.98	0.98	1.08	0.98
Share Housing Units Affordable	Lasso	0.99	1.02	0.70	0.97	0.69	0.70	0.69
	OLS	0.99	1.02	0.69	0.97	0.68	0.69	0.68
	XGBoost	0.99	1.03	0.63	0.95	0.53	0.62	0.52
Variables Included	Regulations	Yes	No	No	Yes	Yes	No	Yes
	Land	No	Yes	No	Yes	No	Yes	Yes
	Income	No	No	Yes	No	Yes	Yes	Yes

Notes: These specification report the RMSE of models predicting the housing market outcomes in Table A3. We include housing regulation measures in columns 1, 4, 5, and 7. We also include land availability controls (the share of land that is flat plains) in columns 2, 4, 6, and 7. We include median household income from the 2022 ACS as a control in specifications 3, 5, 6, and 7. All variables are first demeaned at the MSA level, or when not available at the state level, then transformed into z-scores (mean=0, std=1). We use imputed regulations when the LLM reports "I don't know" to avoid different imputation values across different controls, see footnote of Appendix Table A4 for further details.

Table A8: Housing Regulation and Housing Supply Elasticity

Dependent Variable:	First Principal Component		Second Principal Component	
New Housing Unit Elasticity	-0.43***	-0.12*	-0.74***	-0.03
	(0.05)	(0.07)	(0.05)	(0.17)
Share Land Developed (2001)	-0.87***	-0.17	-2.24***	-0.80**
	(0.11)	(0.20)	(0.12)	(0.36)
Squared Share Land Developed (2001)	0.38***	-0.12	1.26***	0.36*
	(0.08)	(0.15)	(0.08)	(0.20)
Share Land Flat Plains	0.01	0.02	0.19***	0.09
	(0.03)	(0.04)	(0.03)	(0.09)
Log Miles to Metro Center	-0.10***	-0.07*	-0.15***	-0.10**
	(0.02)	(0.04)	(0.02)	(0.04)
Intercept	0.14***	-0.30***	0.06***	-0.84***
	(0.02)	(0.07)	(0.02)	(0.12)
R-squared	0.06	0.21	0.17	0.41
N	3890	3890	3890	3890
MSA Fixed Effects	No	Yes	No	Yes

Notes: This specification has as the dependent variable the first regulatory principal component (value capture, first two columns) and the second regulatory principal component (exclusionary zoning, second two columns). We regress these variables against a range of variables relating to new housing production and land availability. All variables are normalized to z-scores for the regression. 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. MSA fixed effects include state fixed effects for municipalities not in an MSA. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are shown in parentheses. Standard errors are clustered at the MSA level when using MSA fixed effects and are robust otherwise.

Table A9: Housing Regulation Intensity and Distance from Metropolitan Centers

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

Notes: We perform this regression on the set of municipalities within 50 miles of the center of a metropolitan area which accounts for 3,605 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. Asterisks denote significance levels: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Standard errors are shown in parentheses. We cluster standard errors at the metro level.

Table A10: Description of Variables Used in the Study

Variable	Source	Definition
Auto Commute Share	2022 American Community Survey	Percentage of commuters using either cars, trucks, or vans as their primary commute method.
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).

Continued on next page

Table A10 continued from previous page

Variable	Source	Definition
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).
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.

Continued on next page

Table A10 continued from previous page

Variable	Source	Definition
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).

Continued on next page

Table A10 continued from previous page

Variable	Source	Definition
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).
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. This excludes property taxes raised from independent school districts.
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.
Local Revenue Per Student	2022 Annual Survey of School System Finances	The local revenue for a school district divided by the enrollment of that school district. We spatially merge school districts into the Census of Governments. Local revenue for school districts includes property taxes directly raised by the school and transfers from local governments for subordinate school districts.

Continued on next page

Table A10 continued from previous page

Variable	Source	Definition
Building Permits All Units	Building Permits Survey	The number of housing units permitted divided by the population of the local government averaged over 2019–2023.
Year of Incorporation	Goodman, C. B. (2023). Municipal Incorporation Data, 1789-2020 (Version 1.0.0) [Computer software]. https://github.com/cbgoodman/muni-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	Baum-Snow and Han (2024)	The coefficient on the change in new housing units to the change in housing prices, specifically gamma01b_newunits_FMM.
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 B: LLM Replicability

LLMs responses are not fully deterministic and so the results of this study cannot be completely replicated for several reasons. First, LLMs sample tokens from a probability distribution leading to variation in responses to the same query across iterations. Second, the GPT-4 class of models that we use in this study follow a Mixture of Exports (MoE) architecture (see here and here for further details). This means that the specific expert that an LLM query gets routed to varies across API calls depending on supply/demand of experts. In turn, this implies that the underlying probability distribution that tokens are sampled from may change from one query call to the next, depending on the availability of experts. Moreover, OpenAI's models are closed source so there may exist other sources of randomness across API calls that we cannot explain.

In this section we quantify how deterministic LLM queries are in our use case and how ensembling many query calls may help mitigate the issue. In general, we query the LLM to respond with a detailed answer, (i.e. "Think step by step") followed by a structured output (i.e. "Yes" or "No"). Randomness in output for LLMs leads to a high frequency of variation of in open-ended responses, but many of these differences do not change the overall meaning of a response, for example just swapping synonyms. However, we do find some variation in the structured output of LLM responses in our use case.

One potential way to mitigate non-determinism is to request multiple responses from the LLM and then aggregate the answers by majority rule, an ensemble approach. Previous research has also found that ensemble methods can greatly improve LLM performance (Li et al., 2024).

With the OpenAI API there are two ways to ensemble API calls. The first way is to request multiple chat completions for a given query (by setting the API parameter $n > 1$). This effectively samples the distribution of tokens several times. This method is also cost effective because OpenAI only charges the user once for the input tokens regardless of how many iterations of output tokens are requested. However, this approach fails to sample from the distribution of potential experts or other potential sources of variation, for example the hardware of the server in which the LLM was run. A more costly approach is to separately query the LLM for each of the ensemble queries paying for both the input and output tokens used in each call. This second approach more broadly

samples from the various sources of randomness for an LLM response.

We measure replicability with two measures in this analysis. Both measures average pairwise matching rates. We compare the final structured answer from a given LLM query across multiple API calls for all pairs (n choose 2), and take the ratio of the number of pairs that match to the total number of pairs. We call this ratio a consistency score. We measure both internal consistency, scores from requesting multiple chat completions for a given query, and external consistency, scores from comparing separate API calls.

We confirm that lower temperatures create more deterministic responses, even after ensembling. In Table A11, we compare temperatures of 0, 0.5, and 1 (lower temperatures should mean more deterministic responses) as well as whether including a random seed makes responses more deterministic. We use a random sample of 30 municipalities from our national sample and use two questions, a binary one (whether there are permits caps or project phasing) and a continuous one (how many districts there are). For each specification we run the model five times, each time requesting 10 chat completions. We measure internal consistency scores within a model run, and external consistency across aggregated majority rule answers from each model run. We find that lower temperature models are more internally consistent, especially for the continuous question, and have a lower variance of internal consistency. After aggregating responses, we find that external consistency scores are fairly similar for the binary question, but still higher for the continuous one and with lower variance. We also do not find evidence that including a random seed makes responses more deterministic.

We next show in Figure A16 that external consistency grows with ensemble size, especially for the continuous question. This suggests that answers begin to stabilize at larger ensemble sizes, though not fully. We still find that at least five percent of pairwise comparisons do not match even with ensemble sizes of 10 and a temperature of 0. We also find that the zero temperature specification already begins at a fairly high level of external consistency even with an ensemble size of 1.

Next, we ask how informative the internal consistency score is for predicting the external consistency score. In Figure A17 we find that higher temperature model internal consistency scores are highly predictive of external consistency and that this effect grows with ensemble size. However,

lower temperature models internal consistency scores are not very informative for external consistency. This suggests that within model run variation for low temperature models may not reflect the same source of randomness as across model runs. For example, the variation in which expert the query is routed to may be more important than the within expert sampling distribution for predicting external consistency.

We next explore whether performance increases with ensemble size. For this analysis we return to our testing sample of 30 Massachusetts municipalities where we have a clean dataset to compare answers to. We use temperatures of 0 and 1 and request 10 chat completions from each model run. We do not find evidence of increased performance with larger ensemble sizes. In Figure A18 we find a fairly persistent outperformance of the 0 temperature model for the binary question and of the 1 temperature model for the continuous question.

We next ask whether internal consistency scores are helpful for predicting accuracy of questions. If the internal consistency score is highly predictive of the external consistency score then the measure can be used as a model confidence measure. We find in Table A12 that the internal consistency score is somewhat helpful for predicting accuracy with the temperature 1 model but not helpful for the temperature 0 model. With a larger ensemble size of 100, (Wei et al., 2024) found that answer frequency within an ensemble was highly predictive of accuracy, especially for more advanced reasoning models.

We suggest researchers consider one of two specifications. First, a low temperature single shot approach. This approach is cost effective by only requesting one chat completion, has a high degree of external consistency, and is straightforward to explain. If a researcher wants to use an ensemble approach then we suggest using a high temperature model with an ensemble size of at least five. High temperature ensemble models have a high degree of external consistency and have informative internal consistency scores for both predicting external consistency and accuracy. We choose to use the first approach in this paper for the higher external consistency, cost savings, and for simplicity.

Table A11: Internal and External Consistency Varying Temperature and Seed

Question	Temperature	Seed	External Consistency Mean	External Consistency Variance	Internal Consistency Mean	Internal Consistency Variance
Continuous: Number of Districts	0	No	0.91	0.05	0.9	0.05
		Yes	0.89	0.04	0.9	0.04
	0.5	No	0.83	0.09	0.68	0.1
		Yes	0.8	0.11	0.67	0.1
	1	No	0.75	0.11	0.61	0.1
		Yes	0.78	0.11	0.64	0.1
Binary: Whether Permit Caps or Phasing	0	No	0.93	0.03	0.96	0.02
		Yes	0.92	0.03	0.95	0.02
	0.5	No	0.93	0.03	0.88	0.04
		Yes	0.9	0.05	0.89	0.04
	1	No	0.91	0.04	0.88	0.04
		Yes	0.92	0.04	0.86	0.04

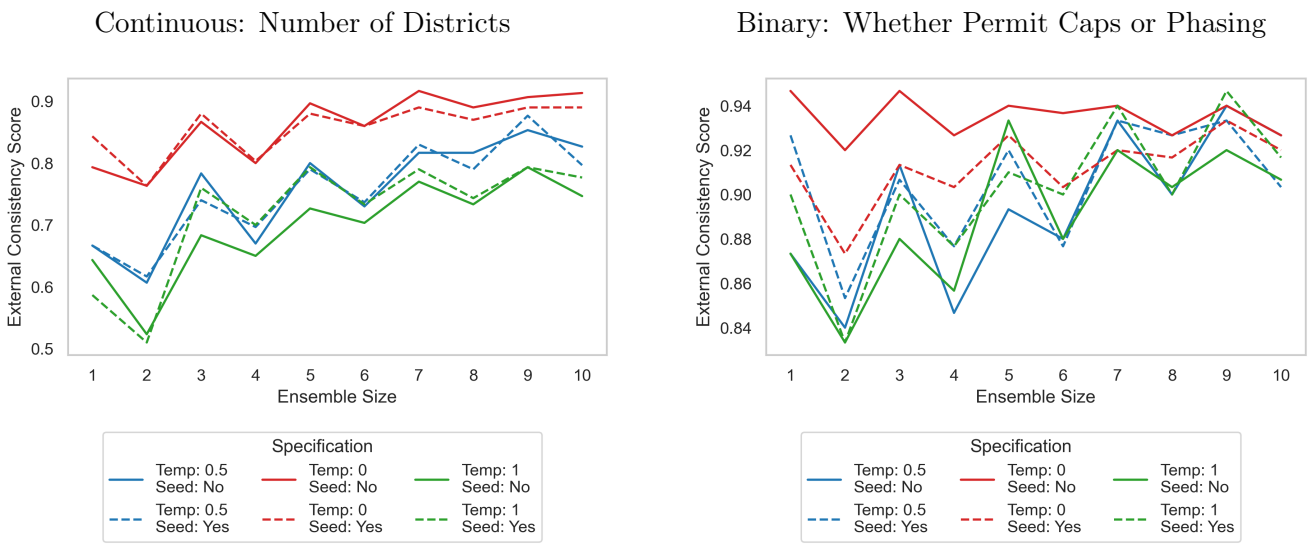
Notes: This Table compares the internal and external consistency of LLM responses across different temperature settings and with/without a random seed. "Temperature" refers to the randomness in the LLM's output (0 being most deterministic, 1 being most random). "Seed" indicates whether a random seed was used for reproducibility. "External Consistency" measures agreement across separate API calls, while "Internal Consistency" measures agreement within a single API call requesting multiple completions. Results are shown for two types of questions: a continuous question about the number of zoning districts, and a binary question about permit caps or phasing. Mean values closer to 1 indicate higher consistency. Lower variance indicates more stable results across trials.

Table A12: Regression of Internal Consistency Score on Whether Correct

	Temp = 0		Temp = 1	
	(1)	(2)	(3)	(4)
Consistency Score		0.2603		0.4175*
		(0.3614)		(0.2097)
Question FE	Yes	Yes	Yes	Yes
R-squared	0.1383	0.1461	0.2188	0.2696
N	60	60	60	60

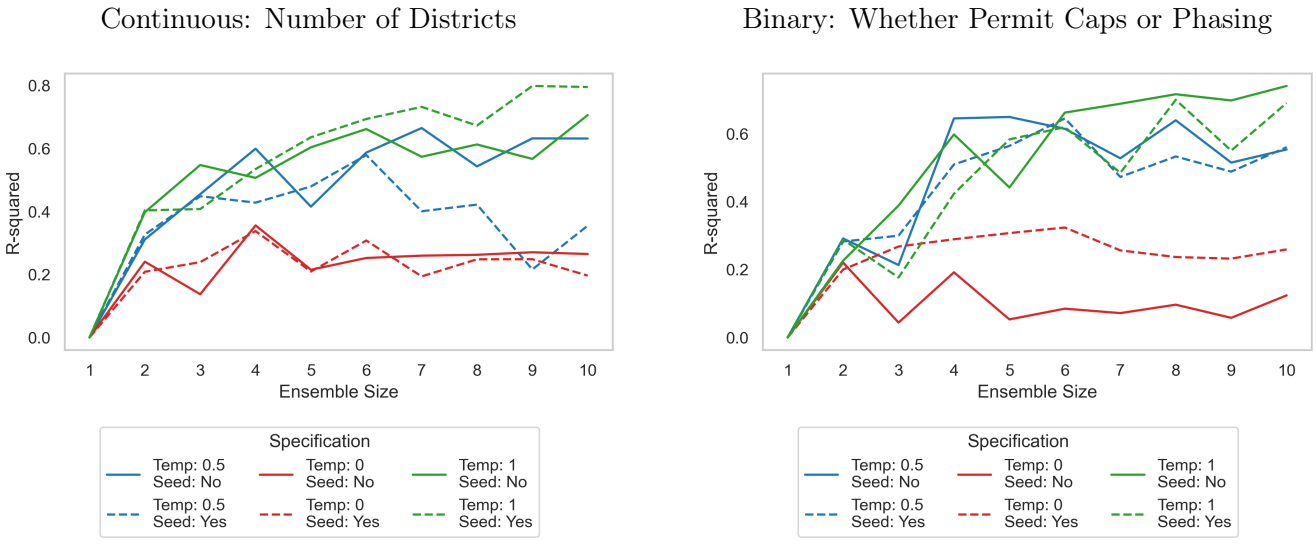
Notes: This table presents regression results examining the relationship between the internal consistency score of LLM responses and their accuracy. The dependent variable is a binary indicator for whether the LLM's response is correct. Results are shown for two temperature settings: 0 (most deterministic) and 1 (most random). Columns 1 and 3 include only question fixed effects, while columns 2 and 4 add the consistency score as an explanatory variable. Standard errors are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure A16: External Consistency vs. Ensemble Size



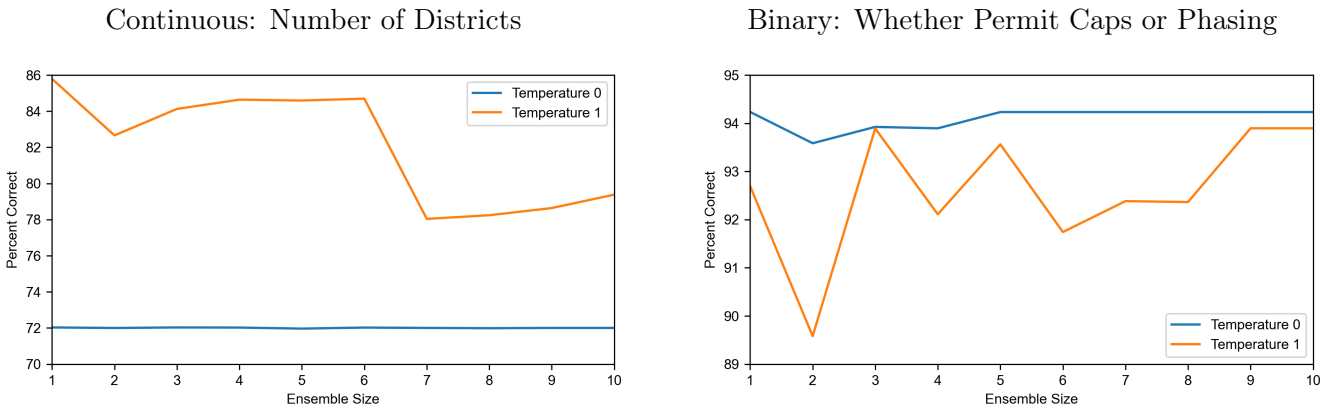
Notes: This Figure shows how external consistency of LLM responses varies with ensemble size for two types of questions: a continuous question about the number of zoning districts (left) and a binary question about permit caps or phasing (right). The x-axis represents the ensemble size (number of model runs aggregated), while the y-axis shows the external consistency score. Different lines represent various temperature settings (0, 0.5, 1) and whether a random seed was used.

Figure A17: Relationship Between Internal and External Consistency vs. Ensemble Size



Notes: This figure demonstrates the relationship between internal and external consistency of LLM responses as a function of ensemble size for two types of questions: a continuous question about the number of zoning districts (left) and a binary question about permit caps or phasing (right). The x-axis represents the ensemble size, while the y-axis shows the R-squared value, indicating how well internal consistency predicts external consistency. Different lines represent various temperature settings (0, 0.5, 1) and whether a random seed was used.

Figure A18: Ensemble Size vs.Correct



Notes: This Figure illustrates the relationship between ensemble size and accuracy of LLM responses for two types of questions: a continuous question about the number of zoning districts (left) and a binary question about permit caps or phasing (right). The x-axis represents the ensemble size (number of model runs aggregated), while the y-axis shows the percent of correct responses. Two temperature settings are compared: 0 (blue line, most deterministic) and 1 (orange line, most random). For the continuous question, accuracy is measured as the percentage of responses within a certain tolerance of the true value, while for the binary question, we use the percentage of correct classifications.

Appendix C: 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 explictly 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 explictly 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 explictly by right in any district? Question 2: Is open space residential design allowed explictly by right in any district? Question 3: Is any type of flexible zoning other than cluster development and open space residential design allowed explictly 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 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 explictly 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 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 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 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 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 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?

Appendix C: Unincorporated Area Construction

Making an Incorporated Area Shapefile

We use the Census of Governments (CoG), specifically the Government Units Survey (GUS) subcomponent as an index for all local governments in the U.S. Local governments are sub-county governments and are categorized into either municipalities or townships. Municipalities correspond with Federal Information Processing Standards (FIPS) place codes and exist in every state. Townships correspond with FIPS county subdivision codes and exist only in the Northeast and Midwest regions of the US. To measure the geographic area of the US covered by local governments we merge all municipalities and townships with their corresponding census shapefiles by merging municipalities with FIPS place shape files and townships with FIPS county subdivision shape files. We then take the spatial union of these shape files as the incorporated area of the US. We call this the incorporated area shapefile.

Incorporated Area in 2000

We use the 2002 GUS subcomponent of the CoG as the index for local governments. We then merge these local governments with the 2000 census TIGER/line shape files. We use the 2000 shape files because they are available in the modern Shapefile format, while the 2002 TIGER/Line files are only available in ASCII formats. In practice, this means that we drop a small number of local governments that were incorporated in 2001/2002 and thus did not yet have corresponding shape files. We also miss a small number of local governments that disbanded in the two years.

Incorporated Area in 2022

We use the 2022 GUS subcomponent of the COG for the index for local governments. We use the 2022 Census shapefiles

Measuring Share Incorporated For Counties

For each county we first isolate the area that is land by taking the spatial difference of each county's shapefile with its corresponding water area shapefiles. We do this because many counties

do not incorporate water areas in their domain and this area is not relevant for housing outcomes. Then, for each county we intersect its land area with the incorporated area shapefile to measure the land area of the county that is incorporated. We then take the ratio of the land area that is incorporated to the total land area of the county as the share incorporated of the county. To measure MSA incorporation share we take the land area weighted average of all counties within an MSA.

Measuring Share Incorporated For Granular Census Geographies

For both census blocks and block groups we use the same methodology to determine the share incorporated. Specifically, we intersect the corresponding shapefile for each block/block group with the incorporated area shapefile. We take the ratio of the area incorporated within the block/block group with total area of the block/block group as the share incorporated. Note that we do not exclude water area for blocks and block groups since they do not tend to include much water area.

Using Share Incorporated For Granular Census Geographies

Most blocks and block groups are either entirely inside of an incorporated area or outside of all incorporated areas, while many are partially located inside of incorporated areas. For partially incorporated blocks/block groups, we make the assumption that variables (i.e., population) are evenly dispersed across space. In other words, we assume that there are no differences in the incorporated and unincorporated portion of a partially incorporated block/block group. As a result, we proceed by dividing up partially incorporated blocks/block groups into sub-blocks and sub-block groups that are entirely incorporated and entirely unincorporated with weights of the percent incorporated. For example, if a block has a population of 1,000 and is 25% incorporated then we say that 250 people in that block live in the incorporated area and 750 people live in the unincorporated area.

Appendix D: Comparison With Wharton Study

The Wharton Residential Land Use Regulation Index (WRLURI) is a survey-based measure of local housing market regulations. In this section we compare our study to the 2018 version of the index (Gyourko et al., 2021). The survey was administered through the International City Managers Association (ICMA) to its 10,949 member municipalities at the time, receiving 2,825 responses for a response rate of 25.8%. In Table A13, we compare our dataset's sample coverage with that of the Wharton index. Our analysis reveals superior coverage across the distribution of various demographic and housing characteristic variables. We especially have greater coverage for low population, low white share, high college share, and high median home value local governments.

In Panel A of Table A14, we compare question level responses focusing on two areas where our question bases overlap: affordable housing mandates and minimum lot sizes. Our measure of affordable housing shows a substantial correlation of 0.38 with the Wharton measure. For minimum lot size requirements, we find smaller but meaningful correlations ranging from 0.18 to 0.37 across different size categories. We next compare the accuracy of each study for the highest minimum lot size by benchmarking responses against independently sourced geographic data from Massachusetts GIS in Panel B. Our approach achieves an average accuracy rate of 77%, substantially higher than Wharton's 50% accuracy rate when compared against the same ground-truth data.

In Table A15, we examine correlations between our principal components and the Wharton index. We find positive but somewhat low correlations ranging from 0.11 to 0.33. Besides differences in question-level accuracy, our study exclusively measures regulation, while WRLURI takes a broader approach. Of its 71 questions, only 56.3% directly measure zoning regulations. The remaining questions cover important aspects of the housing production process, such as approval processes and timelines in practice (22.5%), but also questions less relevant to a strictly regulatory index, like housing market outcomes and market conditions (i.e., whether single-family housing supply meets demand). This broader scope is reflected in the PCA weights in the Wharton index. The highest weights are assigned to process-based measures, such as court involvement, state

Table A13: Sample Coverage Comparison With Wharton Index

	Variable Distribution Percentiles					
Variable	Bottom 5%	5-25th	25-50th	50-75th	75-95th	Top 5%
Population						
Wharton Sample (%)	0.0	0.0	0.0	5.1	25.1	31.6
Our Sample (%)	0.3	0.7	3.8	13.7	40.5	65.3
Difference (p.p.)	0.3	0.7	3.8	8.7	15.4	33.7
Owner-Occupied Share						
Wharton Sample (%)	12.2	13.8	8.8	6.1	3.2	0.1
Our Sample (%)	30.9	29.2	17.7	10.8	5.9	0.2
Difference (p.p.)	18.8	15.5	8.9	4.7	2.7	0.1
65+ Population Share						
Wharton Sample (%)	2.7	9.8	10.5	8.6	4.8	1.4
Our Sample (%)	7.7	21.8	20.5	15.3	10.1	6.0
Difference (p.p.)	5.0	11.9	10.0	6.7	5.3	4.6
Under 18 Share						
Wharton Sample (%)	0.4	5.8	10.3	10.5	7.4	1.2
Our Sample (%)	2.6	12.1	19.2	21.3	15.4	4.5
Difference (p.p.)	2.2	6.3	9.0	10.9	8.0	3.3
White Share						
Wharton Sample (%)	11.2	15.1	9.8	6.0	1.5	0.0
Our Sample (%)	34.5	35.8	18.7	8.0	1.7	0.4
Difference (p.p.)	23.3	20.7	8.9	2.1	0.2	0.3
College Share						
Wharton Sample (%)	0.2	2.3	6.0	8.9	14.3	16.9
Our Sample (%)	1.4	6.9	11.6	15.6	27.0	46.8
Difference (p.p.)	1.1	4.6	5.6	6.8	12.7	30.0
Median Home Value						
Wharton Sample (%)	0.5	3.0	6.5	9.1	13.8	23.3
Our Sample (%)	1.7	6.7	12.1	14.1	27.9	71.4
Difference (p.p.)	1.1	3.7	5.5	5.1	14.1	48.1

political involvement, and local political pressure indices, while direct regulatory measures such as supply restrictions and density restrictions indices receive the smallest loadings[12].

[12]The factor loadings from the principal component analysis used to construct the WRLURI2018 are: Court Involvement Index (0.42), State Political Involvement Index (0.41), Local Political Pressure Index (0.40), Environmental Index (0.28), Approval Delay Index (0.28), Local Project Approval Index (0.29), Local Zoning Approval Index (0.27), Open Space Index (0.24), Affordable Housing Index (0.27), Local Assembly Index (0.17), Supply Restrictions Index (0.12), and Density Restrictions Index (0.09).

Table A14: Correlation With Wharton Index

Panel A: Averages and Correlation For Wharton Questions

Question		Wharton Average	Our Average	Correlation
Affordable Housing		0.20	0.06	0.38
Minimum Lot Size	Less than 1/2 acre	0.50	0.49	0.37
	1/2 to 1 acre	0.17	0.13	0.18
	1 to under 2 acres	0.12	0.17	0.26
	2 acres or more	0.22	0.16	0.27

Panel B: Comparison to Massachusetts GIS Minimum Lot Sizes

	Massachusetts GIS Data	Less than 1/2 acre	1/2 to 1 acre	1 to under 2 acres	2 acres or more	Average Accuracy
Our Model	Less than 1/2 acre	6	2	0	1	
	1/2 to 1 acre	1	12	1	1	
	1 to under 2 acres	0	1	23	2	
	2 acres or more	1	0	4	8	
	Percent Correct	75%	80%	82%	67%	77%
Wharton	Less than 1/2 acre	3	3	1	0	
	1/2 to 1 acre	4	6	7	0	
	1 to under 2 acres	0	4	13	2	
	2 acres or more	1	2	7	10	
	Percent Correct	38%	40%	46%	83%	50%

Notes: The sample overlap between this study and Gyourko et al. (2021) is 1,171 municipalities. We drop municipalities that do not have any minimum lot size requirements. The Affordable Housing questions refers only to affordable housing mandates, not incentives, and the minimum lot size questions refers only to residential districts. Massachusetts GIS data comes from MassGIS (Bureau of Geographic Information). To find the highest residential min lot size from MassGIS we first drop districts with missing/no min lot size info and then drop districts not flagged as single family zoned.

Table A15: Index Level Wharton Correlations

	Wharton Index	PC 1	PC 2	Overall Index
Wharton Index	1.00	0.33	0.11	0.22
PC 1	0.33	1.00	0.07	0.36
PC 2	0.11	0.07	1.00	0.74
Overall Index	0.22	0.36	0.74	1.00

The Overall Index sums the z-scores of all zoning questions.

Appendix E: Economic Model of Housing Regulation and Municipal Zoning

This appendix provides a detailed technical exposition of a model used to analyze municipal zoning regulations, household location choices, and government policy. The model combines a

spatial equilibrium framework with non-cooperative local government optimization to explain the key facts around housing production, choice of regulations, and socio-economic sorting which we observe in our key facts.

A1 Model Setup

Agents and Locations: To capture the key socio-economic differences among agents, we define households as of two types: low productivity (L) and high productivity (H), earning wages w^L and w^H, respectively). These agents decide to live in one of two zones, reflecting the key spatial differences within urban areas as either a city core c and a suburb s, which are differentiated by distinct amenities (α_c, α_s) and regulatory environments. $\eta_i = \frac{N_i^H}{N_i}$ is therefore the share of high productivity workers in each zone.

Government Structure: Local governments within each zone $i \in \{c, s\}$ levy a homogeneous wage tax τ and differentiated housing tax t_i. Governments produce public services s_i using labor G_i, paid for by low productivity workers, and maximize net revenue.

Housing Market: The housing market clears through spatial sorting of workers. Housing rents r_i in each region are upward sloping and depend on local population density:

$$r_i = \bar{r}_i + \beta_i \cdot \log(N_i^L \cdot h_i^L + N_i^H \cdot h_i^H)$$

where $\bar{r}_i$ is the baseline rent fixed in each region, $\beta_i > 0$ governs the supply elasticiy, and N_i^J is the population of type j in zone i.

A2 Household Problem

Households maximize utility by choosing consumption (c), housing (h), and location (choice of i). The utility function for a household of type j in zone i is:

$$U_i^j = \max_{c,h} \gamma \log c + (1-\gamma) \log h + \alpha_i - \log(N_i) + \log(s_i)$$

subject to the budget constraint:

$$c + h \cdot r_i = (1 - \gamma) \cdot w^j.$$

Utility in this framework is decreasing in local population (N_i) due to congestion, and increasing in public services s_i. Housing demand in this framework therefore satisfies:

$$h_i^= \frac{(1-\gamma)(1-\tau)w_j}{r_i}$$

A3 Government Problem

Local governments maximize net revenue:

$$\max_{G_i, t_i, \underline{h}_i} \; t_i \cdot r_i \cdot N_i + \tau \cdot (w^L \cdot N_i^L + w^H \cdot N_i^H) - w^L \cdot G_i,$$

subject to a production function on public services:

$$s_i = \gamma_i G_i,$$

a requirement on minimum housing consumption, which corresponds to exclusionary zoning:

$$h_i^j \geq \underline{h}_i,$$

and a population constraint:

$$N_i^L + N_i^H = N_i$$

The two key government regulatory policies are the choice of housing taxes t_i, which corresponds to value capture or the first principal component, and $\underline{h_i}$, a minimum housing requirement, which corresponds to exclusionary zoning and the second principal component. Governments in both zones choose regulations in a non-cooperative way independent of the other jurisdiction.

A4 Equilibrium and Sorting

The key condition for market clearing is that total population satisfies

$$N_c + N_s = N$$

for each household type.

In spatial equilibrium (Rosen, Roback), households sort across zones until they are indifferent, resulting in the following condition:

$$\alpha_c - \alpha_s + \log \frac{s_c}{s_s} = (1-\gamma)\left[\log \frac{r_c}{rs} + \log \frac{1+t_c}{1+t_s}\right] + \log \frac{N_c}{N_s}.$$

In other words, the relative benefit of being in one zone instead of the other (based on local baseline amenity and government services) is equal to the additional costs (taxes and congestion costs).

A5 Optimal Policy Conditions

The key force behind optimal policy is the share of local workers of a given type, η_i.

Exclusionary Zoning ($\underline{h_i}$) This regulatory choice is binding when raising $\underline{h_i}$) raises revenue via the sorting of high productivity households:

$$\frac{\partial \eta_i}{\partial \underline{h_i}} N_i \left[t_i \frac{\partial r_i}{\partial \eta_i} + \tau \frac{\partial \tilde{w}}{\partial \eta_i}\right] + \frac{\partial N_i}{\partial \underline{h_i}} \left[t_i r_i + \tau \tilde{w}_i + N_i t_i \frac{\partial r_i}{\partial N_i}\right] \geq 0.$$

Value Capture (t_i This regulatory choice is optimal when increased tax revenue offsets population loss:

$$r_i N_i + \frac{\partial N_i}{\partial t_i}\left[t_i r_i + N_i t_i \frac{\partial r_i}{\partial N_i} + \tau \tilde{w}_i\right] \geq 0.$$

A6 Key Model Predictions

In spatial equilibrium, the model features the following spatial relationships:

Theorem 1. *Suburban Exclusion: Suburbs s adopt stricter bulk regulations* $\underline{h}_s$ *to attract high income households.*

Theorem 2. *Urban Value Capture: City cores c prioritize revenue extraction* (t_c) *due to higher demand and amenities.*

Theorem 3. *Income sorting: High income households cluster in suburbs, while low income households concentrate in city cores.*

The model also features the following comparative statics on regulation:

Theorem 4. *PC2 increases with distance from city center*

Theorem 5. *PC1 decreases with distance from city center*

Theorem 6. *PC1 correlates with house prices*