INTRODUCTION TO THE TERNER HOUSING POLICY SIMULATOR VISUALIZATION:

METHODS, USES, AND LIMITATIONS

Author: Stephanie Hawke, Associate Research Director of Land Use and Supply, Terner Center for Housing Innovation

May 2026

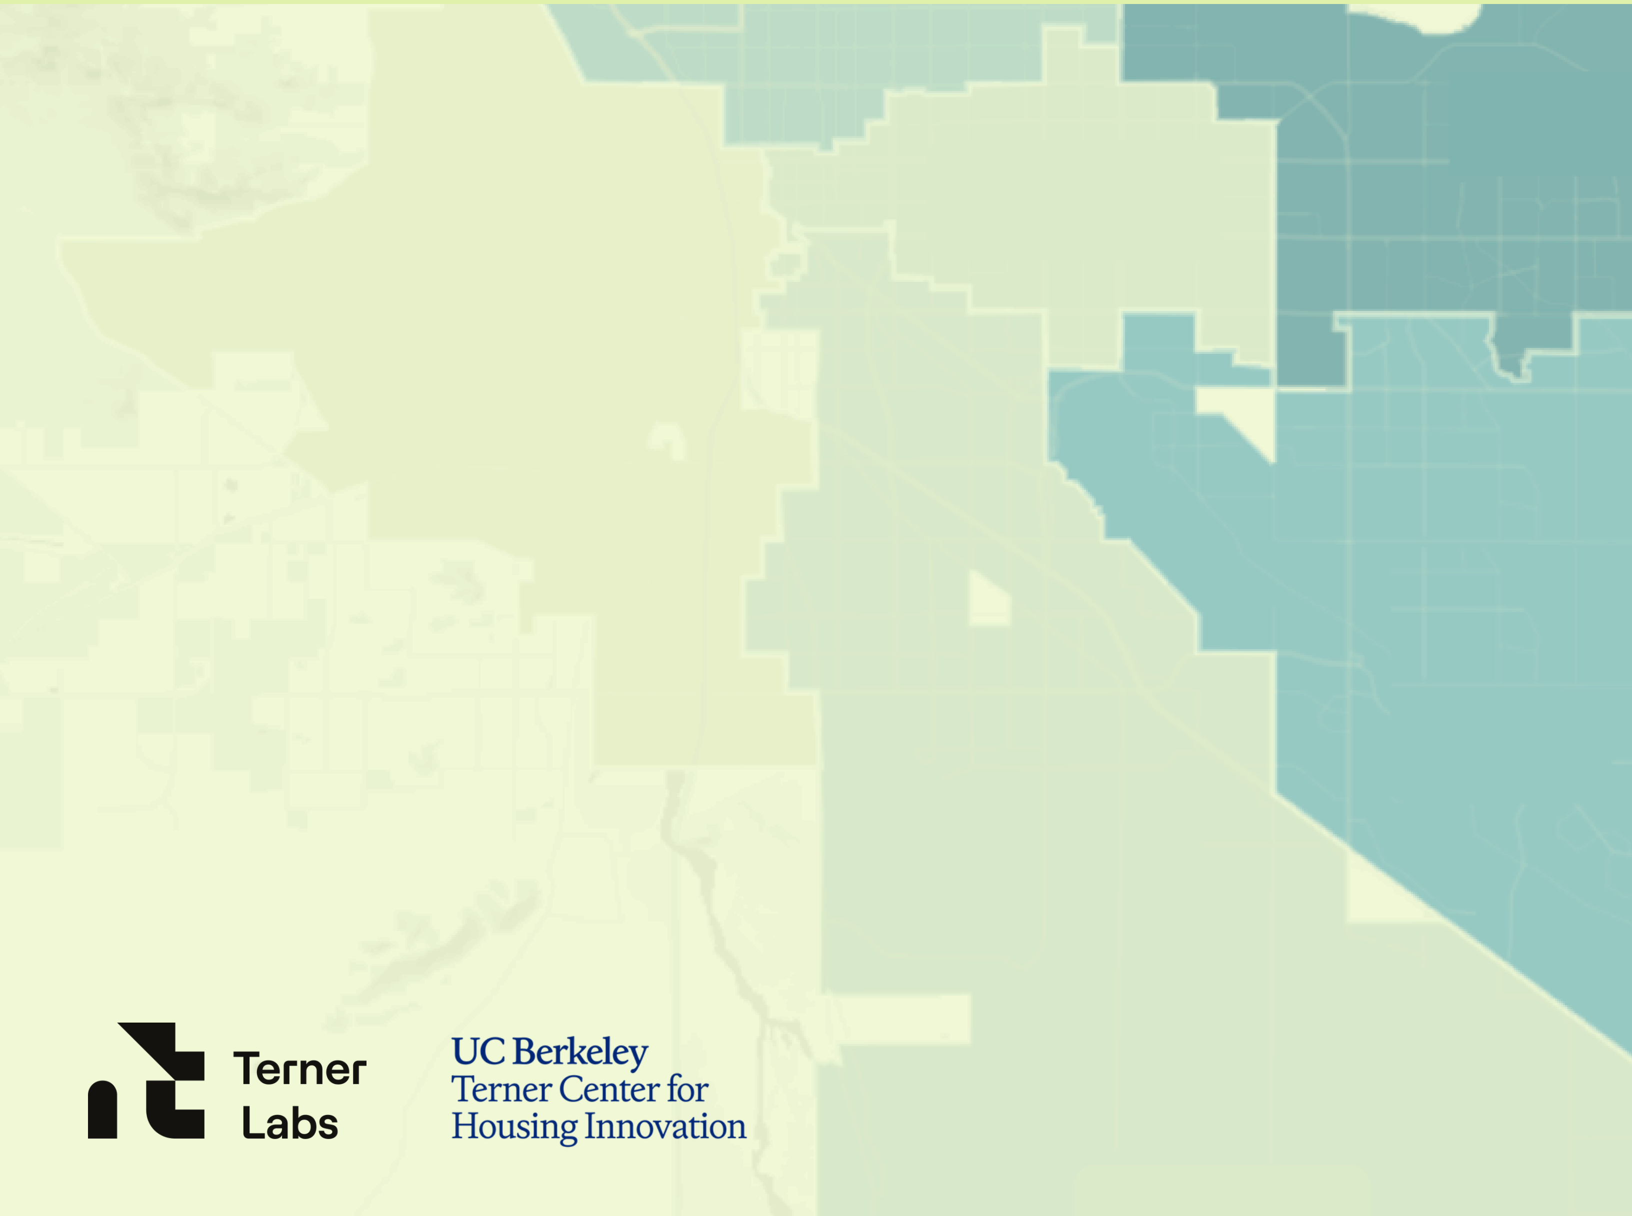

Terner Labs

UC Berkeley
Terner Center for Housing Innovation

BACKGROUND AND PURPOSE

A growing body of research shows that local market conditions—such as land values, construction costs, regulatory environments, and baseline zoning capacity—shape how housing development responds to policy changes. As a result, reforms that meaningfully expand housing supply in one city may have far more limited effects in another. Understanding these differences is critical for policymakers: without accounting for local market dynamics, even well-intentioned housing policies may fail to produce the intended increases in production.

Terner Labs has developed a tool—the Terner Housing Policy Simulator (Simulator)—to help researchers and policymakers understand the implications of adopting housing policies in specific jurisdictions. It uses a complex set of models that forecasts development based on thousands of inputs. The Terner Labs team works closely with research partners to develop and modify the inputs, testing policy details and sensitivities. The Simulator takes into consideration local economic conditions, zoning rules, historical development patterns, and financial metrics of potential projects. Based on these inputs, it estimates the likelihood of a developer building multifamily market-rate housing on a given piece of land and calculates how this might change if a specific policy were applied.

To make the analytics more broadly accessible, Terner Labs has developed a public-facing interactive visualization of the Simulator that allows users to explore how various policy scenarios and economic conditions affect housing development across specific cities, starting with Denver, CO; San Diego, CA; and Tucson, AZ. The Simulator Visualization provides a user-friendly interface, making it easy to compare how housing development responds to different policy and economic conditions.

This brief provides a high-level overview of the models that drive the Simulator, to ensure that the assumptions and limitations of the analysis are clear. More detail about the Simulator's design and underlying assumptions is available in the accompanying methods documentation.

The Simulator Visualization illustrates the capabilities of the Simulator under particular scenarios and provides a useful framework for exploring development scenarios, but it cannot fully capture the complexity of real-world dynamics. Analyses should therefore be interpreted with recognition of the model's scope and limitations. For an example of how the Simulator Visualization's results can be interpreted, see the Terner Center's accompanying analysis.

HOW THE SIMULATOR WORKS

Models

Model to assess zoned residential capacity of each site

Each parcel is unique, and the optimal development project for one parcel is not necessarily the optimal project for another. With this in mind, the Terner Labs team takes a parcel-first approach when developing a Simulator profile for a new city.

The first model evaluates every parcel in the city and incorporates geographic characteristics that may affect development—such as parcel slope, the presence of creeks, and other physical constraints—as well as the city's existing zoning regulations. These zoning rules inform development parameters such as setback requirements, parking requirements, Floor Area Ratio (FAR) limits, height restrictions, and allowable unit counts.

The model then identifies all plausible and legally allowed multifamily building configurations that could be developed on each parcel. This process evaluates different combinations of building types, unit counts, and physical configurations that are consistent with both the parcel's characteristics and the city's zoning rules.

From this step, the model calculates the highest number of units that could theoretically be developed on each parcel in the city. This metric is referred to as Maximum Dwelling Units (Max DU).

Model to assess economic feasibility

Next, the Simulator uses an economic model that mirrors the pro formas used by developers when evaluating potential projects—essentially, a financial calculator used to estimate a project's potential profitability.

A pro forma takes into account expenses—such as land acquisition costs, financing costs, carrying costs and interest rates, demolition and construction costs, and operating costs—and compares them to potential profits. These profits are operationalized as the rents that the developer expects to receive after the units are constructed or the potential resale value of the leased-up property to an investor. The economic model calculates the net present value (NPV), as well as complementary financial metrics, such as the residual land value ratio and NPV per equity dollar, for each plausible scenario identified in the model to assess zoned residential capacity of each site, described above.

The model then identifies the parcel-level development scenario with the optimal financials and the number of dwelling units associated with that configuration. This is referred to as the Optimal Dwelling Units per parcel, or Opt DU.

Model to assess the probability of development

Finally, the economic model is run on city-specific data from the past to understand what is likely to be constructed in the future. Data from a specified year in the recent past is run through the entire process described above to calculate the financial metrics and Opt DU for each parcel. These metrics are paired with one additional data point: whether that parcel was ultimately developed.

Using a trained logistic regression model to calculate the likelihood of development for parcels in the current data, the probability of development for each parcel is calculated.

The probability of development is then multiplied by the Opt DU figure for each parcel to calculate a metric called Expected Dwelling Units, which is then annualized and made comparable across cities.

On its own, parcel-level data is difficult to interpret. As with any model, it is not appropriate to attribute model results to any single observation. However, when considered in aggregate, the model's predicted trends and patterns become meaningful. The Expected Dwelling Units metric is aggregated to the neighborhood level and displayed in the Simulator Visualization. The metric is also aggregated to the city level to calculate the final predicted number of units created by a scenario.

Quality control

To ensure the reliability of the results, the Terner Labs team conducts a series of diagnostic tests at multiple stages of the modeling process to identify outliers, implausible values, or unexpected model behavior at the parcel and city levels. The team reviews these outputs in detail, cross-checking results against known market conditions and development patterns. Where necessary, inputs are refined and manual adjustments are made to correct data issues.

Economic and Policy Scenarios

The process described above is repeated for each economic and policy scenario. First, the team runs the model under a status quo policy environment, combined with a baseline economic scenario, in which economic variables remain roughly consistent with current conditions.

For every policy scenario, the model is evaluated under three economic conditions: unfavorable conditions for development, baseline, and favorable conditions for development. Users can toggle between these scenarios to explore how different policies may affect housing development under different economic environments.

While the Simulator is not designed to predict the exact number of units that will be built in the future, it is a powerful tool for understanding how development feasibility responds to different policy and economic environments.

USES AND APPLICATIONS OF THE SIMULATOR VISUALIZATION

The Simulator Visualization displays results for five policy scenarios and three economic conditions in three cities: Denver, CO; San Diego, CA; and Tucson, AZ. It provides an easily navigable, publicly available interface with clear explanations of each result.

There are a number of ways that the Simulator Visualization can meaningfully inform housing policy debates.

Policy Experimentation

The Simulator Visualization provides a powerful platform for understanding how various policy choices may affect future housing supply. It is often unclear how a proposed reform will play out in a given city, and evaluating policy changes can take years as projects move through the development pipeline. The tool allows researchers and policymakers to approximate how different policy changes may affect development feasibility under current conditions. By testing policy scenarios within the model, users can explore how changes to regulations (such as parking requirements or zoning rules) may shift the economics of development.

The Simulator Visualization also provides a way to explore how different policies interact with one another. Housing development is shaped by a complex set of regulations that operate simultaneously, and the effect of any one policy often depends on the broader regulatory environment.

The Simulator Visualization allows users to see how existing broad upzoning policies in Tucson might interact with a missing middle policy experiment, or how San Diego's decades-long investment in public transit interacts with a transit upzoning policy scenario.

Geographic Comparisons

The Simulator Visualization also makes it easier to compare how similar policy changes play out in different cities. As illustrated in the Terner Center's analysis, these cross-city comparisons can be informative. By applying the same set of policy and economic assumptions across multiple jurisdictions, the Simulator Visualization helps reveal how local conditions shape development outcomes. These comparisons allow researchers to move beyond simply asking whether a state or local policy works and instead explore how its impacts vary across jurisdictions.

The Simulator Visualization can illustrate how a given policy change may affect different parts of a city, enabling neighborhood-level comparisons. This allows researchers to explore how development feasibility shifts across neighborhoods under different policy and economic scenarios. Exploring this nuance within cities also allows policymakers to assess whether the additional supply generated by a policy is aligned with other goals, such as providing access to opportunity and reducing climate pollution.

Isolating Policy and Economic Impacts

The Simulator Visualization can help disentangle the effects of policy from broader economic conditions. The tool allows users to toggle between different economic scenarios—unfavorable, baseline, and favorable conditions—while holding policy assumptions constant.

This structure allows researchers to examine how sensitive policy impacts are to broader economic cycles. Some policies may only meaningfully expand development feasibility when market conditions are favorable, while others may continue to support development even when economic conditions weaken.

By exploring these interactions, the Simulator Visualization helps highlight how policy effectiveness can vary depending on the economic context in which it is implemented.

LIMITATIONS

While the Housing Policy Simulator and the Simulator Visualization are powerful analytical tools, it is important to recognize their limitations. As with any model, the results they produce depend on a set of assumptions and simplifications. Understanding these limitations is essential for meaningful interpretation of the results.

The current version of the Simulator focuses exclusively on multifamily rental housing. It does not model the development of for-sale housing, single-family housing, or subsidized affordable housing. Because the economics of rental and ownership development differ in important ways, the Simulator's results should be interpreted specifically as reflecting rental development feasibility.

Data are aggregated for a number of reasons, most of which are common to statistical models. For example, the model is constrained by omitted variable bias: there are simply too many unobserved variables at the parcel level (i.e., information not captured by the model) to make predictions about specific parcel-level changes. These omitted or unobserved variables include the individual decision-making that ultimately determines whether any given parcel is developed—across developers, landowners, and local building officials and permitting offices.

In aggregate, however, these factors are smoothed into generalizable patterns that help explain outcomes at a broader level. The results should therefore be interpreted as estimates of how development feasibility changes under different scenarios, rather than precise predictions about where specific projects will occur.

The Simulator also relies on a single statistical model to estimate the probability of development across parcels. In reality, different types of projects may respond differently to economic conditions. For example, smaller missing middle developments may face different financial constraints and decision dynamics than larger multifamily projects. To the extent that these differences exist, they may not be fully captured within a single probability model. Additionally, based on the sample size of observed development in the lagged data, overall data quality, and the degree to which past development has followed predictable and consistent patterns, probability models in certain cities perform better than others.

The model also relies on a number of economic assumptions that inevitably affect its results. These assumptions include interest rates, loan-to-construction cost ratios, multifamily cap rates, investors' preferred rates of return, and absorption rates, as well as future rents. While these assumptions are carefully calibrated using available data and through conversations with experts, they do represent a simplified version of the real market dynamics that underpin development decisions.

There are also general equilibrium limitations to the Simulator broadly and the Simulator Visualization specifically. In plain terms, this means that development is predicted without regard to what else might be built. For example, if the amount of development in a city were to double, as is predicted in the "all policies combined" scenario in San Diego, this could strain local permit offices and other resources. It could also drive up construction costs if there are not enough construction workers available. A large uptick in supply may also drive down rents and thus change the financial calculus. Accordingly, policy changes that are closer in range to historically witnessed variations are likely to result in more accurate predictions.

While the model highlights how policies might spur new development, it cannot capture the experiences and perspectives of current residents or a neighborhood's cultural and historical importance—factors that are central to equitable and context-sensitive policy decisions. It is not meant to replace or ignore planning processes focused on advancing other important policy goals, such as reducing vehicle emissions, managing development in areas vulnerable to displacement, or employing strategies to affirmatively further fair housing.

We do not outline these limitations to diminish the value of the Simulator. Rather, we intend to clarify the appropriate role of the tool. It is best understood as a framework for exploring how policy and economic conditions shape development feasibility, rather than as a precise prediction of future housing production.

CONCLUSION

The Simulator opens up new possibilities for understanding how policy choices and economic conditions interact to shape housing development. It can be used to test potential policies, highlight geographic variation, and make the potential outcomes of complex trade-offs more visible. Its outputs are conditional on the model's assumptions and should be interpreted within that context—not as definitive results, and not meant to be used in isolation from equity, environmental, and other key policy considerations. Yet despite these limitations, the Simulator reflects the Terner Center and Terner Labs' commitment to bringing rigorous, data-driven evidence into policy deliberations.

Acknowledgments

The Terner Housing Policy Simulator was developed by Terner Labs. Its analytics are powered by MapCraft.io software. This interactive visualization of some Simulator outputs was designed and built by Graphicacy.

We are grateful to Terner Center faculty research affiliates Nick Marantz and Moira O'Neill for their review of this brief, and to several Terner Center colleagues for informing and guiding this research, including Michelle Boyd, Alex Casey, Arian Dehnow, Sarah Karlinsky, Megan Moore, and Carolina Reid.

The Terner Housing Policy Simulator has received generous support from Airbnb, Arnold Ventures, Ballmer Group, Coefficient Giving, the Conrad N. Hilton Foundation, JP Morgan Chase & Co, Meta, San Diego Regional Policy Innovation Council, San Francisco Foundation, Sobrato Philanthropies, and the Weingart Foundation.