Land subsidence: Environmental risk in housing markets in Mexico City*

Lucy Hackett†, Carolina Rodríguez Zamora‡

March 19, 2026

LINK TO LATEST VERSION

Abstract

We study the costs of and the housing market response to subsidence– the sinking of land areas due to groundwater over-extraction– in Mexico City. We propose an equilibrium model of the housing market that features housing re-development in the face of an evolving environmental hazard that has both realized and expected future impacts to home quality. Our model highlights that while realizations of subsidence attract development by lowering the opportunity cost of re-building units, information frictions affecting the capitalization of future risk lead to an over-supply of housing in risky areas. Guided by model-derived estimating equations, we exploit quasi-random variation in sinking intensity to estimate the impact of both realized and future subsidence on home values and redevelopment. We find that realized subsidence imposes substantial costs, lowering prices by 1.5% on average. However, prices are unresponsive to measures of expected future sinking, and novel survey evidence suggests that information frictions affect the ability of homebuyers to capitalize predictable future risk. Consistent with model predictions, units that have experienced more sinking are *more* likely to be redeveloped, as these have lower opportunity cost of being re-built. Evaluating welfare using our parameter estimates implies that current rates of subsidence cost Mexico City a total of $33 billion USD, about $18 billion of which are due to information frictions that inefficiently increase the housing stock in the riskiest neighborhoods.

*We are grateful to Marco Gonzalez-Navarro, Kirill Borusyak, Lucas Davis, and Joe Shapiro for their guidance and support. We also thank Lorenzo Aldeco, Nano Barahona, Benjamin Faber, Cecile Gaubert, Amir Kermani, Antoine Levy, Francesc Ortega and numerous seminar participants for feedback and suggestions. We thank Dr. Natalia Volkow and the staff members at the Microdata Lab at INEGI and Rodolfo Oviedo Moguel for invaluable data support. Housing data was accessed through the Econlab at Banco de México, who collected and processed the data as part of its effort to promote evidence-based research and foster ties between Banco de México's research staff and the academic community. Inquiries regarding the terms under which the data can be accessed should be directed to: econlab@banxico.org.mx. We gratefully acknowledge funding from the Clausen Center, the Fisher Center, and CEGA at UC Berkeley.

†Department of Agricultural and Resource Economics, UC Berkeley

‡Banco de México

Acknowledgement of support

We acknowledge financial support for this project provided by the Clausen Center, the Fisher Center, and CEGA at UC Berkeley. We confirm that we have no financial or other relationships that could present conflicts of interest in relation to this work.

1 Introduction

As groundwater levels decline in places of high population density worldwide, policymakers are increasingly concerned about the environmental and economic consequences of falling water tables (Carleton et al., 2024; Rodella et al., 2023). One important yet under-studied externality resulting from groundwater decline is land subsidence, the sinking of the Earth's surface due to groundwater depletion. Land subsidence occurs when aquifer levels within a groundwater basin fall, leading to the irreversible loss of elevation for the overlying land, along with any infrastructure located in the area.[1] Subsidence is an externality because aquifers are hydraulically connected: a drop in water levels caused by pumping in one location eventually equalizes over space, causing land to sink in potentially distant areas within the same basin.

While the physical sinking process is slow– with rates in "severe" cases ranging from 10 to 30 centimeters per year of elevation loss– the cumulative impact creates significant hazards for the built environment. Over time, structures endure increasing stress that leads to tilting, cracking, and eventual loss of structural integrity as the foundation shifts below them. Large public infrastructure is also vulnerable to uneven sinking across broad areas, which causes fissures and sinkholes in roads and damages subterranean assets like water mains and subway tunnels by distorting their required slopes (Cigna and Tapete, 2021; Gambolati and Teatini, 2021). These structural and infrastructural issues are particularly acute in large cities overlying severely over-drafted groundwater sources, with prominent examples including Jakarta, Tehran, and Mexico City.[2] These major world cities have experienced cumulative losses of several meters in recent decades (Bagheri-Gavkosh et al., 2021; Werner et al., 2013), resulting in severe consequences for the physical capital overlying sinking areas.[3]

[1] Subsidence can also be induced by intensive mining activities or volcanic eruptions; in this paper we focus on groundwater depletion as this is the predominant cause in cities.

[2] Subsidence "hot spots" can be found on every continent, and there are particularly fast-sinking areas in coastal East Asia, Iran, Mexico, and the United States with areas of intense sinking in 45 states in the United States and more than 95 cities in China (Bagheri-Gavkosh et al., 2021). As Figure B1 illustrates, this issue tends to be particularly acute in developing countries with water management challenges.

[3] In subsiding coastal areas, it is estimated that up to 30% of projected sea level rise is in fact due to

We study the costs of subsidence and the equilibrium housing market response to this evolving hazard in one of the largest and fastest-sinking urban areas in the world: Mexico City. Mexico City has experienced rapid sinking with the worst-hit parts of the city losing approximately a meter of elevation every 3-4 years since 1950 (Chaussard et al., 2021). Figure 1 depicts some examples of what the consequences of subsidence look like in the city: sinking buildings become damaged and potentially uninhabitable, or require extensive repairs to remain usable. Additionally, subsidence has caused damage to roads, water mains, and subway tunnels, especially in areas sinking unevenly.[4]

Slow-moving disasters like subsidence are challenging to conceptualize empirically because they have both *realized* impacts on homes and expected *future* impacts. Therefore we structure our approach to estimating the costs of subsidence by proposing a spatial equilibrium model of the housing market that incorporates three key features of the problem: subsidence damages homes; there may be frictions affecting the capitalization of future damages; and re-development can 'reset' housing quality by erasing the realizations of subsidence. This model guides us in specifying theory-consistent estimating equations, and allows us to explore how frictions affecting the pricing of expected future impacts to a home affect the equilibrium distribution of the housing stock. This mechanism is motivated by our own empirical findings and by a repeated observation in the literature that the capitalization of environmental hazards may lag behind predictable impacts in the absence of sudden and extreme events that correct the market.[5]

Our framework highlights the dual role that evolving hazards that affect housing play in

these areas sinking into the sea (Nicholls et al., 2021), and currently 11 coastal cities are at risk of becoming totally inundated by the end of the century. In a striking example, Indonesia has developed a plan to move the capital city away from Jakarta, which is rapidly subsiding below sea level due to groundwater extraction (Associated Press, 2022).

[4]While sinking is less damaging when it is relatively even, even on a single plot the texture of the underlying geology means that sinking at the rates seen in Mexico City generally cannot occur so evenly as to not damage a structure.

[5]We discuss this literature in detail alongside our contribution, but studies on the capitalization of hazards such as floods and storm damage consistently find that extreme realizations or floods or hurricanes drive the capitalization of risk into home prices and insurance take-up, whereas predictable risk over space otherwise lags behind costs (Hino and Burke, 2021).

the incentive to redevelop. While worsening expectations about future impacts dissuade re-development, the depreciation caused by subsidence lowers the opportunity cost of renewing existing units, leading to increases in the probability of re-development on the extensive margin. Frictions affecting the pricing of expected future damages lead to inefficiently high housing supply in hazardous areas, as developers receive too high a price for units that will sink in the future. The model elucidates that three parameters are central to this relationship: the willingness to pay of residents to avoid subsidence, the extent of pricing frictions, and the price elasticity of housing supply. We estimate these by bringing together original survey evidence on resident beliefs and information about subsidence with administrative data on bank appraisals, data on housing developers, and novel, historical subsidence measures that we construct.

To measure subsidence, we combine original processing of radio wave-based satellite products with pre-existing geophysical data, resulting in the longest, highest-resolution panel of sinking in Mexico City to date. We do so using a process called Synthetic Aperture Radar interferometry, a tool used by geophysicists to measure distortions in the Earth's surface with applications related to subsidence, volcanic eruptions, and glacier movement, among others. This new dataset measures annual sinking rates in Mexico City from 2007-2024 at 100m resolution, and we use these data to infer the subsidence history of individual properties in our appraisal and survey data, measuring how much a specific property has sunk over its life and its sinking post-transaction.

With these measures in hand we take our model to the data, estimating the willingness to pay of home buyers to avoid subsidence and the capitalization of expected future subsidence into home values using the universe of mortgage appraisals in Mexico City. In order to causally identify these parameters, we exploit the fact that sinking on a given plot is caused by non-local groundwater pumping to isolate exogenous variation in sinking intensity, and implement a repeat appraisal design that allows us to flexibly control for all static quality and neighborhood differences between different buildings.

We find that subsidence has important impacts on home values. A meter of sinking causes a loss of 6% in appraised value, and the average home in our sample lost 1.5% of its value to sinking over 10 years (conditional on not being re-developed). Exploring mechanisms of how sinking impacts home values, we find evidence that losses are a product of both physical impacts to the home and increased maintenance costs. Including uneven neighborhood sinking explicitly in estimation shows that a neighborhood's relative sinking compared to its neighbors also matters: The average relative change in the altitude of a neighborhood causes a 2.4% loss in appraised value over 10 years, and is correlated with an almost doubling of the probability that residents report issues with flooding in their neighborhood.

We then turn to the question of whether expected *future* sinking is capitalized into home values, and find evidence of substantial frictions preventing the accurate pricing of future sinking risk. Evidence from the model-derived specification shows that home prices have no detectable relationship with the expected future sinking of the building, and this is robust to multiple ways of specifying expectations. While this estimation approach faces the inherent challenge of accurately specifying expectations, the failure to detect any price response to expected future sinking suggests that home values do not internalize these future impacts.

We implement a novel survey of homeowners and long-term occupiers of their housing units in Mexico City to shed light on the mechanisms driving this lack of capitalization. The survey documents residents' beliefs about subsidence and the direct impacts of subsidence on their home and neighborhood, and our survey findings validate that information frictions about subsidence may play a prominent role in the housing market. We find that residents do not use ex-ante available information to infer future sinking risk. Furthermore, residents of high-sinking plots experience high rates of surprise about the severity of sinking they have experienced and these same residents are the least likely to report having considered sinking when they moved in, a finding robust to rich controls for respondent socioeconomic status and survey measures of risk and discounting preferences. Together, these facts point to the presence of important frictions in residents' ability to price future sinking into home prices,

and provide an estimate of the information friction parameter in our model.

In our last estimation step we estimate the supply elasticity of re-development using a survey of housing developers. We infer that a property was re-developed by matching appraisals to the locations of new builds, and identify the supply elasticity by implementing a control function approach that allows us to instrument for prices using uneven sinking across grid cells. We find that a meter of sinking increases the probability that a plot is redeveloped by 9.1 percentage points, a finding that echos our model intuition that realized sinking lowers the opportunity cost of re-developing a plot. Interpreting our estimates structurally provides an estimate of the supply elasticity.

Finally, we use our estimated model parameters to evaluate welfare under counterfactuals that enable us to quantify the total costs of subsidence and decompose these costs into the direct cost sinking and those arising from information frictions. We estimate that current rates of sinking impose a permanent cost in Mexico City of $33 billion US dollars, which represents 17.5% of the total value of housing in Mexico City. Because homes internalize the full stream of future sinking, this cost is the "permanent" cost of sinking at current rates. These total costs accrue primarily to homeowners, but subsidence also lowers developer profits. Converting this total to dollars-per-unit of groundwater pumped allows us to estimate a potential Pigouvian tax on groundwater pumping of $0.89 per cubic meter. This is a large tax, which if passed on to consumers would imply a 137% increase in the average water bill.

We also estimate the potential gains from policies that address the pricing friction in the housing market. As is the case in many developing and middle income countries, there are no disclosure laws in Mexico requiring the reporting of relevant hazards as in the United States, despite the fact that developers have detailed information about subsidence risk at the plot level. We find that fully mitigating the pricing frictions would offset 55% of the total costs of subsidence and shift up to 5% of the housing stock out of the highest sinking intensity plots. In contrast to the counterfactual that slows subsidence, developer profits fall

when pricing frictions are addressed, as these frictions inefficiently inflate the price of new builds on risky plots.

In our last policy exercise we use these cost estimates to conduct cost-benefit analysis of policies that would permit groundwater recharge, as the costs of subsidence represent the benefits of these investments. We focus on four policies that have already been implemented by the city government to some extent: injecting wastewater into the aquifer, repairing the extensive leaks in the domestic water supply system, investing in new water sources to supply the city, and recycling wastewater. We find that repairing leaks and artificial recharge show the most promise, but caution that our analysis is limited by the use of rough or dated cost estimates.

These findings represent the first comprehensive economic analysis of subsidence. While a large body of work has focused on the costs of environmental externalities such as air pollution, floods, and wildfires, subsidence is a natural hazard that has been relatively understudied by economists despite the fact that it is an important groundwater externality affecting many major cities around the world.[6] We contribute to the sparse economics literature that has studied subsidence (Willemsen et al., 2020; Yoo and Perrings, 2017) by combining improved causal inference with market valuation to estimate costs, and by providing empirical and theoretical analysis of the equilibrium housing market response to sinking. Our original satellite work and survey contribute novel evidence on the physical impacts of sinking on buildings and public infrastructure, which to our knowledge is the first representative measure of damages for a city.

More broadly, our framework provides a way to think about the equilibrium response of housing to evolving environmental hazards, with a particular focus on contexts with pricing frictions or imperfect information. We build on recent empirical findings from work in

[6]Engineering-style estimates of the costs of subsidence suggest that these may be large, especially in urban areas; case studies from China (Lixin et al., 2010), Indonesia (Mahya et al., 2021), Mexico (Fernández-Torres et al., 2025; Novelo-Casanova et al., 2022), and California (Borchers and Carpenter, 2014; Fowler, 1981) find damages in the billions for specific cities or counties. However, the general approach in these studies is to estimate replacement costs for damaged infrastructure without considering market values or forces.

environmental, urban, and real estate economics that study how information frictions about environmental risk impact housing markets (Bakkensen and Barrage, 2022; Bakkensen et al., 2019; Gallagher, 2014; Gourevitch et al., 2023; Hino and Burke, 2021; Holtermans et al., 2024; Lee and Zheng, 2023; Ortega et al., 2025; Ortega and Taşpınar, 2018; Petkov and Ortega, 2025; Wagner, 2022), and go further by embedding these frictions in an equilibrium spatial model of the housing market to quantify the aggregate costs of these frictions.[7]

By modeling the market response to environmental risk in equilibrium, this paper is related to the literature on urban development and adaptation (Gechter and Tsivanidis, 2025; Hsiao, 2023; Ospital, 2024; Ostriker and Russo, 2023), and more generally to a body of work that studies the costs of slow-moving disasters such as climate change (Balboni, 2025; Desmet et al., 2021; Hsiang et al., 2017). Our model contributes to these papers by modeling redevelopment in the context of an evolving hazard, which allows us to carefully examine the different roles that *realizations* and *expectations* of environmental risks play in driving housing supply and re-development.

Our model insights and empirical results provide a general framework for understanding the consequences of market failures that impede the capitalization of environmental hazards into home values. We show that the housing market's ability to respond correctly to environmental threats hinges on the effective capitalization of information– specifically, the ability of homebuyers to understand those threats. While we focus on information frictions, these insights extend to other barriers that prevent efficient capitalization, such as subsidized flood insurance.[8] Furthermore, our work provides rich empirical evidence on the housing market response to an important externality and the value of disclosure in a middle-income context, a setting where the necessary housing value data is rarely available.[9]

[7] Our forward-looking demand framework also draws on insights from recent dynamic hedonic approaches (Rosen, 1974; Benetton et al., 2025; Bishop and Murphy, 2019), and we propose a way of incorporating information frictions into these methods.

[8] Papers that study the importance of beliefs about the environment for adaptation in other contexts or markets include Patel (2024); Zappalà (2024); Mahadevan et al. (2023); Kala (2017).

[9] One important exception is Gonzalez-Navarro and Quintana-Domeque (2016), who also use professional appraiser data to study road paving in Mexico. Related papers on the value of disclosure from the United States include Dranove and Jin (2010); Frondel et al. (2020); Nanda and Ross (2012); Sinha (2022)

The rest of this paper is organized as follows. Section 2 introduces the context and our satellite data for measuring subsidence. Section 3 presents the economic data we use. We present our theoretical framework in Section 4. Section 5 presents the estimation strategy and results from our survey work and the demand and supply estimation. We present policy analysis in Section 6, and Section 7 concludes.

2 Subsidence in Mexico City: Background and measurement

Mexico City has experienced severe rates of subsidence over the last 100 years, which geophysicists have established is caused by sustained groundwater pumping on the Mexico Valley aquifer (Cabral-Cano et al., 2008; Chaussard et al., 2021; Solano-Rojas et al., 2015).[10] This pumping is primarily conducted by the city government and the federal water authority CONAGUA for the purpose of supplying the city with about two-thirds of its total water consumption, the other third of which is supplied by surface water sources.[11] Because the aquifer is hydraulically connected, pumping anywhere affects the water table everywhere (Cigna and Tapete, 2021), which is why the city itself experiences high rates of subsidence despite the fact that there is no groundwater pumping conducted within city limits.[12]

Thus sinking rates are primarily a product of two factors: the height of the water table and the propensity of overlying geologies to compact and sink. Given that the water table is currently far below historic levels, sinking accelerates when the water table is relatively

[10]Mexico City has a long history of water management and scarcity; the Aztec empire had built their capital city around and interwoven with the large lake system in the area, but the lake was subsequently drained by the Spanish colonial government. Because such a large part of the city is built on the soft, silty lake bed of the drained lake, this area is particularly prone to subsidence.

[11]Bringing surface water into Mexico City is very expensive due to how far large surface water sources are and the fact that most sources need to be pumped over the surrounding mountain ranges. The most important surface water source for the city is the Cutzamala system, which involves seven dams and over 330 kilometers of canals and pipes that use over 2 billion kWh of electricity per year to transport water to Mexico City.

[12]The government has limited all pumping on the aquifer to the periphery of the city– a policy decision made after an episode of intense sinking in the 1950s (SACMEX, 2012), which resulted from the fact that water table levels take time to equalize over space. Pumped water is first sent to treatment facilities then to distribution stations, making the relationship where water is pumped and supply of water to a specific area tenuous.

lower (and vice versa), and could only be fully stopped if the aquifer is recharged either naturally or via water injection (Gambolati and Teatini, 2021; Liu et al., 2019). While it is conceivable that the weight of human development could increase sinking, empirical studies in Mexico City have found no correlation between localized sinking rates and population density, pumping rates, or land use type (Chaussard et al., 2021).[13] Our own analysis in Appendix E.1 supports this conclusion, as we find that construction on a pixel does not predict subsequent sinking rates.

2.1 Measuring subsidence

Historically, subsidence was measured by inserting poles deep into the ground at discrete points and measuring the change in how exposed the pole became over time as the ground subsided away from it; this method was used as early as the 19th century by the Spanish colonial government to measure sinking in Mexico City (CONAGUA, 2009). Today, measuring sinking over larger areas at fine geographic scale is possible thanks to the inclusion of Synthetic Aperture Radar (SAR) instruments on satellites.

SAR instruments emit energy at the Earth in the form of radio waves and record a transformation of the number of oscillations that the wave undergoes between the earth and the satellite. While these measurements are meaningless on their own, constructing the difference between two images through a process called interferometry creates high-accuracy measures of vertical displacement.[14] Interferometry is the go-to tool of geophysicists studying events involving land deformation, and because this is a popular and well-established tool the measurement error associated with SAR measurements of subsidence is well-studied, with inaccuracies on the scale of 1-10 millimeters for modern satellites (Wu and Madson, 2024),

[13] Studies have found that the weight of New York City has caused sinking of a couple millimeters, but New York has much higher population density than Mexico City, and this effect is orders of magnitude smaller than the rates observed in Mexico (Parsons et al., 2023).

[14] These calculations are deterministic in nature, and are not the result of machine learning algorithms like those well known in applications such as remote sensing of pollution or applications that use image-based sensors.

an order of magnitude smaller than the subsidence rates we measure in Mexico City.[15]

We apply gold standard tools for conducting interferometry and subsidence time series construction that correct for distortions caused by terrain, tropospheric factors, and orbital error to construct an 18-year panel of subsidence rates at the 100-meter pixel level for Mexico City (Fattahi et al., 2017; Yunjun et al., 2019; Yu et al., 2018b). We do this by combining three series of data, summarized in Appendix Figure B2, performing original interferometry using data produced by the ALOS-1 (2007-2011) and Sentinel-1 (2021-2024) satellites to complement measures created by geophysicists using Sentinel-1 readings from 2014-2020 (Chaussard et al., 2021).[16] By using three separate series, we are able to address the computational challenge presented by calculating a single time series: The raw Single-Look Complex (SLC) SAR data and its processed output for just one orbit direction (ascending) for 5 years contains almost five terabytes of data and takes over three weeks of continuous computational time to process.

In order to characterize the gradient of sinking over larger spatial areas, the principal risk factor for damage to roads and water mains, we define a 500-meter resolution grid (the median census tract size) over Mexico City and construct two complementary measures of how uneven sinking is over space. The first measure is the standard deviation of sinking among pixels overlaying the grid cell, an intuitive measure of unevenness. The second measure takes inspiration from the horizontal gradient used in geophysics studies such as Fernández-Torres et al. (2020) and calculates the grid cell's *relative* displacement between periods t and t' as the absolute value of the accumulated difference in grid cell sinking and

[15]Interferometry achieves higher accuracy in urban areas because the hard reflective surfaces minimize noise generated by signals intercepting vegetation. We provide a detailed discussion of the sources and potential magnitudes of measurement error when conducting interferometry in Appendix E.1.1, but by using state-of-the-art algorithms we address the main sources of error that could affect our estimates, including distortions potentially generated by construction. Intuitively, construction does not have a strong impact on our measurements for two reasons. One, because satellites take measurements at an angle, interferometry has trouble detecting building height in general. Second, because the construction of the time series involves linearly projecting through high-frequency measurements every 6 days by year on a pixel-by-pixel basis, this minimizes the noise from irregular measurements or changes induced by demolition/construction.

[16]Because no SAR measurements were taken for Mexico from 2012-2013, we linearly interpolate data to fill this gap. Finally, we again perform interferometry to create a panel from 2021-2024 using Sentinel-1.

the average of its neighbors:

$$\text{Abs. relative elevation change}_{gtt'} = \sum_{\tau=t}^{t'} |s_{g\tau} - \overline{s}_{g\tau}^{nbrs}| \tag{1}$$

The absolute value allows us to capture the potentially damaging macro-infrastructure effects of areas that are sinking more or less than their neighbors. Because this measure accumulates over time, it captures the cumulative impact of relative altitude change in the environment, which may matter more for infrastructure like water mains and sewers than unevenness alone.

2.2 Sinking is fast and predictable in Mexico City

Our subsidence data paint a detailed picture of how sinking has evolved in Mexico City in recent decades. Figure 2, Panel (a) plots sinking rates across Mexico City in 2010. The highest rates of subsidence across the city are around 20-25 centimeters per year, rates that are among the highest in the world (Werner et al., 2013). These fast rates are strongly localized in the eastern part of the city, which overlies particularly vulnerable geologies prone to subsidence.[17] Panel (b) depicts the change in the rate from 2010 to 2020, and shows that by 2020 the northeast corner of the city actually slows somewhat whereas part of the Western edge of the valley begin to sink. This onset of sinking is due to some destabilizing drought periods and seismic events have weakened these historically resistant areas. Finally, Panels (c) and (d) plot our measures of uneven sinking, showing that high relative altitude change and standard deviation of sinking are concentrated around the edges of the high-sinking area, as well as around relatively fixed geologic features that maintain their altitude as the neighborhoods around them sink.

Despite the slight changes in sinking intensity over time, subsidence on a broad scale is still very predictable over space due to the concentrated geological features that determine

[17] Geophysicists estimate that this part of the city could sink up to 30 more meters given the thickness of the features of the subterranean geology prone to compaction (Chaussard et al., 2021).

risk.[18] Figure 3 plots the R^2 of regressions that predict the sinking rate on a pixel using sequentially more granular spatial fixed effects, including and excluding the lagged rate.[19] Even without the lagged rate, knowing the 5-kilometer grid cell a plot is located on already explains over half of the variation in sinking, and knowing the pixel location explains almost 80% of the variation. If the lagged sinking rate is included, this predictive power goes up to 83%. Perhaps because of this broad systematic predictability, subsidence is uninsurable in Mexico City, and is explicitly listed in home insurance policies as a non-covered event.[20]

2.3 Survey evidence: Sinking is poorly understood by residents

Because housing markets are forward-looking, residents' beliefs about both past and future sinking matter for the market's ability to respond to this predictably hazard. We provide evidence on this using a novel in-home survey representative of long-term Mexico City residents about subsidence impacts and residents' beliefs and information about subsidence.

2.3.1 Survey design

Respondents were restricted to owner-occupiers (73%) or those who had lived in their home for at least five years (27%). The survey was stratified based on subsidence risk, with two-thirds of the sample drawn from the seven high-risk boroughs and one-third from the remaining nine boroughs.[21]

This first-of-its-kind survey asked homeowners about the damages from subsidence to their housing unit and the urban environment, as well as residents' beliefs and attention to the issue when making housing decisions. First, we surveyed residents on the physical

[18] Appendix F provides detailed discussion of the data generating process of subsidence and its predictability.

[19] We plot the unadjusted R^2 to imitate an "out-of-sample" prediction exercise of adding another year of data conditional on location as opposed to an out-of-sample that conceptualizes adding another spatial unit, which would want to adjust for the number of spatial units used in prediction.

[20] Subsidence was recently included in "home quality insurance" policies that cover INFONAVIT mortgages from the federal government, but these policies only cover the first 10 years of the home's life, within which the most severe impacts of sinking would not yet be observed.

[21] More details about survey sample definition and implementation can be found in the Data Appendix E.2.

quality of their home and neighborhood as well as maintenance expenditures and activities before mentioning subsidence. In the second module, we then mentioned subsidence by name, surveying residents on their knowledge about subsidence, how they thought about subsidence before living in their current home, and their experience with sinking since arriving at their current unit. We use the location of respondents' homes to map their property to our satellite measures of sinking and infer the total amount of sinking their home has experienced since construction (or 2007, whichever is later), the amount of sinking on their property over the last 5 years, and the unevenness measures associated with their grid cell.

2.3.2 Residents face information frictions when evaluating future risk

Subsidence is a familiar issue to Mexico City residents: In our survey, before we mention the word subsidence, 56% of respondents included "sinking" in their response when asked about the likely causes of the degraded state of a photograph of a tilted and cracked building. Additionally, only 18.5% of respondents could not provide a guess when asked where in the city has the most issues with subsidence.

Despite this general awareness of the issue, survey respondents use little *ex-ante* available information to infer risk of future sinking to their home. Figure 4a plots the share of different responses to the question, "How can you know if a home will have issues with sinking in the future?", separated by respondents in high and low risk boroughs.[22] 21% of respondents say it is impossible to know in advance whether a home will have issues with sinking, and despite the strong spatial correlation in sinking, only 21% report that you can infer risk by knowing the home is in a risky area.

Grouping these responses into ex-ante ways of knowing, which do not require that a home be affected yet, such as asking an expert or the risk zone, and ex-post ways of knowing (seeing the effects in the home or plot) in Figure 4b, we find that residents of low sinking boroughs are 47% more likely to report using ex-ante available factors to infer risk. This

[22]This question was free response, and we prompted respondents for more answers until they indicated they could not come up with any more.

suggests that residents with less information are more likely to locate in high-sinking homes, a finding consistent with insights from real estate models that highlight that the least informed buyers will sort into and bid up prices in risky places (Bakkensen and Barrage, 2022).

Not only are residents inattentive to *ex-ante* available information about sinking on average, but they show significant sorting on this inattention and surprise about its impacts. Figure 5 plots the relationship between a measure of expected 5-year sinking on a respondent's plot, estimated as five times the average rate over the panel, and two survey measures of information frictions. In Panel (a), this sinking intensity is plotted against the probability that the respondent reports using any *ex-ante* available information to infer sinking risk. Residents living in units sinking at higher intensity are less likely to use this kind of information to infer risk than residents of low-sinking intensity plots.

Complementing this in Panel (b), we find that residents of higher risk plots are correspondingly more likely to report that sinking on their property has been worse than expected, providing evidence that residents do not buy sinking homes "eyes wide open" about the future impacts they will experience. We show that this finding together with the results in Figure 5 are robust to controlling for respondent and housing unit characteristics, including the education of the respondent, in Table 1. Furthermore, Table A1 in the Appendix confirms that these effects are similar though noisier due to the small sample size when limiting to residents who moved within the last 10 years, assuaging concerns that results are driven by selective recall or habituation.

Residents also underestimate how much sinking they have experienced *ex-post*, highlighting the difficulty of accurately perceiving slow-moving hazards. Figure 6 plots respondents' belief about how much their property has sunk over the last five years to our satellite measure, and find that the slope of this relationship is only 0.05, showing that even in retrospect respondents have little idea about how much their property has sunk in quantitative measures.

Finally, we directly test how accurate respondents' information about the spatial

distribution of sinking risk is in Appendix Figure B3, which grades respondents on a question where we ask them to name boroughs that have had issues with sinking, and shows that the median respondent commits substantial Type 1 and Type 2 errors when trying to name high-risk areas.[23]

These findings could be explained by high discount rates or low risk aversion that lead residents of high-risk areas to systematically ignore future sinking.[24] We implement Falk et al. (2016) instruments for measuring risk preferences and discount rates to test whether these sorting patterns are driven by residents with higher tolerance for risk or higher discount rates locating in high-risk units, a first that we are aware of for the literature on environmental risk and housing markets. Appendix Table A2 shows that discount rates are not correlated with either the expected sinking on a resident's plot, the realized sinking on the resident's plot over the last 5 years or the relative altitude change over the last 5 years. We pick up no correlation with plot-level sinking on measured risk aversion, and if anything residents of neighborhoods with highly uneven sinking show higher measured risk aversion, suggesting that our results cannot be explained by residents with high risk tolerance or discount rates buying sinking homes.[25]

Overall these findings suggest that respondents face substantial information frictions about subsidence, and that these frictions are particularly related to the ability to infer the probability of future subsidence absent realized sinking. We also find evidence that residents sort on information about subsidence, with less informed residents being more likely to live in high-sinking properties. These findings echo insights from papers in the United States that find that *realizations* of floods or hurricanes drive discounts of environmental risks in

[23] We grade respondents on two measures: how many boroughs we consider to be high risk that they name, and how many boroughs they name that we would consider to be high risk. For the former, respondents may interpret the question with different degrees of severity and name a subset of all high-risk boroughs. For the latter, respondents may be aware of outlier high-sinking cases in otherwise low-risk boroughs.

[24] This mechanism is unlikely for housing which has an important asset value, and which should tie discounting of future values to other financial products.

[25] It is also possible that some residents simply do not value the impacts of subsidence on their home, causing them to be inattentive to future sinking. However, if this were the case we would not expect these residents to also report that subsidence has been worse than expected, as they should also be inattentive to realized subsidence.

housing prices (Bakkensen et al., 2019; Gibson and Mullins, 2020; Ortega and Taşpınar, 2018), and that less informed residents sort into riskier settings (Bakkensen and Barrage, 2022). Motivated by these stylized findings we incorporate potential underestimation of the severity of future sinking into our theoretical framework.

3 Administrative data on the housing market

3.1 Measuring home prices: Appraisals

We measure housing values using confidential data on the universe of mortgage appraisals for the purpose of home sales performed by the Federal Mortgage Society (SHF, *Sociedad Hipotecaria Federal*) in Mexico City from 2005-2020, which were made available through freedom of information request.[26] We use property characteristics and exact locations to identify identical units from the same building, a matching process that results in 17.4% of our sample as being identified as a repeat appraisal.[27] We assign sinking to transactions by matching appraisals spatially to pixels using the exact location to calculate the total amount of subsidence that the unit has experienced since the year of construction.

An advantage of our data is that mortgages represent relatively arms-length transactions, which are less likely to be affected by relational contracting or fraudulent reporting for the purpose of tax evasion, two common issues affecting price data from other sources in Mexico (Gonzalez-Navarro and Quintana-Domeque, 2016). However, in order for appraisals

[26] An SHF appraisal is required as a condition of obtaining a mortgage from a private bank or the public mortgage financier, and all SHF appraisals must be done by an SHF-certified appraiser. Thus, these data contain the universe of SHF appraisals, but do not include unofficial appraisals that may have been done for cash purchases of homes. The 2020 National Housing Survey suggests this restriction is important: only 30% of home purchase in Mexico City in 2020 were financed with some form of formal credit, with the majority of homes obtained via self-financing of own construction.

[27] We exactly match appraisals on the built square footage, the square footage of the property, the number of rooms and bathrooms, number of floors, number of parking spaces, and whether the building has an elevator, and require the coordinates of the appraisal to imply a less than 100-meter distance between them. Table A3 in the Appendix describes how repeat sales properties differ from single sales in the data; properties with repeat sales are much more likely to be apartment as opposed to single family homes, and are generally newer, smaller, and less valuable than the full sample.

to capture the true valuation of subsidence by homebuyers as opposed to the appraiser's algorithm, they must contain a strong signal of the transaction price ultimately paid. Conceptually, this is likely as appraisals are used for the purpose of valuing a mortgage, and both appraisers and the buyer have a strong incentive to appraise the home at the price a buyer is willing to pay. We validate this using a confidential dataset of appraisal and transaction prices called the R04, whose access was provided by the *Banco de México*. The R04 data only contain approximate locations, but allow us to directly compare appraisals and transaction prices.

Using these data, we find that 73% of observations have identical appraisal and transaction values, and that appraisal values closely track transaction prices. Figure 7 plots transaction prices against appraisals, showing that points are closely centered around the 45° line.[28] Appendix Table A4 reports coefficients and goodness-of-fit statistics for regressions of transaction prices on appraised values with and without an intercept. We estimate a coefficient of determination of 0.97 (Column (1)), highlighting the close relationship between appraisals and transaction prices. Importantly, Column (3) shows that deviations of appraisal values from sales prices are not correlated with subsidence; total sinking on the locality from 2007-2020 is not a statistically or economically significant predictor of the absolute deviation of appraised values from transaction prices.

3.2 Housing development

We examine the supply side of the market using the DIME (*Dinámica del Mercado Inmobiliario*) dataset, which is a quarterly survey of new housing developers. These data are compiled by a private real estate consulting firm and access was granted through *Banco de México*. From this survey we construct a dataset of new housing developments from 2007-2020, which are new construction projects that comprise one or more housing units.[29]

[28] Figure B4 in the appendix reports the distribution of deviations between appraised and transaction values, showing the full distribution of differences.

[29] Given that new homes comprise 70-80% of the housing market in Mexico and that real estate developers built 60% of new homes, the DIME data observes 42-48% of the housing market in a given quarter

We use these data in two ways. First, we use a combination of address matching and spatial methods to match new developments to appraisals previous to when construction started to identify appraisals that were sold to developers. Because addresses are not always consistently recorded in the data and may change when a plot is redeveloped, we likely miss some appraisals that subsequently became redeveloped in our data. We address this in an alternative analysis that characterizes the geography of new development by creating a panel at the pixel level that measures the number of units built and the probability of new development.

4 Theoretical framework

In order to structure estimation of a hazard that has both *realized* and *anticipated* impacts, we propose a model of the housing market that incorporates forward-looking housing demand that values subsidence as a negative home attribute. Our model yields welfare expressions that allow us to quantify the costs of subsidence and estimating equations for the key parameters needed to conduct welfare analysis.

Time is discrete, and in each period households rent housing from property managers who serve as intermediaries between renters and a housing developer who builds the housing stock. The developer purchases existing housing from property managers and builds new housing that is subsequently sold back to property management firms. Each period begins with an existing continuous measure of housing on each plot j, H_{jt-1}, which is owned by property management firms. There is an initial period $t = 0$ with a measure H_{j0} of housing on each plot. In each period t, the following sequence occurs:

1. The developer draws cost shocks for each plot and all housing units go up for sale on the

(Rodríguez Zamora, 2010). This is likely a lower bound for Mexico City, which has a more formal construction sector than other areas of the country. Appendix E.4 discusses coverage in more detail. There were 20,380 new developments over the study period, which contain 977,221 new housing units. The median housing development included 20 new units, but some developments have hundreds of new units.

market to own housing.[30] Units may be purchased by another property management firm or the developer. If they are purchased by the developer, then new units are built, which have zero realized sinking and re-enter the market for purchase by property managers.[31] The ownership market clears on the price of owning a home p_{jt}, and each plot has a housing stock H_{jt} that is fixed until the next period.

2. Renters draw iid preference shocks over plots and make a discrete choice over which plot to live on. Property management firms supply rentals from the housing they purchased in the previous step. Equilibrium rents R_{jt} adjust to clear the rental market.[32]

We describe these steps in reverse order, as the decision of property management firms to purchase is a function of their expectations over the rental market.

4.1 Renters

A continuum of renters in Mexico indexed by i with mass $\overline{L}$ decide which discrete plot j to rent a unit of housing on. They draw preference shocks over plots in each period, and can choose from plots in Mexico City or choose the outside option of living outside the city ($j = 0$), where rental units are supplied perfectly elastically at zero price. All rental units on a given plot are identical, and they derive utility/disutility from paying rent R_{jt}, how much the unit has sunk to date, S_{jt}, other observable attributes Z_{jt} and unobserved attributes e_{jt}, each of which are unaffected by subsidence. Total subsidence S_{jt} is the accumulation of sinking rates on the plot $s_{j\tau}$ from the time of construction to t:

$$S_{jt} = \sum_{\tau=\text{Year built}_j}^{t} s_{j\tau} \tag{2}$$

[30] By putting the full stock on the market in each period, we abstract from housing inventory and matching buyers to available housing.

[31] We assume the market lasts long enough for these units to compete with a meaningful share of the stock when put back on the market.

[32] We do not take a stand on the length of a period in terms of a fixed rental contract; we allow renters to be mobile so that they have no need to form expectations about sinking, and this reflects the fact that rental contracts are generally unwritten and fluid in Mexico City.

Their utility is given by:

$$V_{jt}^{i} = -\log(R_{jt}) + \gamma S_{jt} + Z_{jt}\delta + e_{jt} + \epsilon_{jt}^{i} \tag{3}$$

We assume $\epsilon_{jt}^{i} \sim$ iid EV(1) with variance 1, and normalize mean utility from the outside option to zero: $V_{0t}^{i} = \epsilon_{jt}^{i}$.

The probability that a renter chooses plot j is:

$$\lambda_{jt} = \frac{\exp(-\log(R_{jt}) + \gamma S_{jt} + Z_{jt}\delta + e_{jt})}{1 + \sum_{l} \exp(-\log(R_{lt}) + \gamma S_{lt} + Z_{lt}\delta + e_{lt})} \tag{4}$$

and total rental demand is given by:

$$H_{jt}^{D} = \overline{L}\lambda_{jt}(R_{jt}) \qquad \forall\, j \tag{5}$$

4.2 Property management firms

In each period a mass of competitive property management firms, indexed by $k \in \Omega$, purchase housing (either used or newly developed) and rent units to renters. Their expected profits on a plot j from holding a measure of housing $H_{jt}(k)$ are given by:

$$\tilde{\mathbb{E}}_t(\pi_{jt}(k)) = H_{jt}(k)\tilde{\mathbb{E}}_t\left(\sum_{\tau=t}^{\infty} \rho^{\tau-t} R_{j\tau}\right) - p_{jt}H_{jt}(k) \tag{6}$$

Expected revenue is the present value of expected rents accruing from renting $H_{jt}(k)$ of the housing measure. These firms form subjective expectations $\tilde{\mathbb{E}}_t()$, which may differ from rational expectations $\mathbb{E}_t()$. Each unit of $H_{jt}(k)$ costs p_{jt} in a lump-sum payment to purchase from another property manager or the developer.[33] Their total expected profits sum over

[33] We abstract from leverage or other financial markets by assuming that they self-finance their capital investment into the property at zero opportunity cost.

their holdings on all plots:

$$\tilde{\mathbb{E}}_t(\pi_t(k)) = \sum_j \tilde{\mathbb{E}}_t(\pi_{jt}(k))$$

The separation of the owners of housing from those who occupy it creates a helpful accounting separation between the agent who holds the asset value of the home and the agent with preferences over attributes, but this nests the more common case in the data of owner-occupiers (65% of residents in Mexico City), which would correspond to a property management firm that rents to themselves.[34]

4.2.1 Frictions affecting evaluation of future risk

Expectations over future sinking play an important role in (6); however, frictions may complicate the valuation of yet-unrealized damages, especially in the case of atomistic home buyers who cannot hedge over a sophisticated real estate portfolio (Bakkensen and Barrage, 2022; Gourevitch et al., 2023; Gallagher, 2014). We incorporate the potential for systematic mis-valuation of future risk by modeling expectations of future sinking rates as a function of rational expectations such that:

$$\tilde{\mathbb{E}}_t(s_{ju}) = \theta \mathbb{E}_t(s_{ju}), \qquad u > t \tag{7}$$

In this expression, θ parameterizes the extent of frictions or myopic behavior on forming expectations about future sinking; when $\theta = 1$, property managers have rational expectations about future sinking. As $\theta \to 0$, buyers shade their expectations about future sinking and prices become less responsive to expected future sinking, remaining inefficiently high as property managers overpay for homes that will sink in the future.[35] This parameter is quite

[34] In this case, owner-occupiers would face no frictions to moving (as renters do not face these frictions), but they do incur asset losses arising from the under-valuation of the impacts from future sinking. That is, an owner-occupier must form expectations about future sinking even if they may move frequently because they internalize the the price they will get for the home when they sell.

[35] In this model, θ does not update with realized sinking, but the total welfare impacts of these frictions are the same in an alternative model in which $\theta < 1$ only for the period that the home is first sold as a new unit, which changes only which property developers experience negative profits. We model that θ is

general, and captures the net impact of factors such as information frictions, myopy, or other behavioral mechanisms that lead to systematic under- or over-estimation of future risks.

4.3 The housing developer

In each period the developer receives plot-specific cost shocks ξ_{jt} and decides which plots to purchase and redevelop. The developer operates only within Mexico City. Purchasing a plot means purchasing the full measure of housing contained on it, H_{jt-1}. Developing involves building new units which have zero realized sinking; the key features of development are that it allows the developer to adjust the housing stock on the plot and it "resets" realized sinking to zero. Their plot-specific profits conditional on developing are:

$$\pi_{jt} = \max_{H_{jt}} \underbrace{p_{jt}(0)H_{jt} - \xi_{jt}H_{jt}^{\frac{1+\eta}{\eta}}}_{\text{Variable profits from developing}} - \underbrace{p_{jt}H_{jt-1}D}_{\text{Cost of acquiring, renewing plot}}, \qquad \eta > 0 \tag{8}$$

where η is the supply elasticity, and $p_{jt}(0)$ is shorthand for the price that a developer receives for new units that have experienced no sinking. D is renovation cost of the plot, which is proportional to the total value of the plot.[36] Because the developer competes with property managers (the used housing stock) both when selling their units and when purchasing plots, they take market prices as given but may make positive profits arising from favorable productivity shocks.

Conditional on developing, the developer chooses how many units to build optimally.

uniform among property managers, but note insights from papers such as Bakkensen et al. (2019) that show that only some mis-informed participants are needed to inflate prices. This behaves similarly to the β in the beta-delta behavioral model (Laibson, 1997; DellaVigna, 2018), except that the discounting is only over future sinking and not all future rents.

[36] This allows demolition costs to increase with buildings that are built with better materials or that are larger, qualities that would be correlated with higher value.

Solving this problem gives:

$$H_{jt}^{opt} = \left(\frac{\eta}{(1+\eta)\xi_{jt}} p_{jt}(0)\right)^{\eta}$$
$$\pi_{jt}^{*} = \max\{0, \tilde{\eta}\xi_{jt}^{-\eta} p_{jt}(0)^{\eta+1} - p_{jt} H_{jt-1} D\}, \qquad \tilde{\eta} = \left(\frac{\eta}{1+\eta}\right)^{\eta}\left(\frac{1}{1+\eta}\right)$$

The developer will re-develop a plot if variable profits are larger than the cost of acquiring the plot:

$$d_{jt} = 1\left(\tilde{\eta}\xi_{jt}^{-\eta} p_{jt}(0)^{\eta+1} \geq p_{jt} H_{jt-1} D\right) = 1\left(\tilde{V}_{jt}^{S} \geq 0\right)$$
$$\tilde{V}_{jt}^{S} = \log(\tilde{\eta}) + (\eta+1)\log(p_{jt}(0)) - \log(p_{jt}) - \log(H_{jt-1}) - \log(D) + \tilde{\xi}_{jt}$$

where $\tilde{\xi}_{jt} = -\eta \log(\xi_{jt}^{S})$, and we assume $\tilde{\xi}_{jt} \sim \mathcal{N}(\mu, \sigma^2)$ for calculation of expected surplus and estimation. Intuitively, higher prices for new units increase $\tilde{V}_{jt}^{S}$ by $\eta + 1$ because this latent value is a function of log revenues (as opposed to quantities). We transform this expression into a function of observable data using the price decomposition given by the demand structure to separate realized subsidence from the rest of the price: $\log(p_{jt}(0)) = \log(p_{jt}) - \gamma S_{jt}$:

$$d_{jt} = 1\left(\tilde{V}_{jt}^{S} \geq 0\right) \tag{9}$$
$$\tilde{V}_{jt}^{S} = \log(\tilde{\eta}) + \eta \log(p_{jt}) - \gamma(1+\eta) S_{jt} - \log(H_{jt-1}) - \log(D) + \tilde{\xi}_{jt} \tag{10}$$

and write housing supply as:

$$H_{jt}^{S} = d_{jt} H_{jt}^{opt} + (1 - d_{jt}) H_{jt-1} \qquad \forall\, j \tag{11}$$

4.4 Equilibrium

Given exogenous parameters $\{\gamma, \theta, \eta, \overline{L}, \{H_{j0}\}_j, \delta, \rho, D\}$ and realizations of cost shocks $\{\xi_{jt}\}_{j>0}$, the equilibrium consists of a set of prices $\{p^*_{jt}\}_j$, rents $\{R^*_{jt}\}_j$, and housing $\{H^*_{jt}\}_j$ that satisfy the following conditions in each period:

- Renters choose the plot to rent on optimally. Property managers and the developer maximize profits.
- Property management firms make zero subjectively-expected profits:[37]

$$p^*_{jt} = \tilde{\mathbb{E}}_t \left(\sum_{\tau=t}^{\infty} \rho^{\tau-t} R^*_{j\tau} \right) \qquad \forall\, j \tag{12}$$

- The rental market clears on rents R^*_{jt} in each period such that rental demand given by (5) equals (11).[38]
- The market for new homes clears on home prices p^*_{jt} according to (12) and (11) given the price function $\log(p^*_{jt}(0)) = \log(p^*_{jt}) - \gamma S_{jt}$.[39]

Figure 8 visualizes these two markets. Demand for owning housing inherits its shape from the rental market, and may be too high if pricing frictions inflate the value of future rents. If re-development happens, then supply is upward-sloping from the developer's problem; otherwise it is inelastically set at the previous housing quantity. Once the housing stock has been set by this equilibrium, the rental market takes this as given and clears on rents so that demand for a plot equals the supply of housing on that plot. As sinking occurs, rental prices fall on plots that have not been redeveloped to maintain all housing on the plot occupied.

[37] This follows because these firms are competitive and we assume that there is no private information or other advantage that firms can leverage.

[38] Given that (i) development decisions and the housing stock on each plot are determined before the rental market occurs, and (ii) conditional on ownership, property managers supply rentals at zero marginal cost, rental supply is inelastically supplied according to (11).

[39] Because there is only one developer competing with many property managers to purchase units, the transaction price will be set by property manager demand. This also implies that the expectation that property managers form of the future stream of rents will be set by the value of those rents conditional on *not* redeveloping, as that is the price a developer will pay them for the asset even if the plot is subsequently redeveloped.

Deadweight loss in the housing market comes from the over-supply of units in sinking areas that result in losses for property managers as the flow of rents is lower than expected on those units.

4.5 Welfare and Model predictions

The total change in welfare in the housing and rental markets in Mexico City between counterfactuals corresponds to the total change in developer, property manager, and renter surplus between alterative equilibria.[40] Per-period developer's surplus is their total expected profits across all plots, taking expectations over the distribution of cost shocks:

$$PS_t = \sum_{j>0} \mathbb{E}_\xi(\pi^*_{jt}) \tag{13}$$

Objectively-expected property managers' profits are derived by substituting (6) into (12):

$$\Pi_t = \sum_j H^*_{jt} \left(\sum_{\tau=t}^{\infty} \rho^{\tau-t} \left(\mathbb{E}_t(R^*_{j\tau}) - \tilde{\mathbb{E}}_t(R^*_{j\tau}) \right) \right) \tag{14}$$

Under rational expectations, these are zero; in the presence of information frictions (rationally) expected profits will be negative.

To account for using log rents in utility, we calculate the equivalent variation Δ^D_t that equalizes average consumer surplus between counterfactuals to characterize the change in renter welfare. The present value of the difference in total welfare between counterfactuals is:

$$\Delta W = \Delta\Pi_t + \frac{\Delta^D_t}{1-\rho} + \frac{\Delta PS_t}{1-\rho} \tag{15}$$

[40] Derivations of the welfare terms and comparative statics in the next section can be found in Appendix G.4.

4.5.1 Model predictions

Examining how model outcomes and welfare behave in equilibrium sheds light on how the housing market responds to subsidence and information frictions; we list the main propositions here, and provide proofs in Appendix G.4. For each experiment, we consider the "long-run" difference between equilibria under different parameter values, allowing prices and the housing stock to fully adjust.

Proposition 1: Pricing frictions leading to under-valuation of future sinking ($\theta \to 0$) increase expected housing supply and producer surplus on sinking plots, but have no effect on the probability of re-development.

Proposition 2: Realized sinking increases expected producer surplus and the probability of re-development, but expected future sinking decreases expected producer surplus and has no impact on the probability of redevelopment.

Proposition 3: The profits of property managers become more negative as $\theta \to 0$.

Intuitively, under-valuing future sinking leads to an inefficiently high price being offered for new units in sinking areas, which implies an oversupply of housing conditional on developing and increased profits. However, because pricing frictions inflate housing density in sinking areas over the long-run, both the developer's fixed cost of buying the plot (which is proportional to housing density) and revenue (a function of expected future sinking) are too high under information frictions, so re-development happens with the same probability regardless of the value of θ.

Expected future sinking has an analogous though opposite effect: future sinking also decreases both cost and revenues equally and therefore has no net effect on redevelopment, but it decreases prices and therefore the number of units built conditional on developing. Realized sinking works through a different mechanism; by directly lowering the costs of buying up a plot, realized sinking increases the probability of redevelopment by increasing the probability of positive profits.[41] Property managers make negative profits under pricing

[41]The impact of these policies on renters is not a direct function of sinking, as they are perfectly

frictions, as they over-pay for units that will sink more than they expect.

In this model three parameters emerge as crucial for quantifying the total costs of subsidence: γ, the per-period direct costs of subsidence, θ, the extent of pricing frictions, and η, which governs the shape of the supply curve and thus the responsiveness of supply to the market distortion. In the next sections, we derive estimating equations from the model for these parameters and estimate them using our data on mortgage appraisals and development.

5 Estimation of model parameters

5.1 The willingness to pay to avoid subsidence

Our granular data permits a direct approach to estimating the willingness to pay to avoid subsidence and capitalization of future risk that leverages within-building variation to isolate plausibly exogenous variation in sinking.[42] We solve Equation (4) for rents:

$$R_{jt} = \exp\{\gamma S_{jt} + \delta Z_{jt} + e_{jt} - \tilde{\lambda}_{jt}\}$$

where $\tilde{\lambda}_{jt} = \log(\lambda_{jt}) - \log(\lambda_{0t})$, and substitute this expression into (12) to derive a relationship between log prices and sinking:

$$\log(p_{jt}) = \gamma S_{jt} + \log\left(\exp(-\tilde{\lambda}_{jt} + Z_{jt}\delta + e_{jt}) + \tilde{\mathbb{E}}_t(\sum_{\tau=t+1}^{\infty} \rho^{\tau-t} \exp\{\gamma S_{j\tau}^{t+1} + Z_{j\tau}\delta + e_{j\tau} - \tilde{\lambda}_{j\tau}\})\right)$$

where $S_{j\tau}^{t+1} = \sum_{u=t+1}^{\tau} s_{ju}$ is the cumulative sinking from $t+1$ to τ.[43]

compensated for sinking through decreased rents. However, they may experience benefits in alternative equilibria where rents are lower on average due to a higher housing stock, and vice versa.

[42] Our approach also allows us to recognize that sinking is an individual property attribute in our appraisals data. We also note the equivalence of hedonic and discrete choice estimators when the characteristic is continuously "supplied" in the market, as is the case of subsidence in our setting (Bayer et al., 2007).

[43] We are able to factor out realized sinking from the expression because (i) its impact on homes is persistent, and (ii) because the transaction price is set by property manager's demand, expected future rents are those of the non-redeveloped property.

Stronger assumptions are required to make progress towards an estimating equation that teases expected sinking from the residual. We first assume that $Z_{j\tau}$ are time-invariant and can be written as Z_j. Next, we take two first-order approximations: first of $\mathbb{E}_t(\exp(\gamma S_{j\tau}^{t+1} + e_{j\tau} - \log(\lambda_{j\tau}) + \log(\lambda_{0\tau})))$ around the expectations of these variables, and then of the log function around $\mathbb{E}_t(e_{j\tau}) = e_{jt}$, $\mathbb{E}_t(\tilde{\lambda}_{j\tau}) = \tilde{\lambda}_{jt}$, and $\mathbb{E}_t(S_{j\tau}^{t+1}) = 0$ to arrive at:[44]

$$\log(p_{jt}) \approx \gamma S_{jt} + \gamma\theta \sum_{\tau=t+1}^{\infty} \rho^{\tau-t}\mathbb{E}_t(s_{j\tau}) + \delta Z_j - \log(H_{jt}^S) + \log(H_{0t}^S) + u_{jt}, \tag{16}$$

$$u_{jt} = e_{jt} + (1-\rho)\sum_{\tau=t+1}^{\infty} \rho^{\tau-t}\left(\mathbb{E}_t(e_{j\tau}) - e_{jt} + \tilde{\lambda}_{jt} - \mathbb{E}_t(\tilde{\lambda}_{j\tau})\right) - \log(1-\rho)$$

and formulate the estimating equation we take to the data as:

$$\log(p_{jt}) = \gamma S_{jt} + \gamma\theta \sum_{\tau=t+1}^{\infty} \rho^{\tau-t}\mathbb{E}_t(s_{j\tau}) + \alpha_j + \alpha_{age} + \alpha_{z(j)t} + Z_j \cdot t\beta + \tilde{u}_{jt} \tag{17}$$

By including fixed effects for identical units in the same building α_j, we flexibly control for both observed and unobserved static quality differences between appraised units (Z_j, e_j), a concern that is first-order in many hedonic estimation strategies. The α_j also absorb $\log(H_{jt}^S)$, as housing supply on the plot is fixed between appraisals. Age fixed effects α_{age} control for the fact that realized sinking will be mechanically higher on older properties, and geography by year fixed effects $\alpha_{z(j)t}$ absorb housing supply outside Mexico City ($\log(H_{0t}^S)$), and allow demand shocks to vary flexibly across housing markets.[45] The inclusion of linear trends by a vector of building characteristics $Z_j \cdot t$ approximates the $\mathbb{E}_t(e_{j\tau}) - e_{jt}$ terms by allowing housing of different types to experience different market trends.[46]

[44] Appendix G.1 provides more details and derivations. The second approximation is as if homebuyers expect demand shocks and relative rental demand to follow a random walk. In the case of the $\tilde{\lambda}_{jt}$ this approximation comes at a cost in terms of model consistency, as these relative choice probabilities are equilibrium objects that have a distribution that is not likely to be a random walk. However, this approximation is chosen to reflect the fact that $\tilde{\lambda}_{j\tau}$ is fixed until re-development; because re-development is rare, with a per-period probability of 2.7%, $\mathbb{E}_t(\tilde{\lambda}_{j\tau}) = Pr(d_{jt} = 0)\lambda_{jt} + Pr(d_{jt} = 1)\mathbb{E}_t(\tilde{\lambda}_{j\tau}|d_{jt} = 1) \approx \lambda_{jt}$.

[45] We show robustness to specifications that define $z(j)$ as the borough and seismic zone. There are 16 boroughs and three seismic zones in Mexico City which delineate seismic risk and construction code requirements.

[46] Our most stringent specification includes linear trends that vary by the borough and type (single family,

Identification in estimating (17) requires that the expected amenities and rental demand in the residual be independent of sinking and expected sinking conditional on the controls. We argue that sinking is quasi-randomly assigned conditional on these controls because local forces cannot directly influence localized sinking, and we conduct tests for parallel trends and placebo tests to validate this argument.

First, we note that property-level fixed effects break up the strong cross-sectional correlation in both sinking and prices, as can be seen in a map of residualized sinking and prices in Figure 9, so our identifying variation is not broadly related to spatial trends affecting housing markets in different parts of the city. Next, we test for pre-trends using the microdata from the Income and Expenditure Survey (ENIGH) test for whether future sinking predicts current rents. Rents need not be forward looking, so if they predict future sinking conditional on current sinking that would provide evidence that local housing market shocks are correlated with subsidence. Table 2 reports coefficients of the impact of sinking and the present value of future sinking on rents, where the present value of future sinking aggregates observed sinking from the survey wave to 2020 at the location of the respondent and an estimate for sinking beyond 2020. Future sinking has no predictive power for log rents, providing reassurance that property values are not evolving systematically before sinking occurs.

We also test for parallel trends by examining the relationship of the composition of the housing market with subsidence. While our approach directly controls for quality, a correlation of the characteristics of units on the market with subsidence could be an indication that the market is evolving in other unobserved ways that confound our estimates. Figure 10 reports the coefficients and 95% confidence intervals from regressions of the Z-scores of unit amenities on sinking on the property. Sinking is uncorrelated with these amenities, and all point estimates imply a treatment effect less than 0.025 standard deviations in size. For the hedonic index that aggregates all of these, the (insignificant) treatment effect of a

apartment, etc.), but we show that our results are robust to using subsets of these controls.

meter of subsidence is less than 0.01 standard deviations. These results are reassuring that the housing markets that sinking plots are located in are not becoming systematically lower quality.

Thus far we have found no evidence that housing markets in sinking areas have confounding trends that are correlated with sinking. In Section 5.3.1 we discuss other pre-trends tests and conduct placebo and other robustness exercises, but now proceed with the presentation of our results estimating the costs of subsidence.

5.1.1 Results

Table 3 reports the results of estimating Equation (17). Standard errors are clustered at the census tract. We formulate expectations as both an AR(1) process estimated at the plot level in Columns (1) and (2), and as a "perfect foresight" measure that uses observed sinking on the plot post-transaction in expectations in Columns (3) and (4). Regardless of the manner of specifying expectations or the discount rate chosen, realized sinking has a consistent and negative impact of appraised values, with a meter of sinking causing a 6% fall in appraised value. Given that the average property in our sample sank 25 centimeters over 10 years, this implies a 1.5% average decrease in values per decade.

We do not detect a role for future sinking in impacting prices in any specification, and in all cases estimate a value of θ that is statistically indistinguishable from zero. However, these estimates will be biased towards zero if we mis-specify the formulation of expectations or if the specification directly absorbs expectations (Dickstein and Morales, 2018). For example, we would expect to estimate $\widehat{\gamma\theta} = 0$ is the true data generating process is that of expectations that model the sinking rate is a constant on the plot, since this would be absorbed by our fixed effects.

We address this by relaxing our spatial controls in Table 4, which shows the results of estimating (17) only for new builds, replacing our property fixed effects with 500-meter grid cell fixed effects which permit more systematic variation in future expectations to enter

estimates. We again find that expected future sinking has no statistically detectable or economically meaningful impact on appraised values in both ways of specifying expectations. However, given that we are limited in our ability to mitigate the attenuation bias that may affect these estimates, we complement these findings with qualitative survey evidence that households struggle to draw accurate inference about future subsidence risk.

5.1.2 Discussion: Potential sources of pricing frictions

What drives the lack of capitalization of future sinking? Our survey evidence in Section 2.3 points to substantial information frictions affecting the internalization of future risk and sorting on these frictions on sinking risk. Consistent with this idea, we find that capitalization of future sinking into the price of new builds is higher in boroughs with slower sinking rates on average. Appendix Figure B5 plots estimates of $\gamma\theta$ by borough, showing that these coefficients are more likely to be negative and near γ in low-sinking boroughs.

Our survey evidence showed that differential discounting or risk preference are unlikely explanations for this lack of pricing. We can also rule out that future sinking is not priced because some residents do not value it as we find that prices do respond to past sinking, meaning that sinking is generally valued in the market. Finally, we note that Mexico lacks large-scale programs such as subsidized insurance or bailouts that have been found to be important drivers of under-valuation in other settings (Bézy, 2025; Wagner, 2022). While we cannot exclusively attribute this lack of capitalization to any one friction, our evidence suggests that information is an important driver of the mis-pricing we find.

5.1.3 Mechanisms: How does realized sinking affect homes?

Our findings in Table 3 show that plot-level sinking imposes important losses to home values, but what is driving the losses from realized sinking? Sinking damages structures, but it is possible that homeowners are able to invest in maintenance and other remedies that offset these damages. We explore these mechanisms in Table 5, where we simplify estimation by

only including realized subsidence. Column (1) reports the result of estimating our repeat appraisal specification, but replacing the outcome with the appraised years of life remaining for the home.[47] This measure captures the appraiser's estimate of the years of inhabitable life left for the structure, and provides an estimate of how degraded or depreciated the physical building is. We find that sinking has a pronounced negative impact on the appraised life remaining of the home: a meter of sinking causes a reduction of 4.4 years of remaining life, and the average home lost 1.1 years of life at the average 10-year sinking of 25 centimeters compared to a home that experienced no sinking.

Our survey evidence on the relationship between the physical state of homes and sinking also provides correlational evidence that sinking depreciates structures. Columns (2) through (5) regress indicators for whether a respondent reports different structural issues in their home that could be caused by subsidence on our satellite measure of the total sinking on the survey respondent's home. Column (2) shows that a meter of sinking on the respondent's home increases the probability of reporting severe structural issues by 17% relative to the mean in the low-sinking strata and increases the probability of reporting issues with levelling and cracks in walls by 20% and 28%, respectively.

Do residents invest in maintenance in response to these damages? Table 6 reports the correlation of survey measures of maintenance investments with sinking on the respondent's plot. While we do not find that sinking is associated with extensive margin increases on the probability of incurring maintenance expenses, we do find that a meter of sinking is correlated with a large increase in spending on maintenance by 44% conditional on incurring any spending. Sinking is also associated with an 18.8% fall in the respondent's assessment of when they will have to make their next "large investment" in the home, suggesting that residents are pulling forward investments in the home.[48]

Finally, we look for evidence that subsidence induces complementary investments in

[47] Appendix Table A5 reports the result of all specifications for this outcome.

[48] We find a similar result using spending on maintenance from the income and expenditure survey in Appendix Table A6, though in these data we are only able to detect extensive margin increases in the probability of incurring maintenance costs.

maintenance or home improvement activities that are not related to subsidence, as these could bias our estimates towards zero if they increase the market value of the home. We test for this in the last column of Table 6 by regressing an indicator for the respondent incurring any spending on non-sinking related issues (additions, remodeling, electric work, painting, or waterproofing) on sinking on the home. We find no statistically significant or economically meaningful relationship between sinking and complementary investments, assuaging concerns that these cause us to underestimate the direct costs of subsidence.

5.1.4 Mechanisms: Uneven sinking across neighborhoods

So far we have focused on the impact of plot-level sinking on home values, but uneven sinking across neighborhoods also has the potential to reduce the desirability of a property through its impact on public infrastructure. We study this by introducing our measures of the unevenness of sinking over space- the cumulative relative sinking of the grid cell relative to its neighbors and the standard deviation of sinking- into our main specifications.

Table 7 reports results. Columns (1) and (2) report the result of estimating the repeat appraisal design on plot-level sinking and our measures of uneven sinking. Column (1) reports results from including the cumulative relative altitude change across 500m grid cells and Column (2) includes the standard deviation of sinking within the 500m grid cell. Both measures have a negative impact on price, but changes in the relative altitude have a larger effect on values: The average change in relative altitude across grid cells causes a 1% fall in appraised values. The impact of locally uneven sinking, as measured by the standard deviation *within* the grid cell, is much smaller, with the average change in the standard deviation of sinking causing a 4.7% fall in appraised values.

In Columns (3) through (6) we use survey reports of infrastructure issues in the neighborhood to estimate the physical impacts of uneven sinking on the probability that a neighborhood experiences flooding or fissures in pavement or roads. First, we see that plot-level sinking has no significant relationship with the probability of these infrastructure

issues, underscoring the intuition from civil engineering studies (Hu et al., 2013) and the Mexico City government's own qualitative reports (Secretaría de Gestión Integral de Riesgos y Protección Civil, 2018) that the damages caused by subsidence on macro-infrastructure are driven by uneven sinking across neighborhoods, not just localized sinking.

Column (3) shows that a respondent is 2.85% more likely to report that their neighborhood has issues with flooding at the average relative altitude change of 2.8 centimeters. By contrast, the standard deviation of sinking within the grid cell is not predictive of this outcome (Column (4)). Fissures in roads or sidewalks show the opposite pattern: relative altitude change does not predict fissures, but mean unevenness as measured by the standard deviation of sinking within the cell is correlated with a 0.5% increase in the probability that a respondent reports issues with fissures. These findings suggest that which of unevenness matters for infrastructure depends on the which types of infrastructure one is interested in; while relative altitude change, which measures the change in relative slopes, affects the probability of flooding, it seems that local unevenness, as measured by the standard deviation, is more predictive of fissures in roads.

Summing up We have found that subsidence has large direct costs and that home values do not show evidence of strong capitalization of future risk. For welfare analysis, we set conservative values for our demand parameters, setting the costs of subsidence to $\gamma = -0.06$ and the extent of information frictions $\theta = 0.08$, the highest value estimated in any specification (Column (4) of Table 4). In the next section we leverage the fact that uneven sinking across grid cells also shifts demand, using this variable as an instrument for prices in supply estimation.

5.2 Supply response to subsidence

The supply model highlights the role that realized subsidence plays in increasing the probability of development by lowering the opportunity cost of redeveloping existing,

degraded units. In this section, we exploit this relationship to estimate the supply elasticity using the responsiveness of re-development to realized sinking and prices.

5.2.1 Reduced-form evidence for model predictions

We begin by testing our model prediction that realized sinking draws in development in two ways. First, we use our matched development-appraisal data to test whether homes that have sunk more are more likely to be sold for redevelopment. Table 8 reports the results of regressing of an indicator for an appraisal being redeveloped on the sinking it has experienced to date and a battery of property characteristics, including the age of the building. Strikingly, Columns (2) and (3) show that sinking strongly predicts whether an appraisal is for the purpose of re-development, with a meter of sinking increasing this probability by 0.5 percentage points over a mean probability of 0.89%. Consistent with irregular neighborhood sinking making neighborhoods less attractive, changes in relative altitude decrease the probability of redevelopment conditional on plot-level sinking, though noisily so. Including the appraised price as a control in Column (3) does not change these conclusions.

As an alternative test, we use our panel of pixels to estimate the impact of sinking on the probability of that construction begins on a new development in pixel j in year t. This approach allows for less detailed analysis of the exact characteristics of the unit, but does not require us to rely on accurate identification of transactions that were later redeveloped. Table 9 reports results for varying ways of controlling for time trends, and our most flexible specification in Column (3) implies that a meter of sinking on a pixel from 2007-t increases the probability that construction will begin on a new development in that pixel by 8.2%.

5.2.2 Estimating the supply elasticity

We estimate the supply elasticity using the model-derived redevelopment decision given by (9) and (10), which tracks the response of development to subsidence. However, estimation

is challenged by the endogeneity of cost shocks $\tilde{\xi}_{jt}$ with prices. Specifically, supply shocks cause movement along the demand curve in equilibrium, which biases elasticity estimates that ignore this effect toward zero. Our solution is a control function approach. Following Petrin and Train (2010), we impose additional structure on the distribution of cost shocks and identify a demand shifter as an instrument for prices, which allows us to derive a control function that absorbs the endogenous part of the error term.[49] We propose the following first stage equation derived from our demand model:

$$\log(p_{jt}) = \Pi \text{Rel. Altitude Change}_{jt} + \gamma S_{jt} + \alpha_{n(j)} + \alpha_t + X_{jt}\beta_1 + \tilde{u}_{jt} \tag{18}$$

where Rel. Altitude Change$_{jt}$ is the relative altitude change of the grid cell containing j from 2007 to t, $\alpha_{n(j)}$ are census tract $n(j)$ fixed effects, α_t are year and calendar month fixed effects, and X_{jt} are appraisal controls including age, type, number of bedrooms, and number of bathrooms. We assume that $\tilde{\xi}_{jt}$ is distributed joint normally with log prices, with constant variance σ and marginal expectation $\mathbb{E}(\tilde{\xi}_{jt}|\tilde{u}_{jt}) = CF(\tilde{u}_{jt};\phi)$ for some parameter ϕ. Then we can re-write our redevelopment condition in (10) by adding and subtracting the control function and adding the controls from the first stage:

$$\begin{aligned} d_{jt} &= 1\left(\frac{1}{\sigma^*}\tilde{V}^S_{jt} \geq 0\right) \\ \frac{1}{\sigma^*}\tilde{V}^S_{jt} &= \frac{\eta}{\sigma^*}\log(p_{jt}) - \frac{\gamma(\eta+1)}{\sigma^*}S_{jt} + \frac{1}{\sigma^*}CF(\tilde{u}_{jt};\phi) + \tilde{\alpha}_{n(j)} + \tilde{\alpha}_t + X_{jt}\beta_2 + \frac{\tilde{\xi}^*_{jt} - \mu^*}{\sigma^*} \end{aligned} \tag{19}$$

where our distributional assumptions imply that $\tilde{\xi}^*_{jt} = \tilde{\xi}_{jt} - \mathbb{E}(\tilde{\xi}_{jt}|\tilde{u}_{jt})$ is normally distributed with mean μ^* and standard deviation σ^*. Thus this is estimated via Probit, including the control function as a regressor.[50] Because of the assumption of joint normality, the control

[49] Control functions are numerically equivalent to IV estimation in the linear case, and can be extended to non-linear estimation under certain conditions. Appendix G.5 provides a proof of the validity of using a control function in this context (Blundell and Matzkin, 2014).

[50] We prove that the new error term $\frac{\tilde{\xi}^*_{jt} - \mu^*}{\sigma^*}$ is mean independent of prices given a valid instrument in Appendix G.5.

function is a linear function of the first stage residuals, i.e. $CF(\tilde{u}_{jt};\phi) = \phi\widehat{\tilde{u}_{jt}}$, and we show robustness to mis-specification by showing that our results are unchanged when including higher order polynomial approximations of this function.

To be considered a valid demand shifter, uneven grid sinking must satisfy the exclusion restriction: it must influence redevelopment solely through its impact on prices. This means the sinking cannot have a direct effect on development costs, conditional on the census tract and other controls we include. Evidence from structured interviews with civil engineers working for construction firms in Mexico City corroborates this assumption. Engineers confirm that construction code regulations mandate plot-specific soil mechanics tests, which are not designed to predict future subsidence or account for sinking in adjacent locations.[51] The inclusion of census tract fixed effects further controls for unobserved, location-specific supply factors, such as underlying soil conditions and local construction code requirements which vary by seismic zone. Finally, we provide empirical evidence in Section 5.3.2 that trends in uneven sinking are uncorrelated with other relevant supply trends.

Table 10 reports the results of estimating (19). Confidence intervals are estimated via clustered bootstrap to account for the fact that the control function is estimated from the first stage. Column (1) reports the first stage results from regressing log prices on uneven grid cell sinking and controls. Uneven sinking lowers prices conditional on tract fixed effects and unit-level controls, and with an F-statistic of 10.6 the first stage is sufficiently powered.

In line with model predictions, both linear IV (Column (2)) and Probit estimation with the control function (Column (3); average marginal effects reported in Column (4)) find that redevelopment responds positively to prices and realized sinking. While the magnitude of the coefficients in Column (3) is not interpretable due to the unidentified scale parameter, Equation (10) implies that we can estimate the elasticity as $\eta = \widehat{\varepsilon}/(1-\widehat{\varepsilon})$, where $\widehat{\varepsilon} =$

[51] The purpose of soil mechanics tests is to ensure that the structure itself does not induce sinking, and thus these report static characteristics according to the geology of the plot. Interviews with civil engineers operating in Mexico City suggest that these tests provide limited information about how much sinking has occurred beyond the static risk implied by the geology, and do not make predictions about how much future sinking will occur unrelated to the structure's impact.

$-\beta^p/\beta_S \times \gamma$,[52] which implies a supply elasticity of 0.28.

This estimate is robust to including a more flexible polynomial approximation for the control function in Column (5). We see that the control function is crucial for addressing bias by comparing these results to estimates that exclude the control function in Column (6), which implies an elasticity near zero. Our estimate of η =0.28 is in line with recent papers that measure within-city housing supply elasticities for the United States (Baum-Snow and Han, 2024; Rollet, 2025).

5.3 Validation and robustness

5.3.1 Estimates of the cost of subsidence

Our estimates of the costs of subsidence are robust to using subsets of our fixed effects and linear trends. Table A7 in the Appendix reports estimates of the effect of subsidence on appraisal values for different relaxations of these controls, and we estimate similar magnitude effects in all cases. Results are also robust to modeling potential measurement error in subsidence; Figure B6 plots the distribution of p-values from a Monte Carlo simulation that models noise in the calculation of the subsidence rate distributed normal with mean zero and standard deviation of 10 millimeters, a conservative estimate of the measurement error for SAR Interferometry in other settings (Liu et al., 2020; Bawden et al., 2023), showing that accounting for this uncertainty does not change results.

We next conduct a placebo test by testing for whether the subsidence of a plot's neighboring pixels affects prices. This serves as a test for whether the increased probability that my neighbors are re-developed introduces a confounder, as sinking neighborhoods could experience gentrification induced by new builds. In order to address the fact that we cannot observe the ages of neighboring properties, we run an analogous specification to (17) but in differences, as we can observe the change in sinking for neighbors. Table A8 in the

[52] Confidence intervals are calculated over estimates of η in each bootstrap sample, and in each sample γ_b is drawn from a normal distribution centered around the point estimate with standard deviation equal to the standard error of this estimate.

Appendix shows that sinking on neighboring plots is not correlated with housing values, giving confidence that we are measuring the impact of damage to the home itself, not other trends in the neighborhood.

Subsidence and seismic risk are very correlated in the cross section. If housing market trends related to seismic concerns are evolving in way that is correlated with shocks to subsidence, then our estimates may pick up these trends. To test for this, we estimate Equation (17) excluding years after 2017, as the major earthquake that hit Mexico City in 2017 may have impacted the evolution of the housing market in ways that are correlated with sinking. Table A9 in the Appendix reports results, which if anything are stronger when limiting to the pre-2017 period, reflecting that our estimates pick up changes in sinking, not broad trends associated with the Eastern part of the city.

Our main pre-trends test looked for evidence that future sinking was correlated with rents, but we can also use our price data directly to test if changes in home prices predict changes in sinking. Figure B7 compares the relationship between price changes between appraisals and sinking between the appraisals (blue), and price changes and future sinking (pink). While sinking and price changes over the same period are negatively correlated, there is no significant relationship between lagged price changes and sinking between the next set of appraisals. Comparing coefficients on a regression of these, lagged prices are an insignificant predictor of sinking, and the point estimate on lagged prices is 66% that of the point estimate on contemporaneous sinking.[53]

We verify that results are not driven by mechanical appraiser valuation of subsidence by dropping localities where the median difference between appraised and transaction values in the R04 data is greater than 1%. Table A10 reports results, which are quantitatively and qualitatively very similar to our main estimates.

The significance of our results is robust alternative ways of estimating standard errors.

[53]Note that if home buyers use past sinking as a signal of future sinking even conditional on the home's location, then this test would not be model-consistent as forward-looking home buyers should incorporate their expectation of future sinking into prices today.

Table A11 reports the results from our different specifications using a 2-kilometer radius for calculating Conley standard errors. Estimates are generally more precise when using Conley standard errors, as differencing within property reduces the strong spatial correlation in outcomes.

Finally, we test for whether changes in appraised values are driven by mortgage penalties on sinking homes. If banks charge higher interest rates on homes that are prone to sinking, then the changes in values we estimate may represent discounts that the market offers to offset these rates. Table A12 in the Appendix estimates the cross-sectional correlation of the sinking intensity of the census tract with the interest rate charged on mortgages on that tract using the R04 data, finding no evidence that banks assign differential interest rates to high sinking areas.[54]

5.3.2 Supply estimation

We investigate whether new developments in sinking areas are systematically different from developments in non-sinking areas. This test addresses the possibility that developers may adapt to sinking— for example, by building at lower heights or including different amenities— choices that could directly affect their costs. However, Tables A13 and A14 in the Appendix show that we cannot reject the hypothesis that developments in sinking areas are built with similar heights and amenities as those in other areas; in terms of observables, developments appear similar regardless of past and future sinking.

Identification in our supply estimation requires that redevelopment trends similarly across areas, both in terms of plots experiencing high versus low sinking and in neighborhoods with uneven versus uniform sinking. To test this assumption, we compare pixel-level development trends between areas with high and low sinking intensity. Intuitively, this comparison focuses on comparing the trends of pixels that were never sinking quickly with those that started

[54] A limitation of this test is that we cannot observe mortgages that were not approved by the bank, so there may be extensive margin mechanisms through which the mortgage market addresses subsidence that we cannot identify.

sinking quickly later in the panel. To formalize this, we classify a pixel as "high sinking" using a cut-off based on the 2020 distribution of sinking rates and define the pixel as "treated" starting in the first year its sinking rate reaches that threshold.

Figure B8 shows the results from the Sun and Abraham (2021) event-study estimates, which exclude pixels that were always fast-sinking. The results confirm that the not-yet-sinking pixels were trending similarly to the never-sinking pixels before the sinking accelerated. Crucially, it is only after pixels begin showing higher rates of sinking that the high-sinking pixels demonstrate an increased probability of development compared to the never-sinking group.

6 Counterfactuals and policy implications

6.1 The total cost of sinking, benefits of alleviating frictions

We estimate the costs of subsidence and the benefits of alleviating information frictions by calculating the difference in welfare between the observed equilibrium and a counterfactual equilibrium according to (15). First, we estimate the total costs of subsidence by comparing welfare in the observed equilibrium to one in which all expected future sinking is set to zero ($\mathbb{E}_t(s_{j\tau}) = 0 \, \forall j, \tau > t$). We then turn to the benefits of disclosure laws that reveal the rationally-expected value of future subsidence to property managers through two policy experiments. The first estimates the gains from implementing a policy that would alleviate the pricing friction without addressing subsidence, which we model as a counterfactual in which sinking is unchanged but $\theta = 1$. For example, if information frictions were the only friction affecting the market and disclosure was fully de-biasing, this would be the impact of disclosure. Second, we estimate how costly subsidence would be in an environment of rational expectations by comparing a counterfactual equilibrium with $\theta = 1$ but sinking is

as in baseline to one in which $\theta = 1$ and expected future sinking is zero.[55]

Figure 11 depicts the gains from these counterfactuals, expressed in billions of 2018 USD. The leftmost red bar depicts the welfare benefits from eliminating future subsidence. We estimate that current rates of subsidence– the full stream of which is capitalized into the housing market– cost Mexico City $33 billion USD, which represents about 17.5% of the total value of housing in Mexico City. Almost all of these costs are borne by property managers, who lose out on lower-than-expected rents; these losses represent the "bad investment" losses that home owners make by over-paying for sinking homes that degrade faster than they expected. Subsidence also costs the developer, but only slightly: Expected future sinking depresses the value of the housing stock and therefore their profits, but this has a muted impact when information frictions preclude capitalization of these gains. Renters enjoy small gains as well, as the higher housing stock in the city lowers rents.

The blue bars depict the changes in welfare from alleviating pricing frictions. Eliminating these frictions generates a gain of $18.2 billion USD, which represents 55% of the gains from completely eliminating subsidence. With no information frictions, property managers again gain as they go from making negative profits to zero profits. Housing developers lose out in the no information frictions counterfactual relative to baseline, reflecting the downward adjustment in the inefficiently high housing stock that they benefited from under information frictions. Renters dislike the higher rents they pay when housing supply falls.

Finally, we can compare the total costs of subsidence (red bars) to what the costs of subsidence would be in a counterfactual with no pricing frictions, which is represented by the yellow rightmost bars. We see in this comparison that property managers are indifferent to the change in subsidence, as they make zero profit in both scenarios. The developer gains from the mitigation of subsidence, as they benefit from the increased value of the housing they build. It is also of note that policies that mitigate subsidence and policies that mitigate pricing frictions create very different set of winners and losers over space. Figures B9 and B10

[55] Appendix Table A15 summarizes the parameters needed to calculate welfare and how each is estimated or calibrated, and details on how we find counterfactual equilibria are available in Appendix G.6.

in the Appendix show that mitigating subsidence leads to housing value gains for incumbent property managers in sinking areas and increases in density, whereas mitigating information frictions alone leads to price depreciation in sinking areas.

6.2 Using model estimates to value alternative policies

The previous model counterfactuals provide an estimate of the total cost of subsidence. In this section we use this cost estimate to value alternative policies outside of the model framework that could address groundwater depletion.

6.2.1 The external cost of groundwater pumping

These estimates provide a measure of the total cost of subsidence at current levels; in order to price groundwater pumping as the externality-generating action, we must convert this to dollars per unit of water pumped. We do this using calibrated hydro-physical relationships between aquifer storage and subsidence in the Mexico City Valley; Appendix H provides details. Combining our total cost estimate with the elevation-to-pumped volume conversion implies that subsidence costs the city $0.89 per cubic meter pumped.

Comparing this tax to the water tariffs Mexico City residents pay in Table 11, which are heavily subsidized for most residents, this amounts to an 52-293% increase over the average marginal rate paid by residents, and would add 32-299% to the average total bill.[56] Such increases are likely politically infeasible, and we study the cost-effectiveness of alternative policies in the next section.

6.2.2 Cost-benefit of other recharge policies

If significantly increasing water tariffs is not possible, are there other cost-effective options for abating subsidence? We conduct cost-benefit analysis to value the effectiveness of four

[56] We assume that 60% of the water bill is taxed, the average share of water in Mexico City that comes from groundwater extraction.

policies that have been implemented at some point in Mexico City that would permit offsetting groundwater pumping: investment in wastewater injection, repairing leaks in the water delivery infrastructure, investing in new sources of water, and wastewater recycling plants.[57] These analyses should be taken as suggestive; while policy interest and previous investments in each of these options over the last 30 years in Mexico City provide some idea of potential costs (SACMEX, 2012), current and accurate cost data for these types of investments would improve these estimates.[58]

Figure 12 depicts the results of these analyses, plotting interventions as a merit order function. In this plot we see that a mixed menu of interventions that promote recharge would be optimal for Mexico City, and that implementing all interventions with a benefit to cost ratio above one would imply offsetting almost half of the city's annual groundwater pumping. While wastewater recycling has the potential to offset the largest volume of water, the cost of building and operating these facilities means they are the least cost-effective option, and this is driven by our estimate of how expensive it is to treat grey water. Repairing leaks, on the other hand, saves the least amount of water, but its benefits dramatically outweigh the costs assuming that these repairs do not take away from recharge. Wastewater injection also shows strong promise in terms of cost benefit analysis, and its benefits would be even larger when spatial targeting is taken into account.[59]

[57] We assume that all leaks are repaired on surface water sources *before* entering the basin, so that all water diverted from leaks would not have otherwise contributed to aquifer recharge. This is an important assumption, as CONAGUA (2024) estimate that most groundwater recharge on the Mexico Valley aquifer comes from water infrastructure leaks, so the groundwater impacts of repairing leaks within the city are unclear.

[58] Detailed documentation of the assumptions made and cost sources for each of these are available in Appendix H.1.

[59] We assume that wells evenly affect subsidence across the city in these calculations, but if wells are built near areas that are subsiding quickly then the imperfect connectivity of the aquifer means that this would have a larger local impact on subsidence, potentially increasing the gains even further.

7 Conclusion

We conduct a comprehensive valuation of the costs of subsidence in Mexico City, finding that subsidence costs Mexico City over 33 billion USD, 55% of which are due to information frictions that cause an over-supply of housing in hazardous areas. These losses are driven by accelerated depreciation of sinking homes, which experience loss of useful life and increased structural risk and cracks. Neighborhood impacts from uneven sinking over space also play an important role in lowering home values, as neighborhoods see significantly more flooding in response to uneven subsidence relative to surrounding areas.

Importantly, our estimates exclude city-wide costs that are shared among residents. For instance, our estimates do not capture expenditures incurred by the city government to repair water main breaks or sinkholes or replace macro-infrastructure as the main sewer line servicing the Mexico City that had to be replaced when subsidence altered the slope beyond repair (SACMEX, 2012). While these should be included for a full cost analysis, our estimate likely captures the largest part of the costs of subsidence as it is orders of magnitude larger than the costs of the entire budget of the water utility ($9.3B USD in present value), the cost of the Cutzamala surface water system ($0.58B USD in present value), or the cost of the recent drainage system replacement ($65 million USD).

More generally, our framework provides a systematic way of thinking about re-development in places that experience evolving environmental hazards such as climate impacts, erosion, or floods. Our model highlights the different economic role that *realizations* of environmental hazards play in driving re-development, as compared to expected *future* hazards that must be capitalized correctly into home values in order for the aggregate amount of housing in harm's way to be efficient. We find that information frictions about future realizations of environmental hazards, which have been found to be widespread in different contexts, can severely exacerbate the costs of these types of externalities by over-supplying housing in risky areas.

References

Associated Press (2022). Indonesia's capital is rapidly sinking into the sea. *NPR*. Published 26 January 2022. Accessed 23 February 2023.

Bagheri-Gavkosh, M., Hosseini, S. M., Ataie-Ashtiani, B., Sohani, Y., Ebrahimian, H., Morovat, F., and Ashrafi, S. (2021). Land subsidence: A global challenge. *Science of The Total Environment*, 778(15).

Bakkensen, L. A. and Barrage, L. (2022). Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics. *The Review of Financial Studies*, 35(8):3666–3709.

Bakkensen, L. A., Ding, X., and Ma, L. (2019). Flood Risk and Salience: New Evidence from the Sunshine State. *Southern Economic Journal*, 85(4):1132–1158.

Balboni, C. (2025). In Harm's Way? Infrastructure Investments and the Persistence of Coastal Cities. *American Economic Review*, 115(1).

Baum-Snow, N. and Han, L. (2024). The Microgeography of Housing Supply. *Journal of Political Economy*, 132(6).

Bawden, G., Sneed, M., Stork, S., and Galloway, D. (2023). Measuring Human-Induced Land Subsidence from Space. Technical report, USGS.

Bayer, P., Ferreira, F., and McMillan, R. (2007). A Unified Framework for Measuring Preferences for Schools and Neighborhoods. *Journal of Political Economy*, 115(4):588–638.

Benetton, M., Emiliozzi, S., Guglielminetti, E., Loberto, M., and Mistretta, A. (2025). Does Climate Change Adaptation Matter? Evidence from the City on the Water. *SSRN 4313434*.

Bishop, K. C. and Murphy, A. D. (2019). Valuing Time-Varying Attributes Using the Hedonic Model: When Is a Dynamic Approach Necessary? *The Review of Economics and Statistics*, 101(1):134–145.

Blundell, R. and Matzkin, R. (2014). Control functions in non-separable simultaneous equations models. *Quantitative Economics*, 5:271–295.

Borchers, J. and Carpenter, M. (2014). Land Subsidence from Groundwater Use in California. Technical report, California Water Foundation.

Bézy, T. (2025). Flood risk pricing and the housing supply: Evidence from insurance pooling in France. *mimeo*.

Cabral-Cano, E., Dixon, T. H., Miralles-Wilhelm, F., Díaz-Molina, O., Sánchez-Zamora, O., and Carande, R. E. (2008). Space geodetic imaging of rapid ground subsidence in Mexico City. *GSA Bulletin*, 120(11-12):1556–1566.

Carleton, T., Crews, L., and Nath, I. (2024). Is the World Running Out of Fresh Water? *AEA Papers and Proceedings*, 114:31–35.

Chaussard, E., Havazli, E., Fattahi, H., Cabral-Cano, E., and Solano-Rojas, D. (2021). Over a Century of Sinking in Mexico City: No Hope for Significant Elevation and Storage Capacity Recovery. *Journal of Geophysical Research: Solid Earth*, 126(4).

Cigna, F. and Tapete, D. (2021). Present-day land subsidence rates, surface faulting hazard and risk in Mexico City with 2014–2020 Sentinel-1 IW InSAR. *Remote Sensing of Environment*, 253:112161.

CONAGUA (2009). Situación del Subsector de Agua Potable, Alcantarillado y Saneamiento. Technical report, CONAGUA.

CONAGUA (2024). Actualización de la Disponibilidad Media Anual de Agua en el Acuífero Zona Metropolitana de la Ciudad de México. Technical report, Subdirección General Técnica, Gerencia de Aguas Subterráneas, CONAGUA, Ciudad de México.

DellaVigna, S. (2018). Structural Behavioral Economics. *Handbook of Behavioral Economics: Applications and Foundations 1*, 1:613–723.

Desmet, K., Kopp, R. E., Kulp, S. A., Nagy, D. K., Oppenheimer, M., Rossi-Hansberg, E., and Strauss, B. H. (2021). Evaluating the Economic Cost of Coastal Flooding. *American Economic Journal: Macroeconomics*, 13(2):444–486.

Dickstein, M. and Morales, E. (2018). What do Exporters Know? *The Quarterly Journal of Economics*, 133(4):1753–1801.

Dranove, D. and Jin, G. Z. (2010). Quality Disclosure and Certification: Theory and Practice. *Journal of Economic Literature*, 48(4):935–63.

Falk, A., Becker, A., Dohmen, T., Huffman, D., and Sunde, U. (2016). The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences. *IZA Discussion Paper No. 9674*.

Fattahi, H., Simons, M., and Agram, P. (2017). Insar Time-Series Estimation of the Ionospheric Phase Delay: An Extension of the Split Range-Spectrum Technique. *IEEE Transactions on Geoscience and Remote Sensing*, 55:5984–5996.

Fernández-Torres, E., Cabral-Cano, E., Solano-Rojas, D., Havazli, E., and Salazar-Tlaczani, L. (2020). Land subsidence risk maps and InSAR based angular distortion structural vulnerability assessment: an example in Mexico City. *Proceedings of the International Association of Hydrological Sciences*, 382:583–587.

Fernández-Torres, E. A., Cabral-Cano, E., Salazar-Tlaczani, L., and Solano-Rojas, D. (2025). Economic risk of differential subsidence in Mexico City (2014-2022). *Natural Hazards*, 121(3):2507–2534.

Fowler, L. C. (1981). Economic Consequences of Land Surface Subsidence. *Journal of the Irrigation and Drainage Division*, 107(2):151–159.

Frondel, M., Gerster, A., and Vance, C. (2020). The Power of Mandatory Quality Disclosure: Evidence from the German Housing Market. *Journal of the Association of Environmental and Resource Economists*, 7(1):181–208.

Gallagher, J. (2014). Learning about an Infrequent Event: Evidence from Flood Insurance Take-Up in the United States. *American Economic Journal: Applied Economics*, 6(3):206–233.

Gambolati, G. and Teatini, P. (2021). *Land Subsidence and its Mitigation*. The Groundwater Project.

Gechter, M. and Tsivanidis, N. (2025). Spatial Spillovers from High-Rise Developments: Evidence from Mumbai Mills. *mimeo*.

Gibson, M. and Mullins, J. T. (2020). Climate Risk and Beliefs in New York Floodplains. *Journal of the Association of Environmental and Resource Economists*, 7(6):1069–1111.

Gobierno del Distrito Federal (2007). Programa de Manejo Sustentable del Agua para la Ciudad de México. Technical report, SEGUIAGUA.

Gonzalez-Navarro, M. and Quintana-Domeque, C. (2016). Paving Streets for the Poor: Experimental Analysis of Infrastructure Effects. *The Review of Economics and Statistics*, 98(2):254–267.

Gourevitch, J. D., Kousky, C., Liao, Y. P., Nolte, C., Pollack, A. B., Porter, J. R., and Weill, J. A. (2023). Unpriced climate risk and the potential consequences of overvaluation in US housing markets. *Nature Climate Change*, 13(3):250–257.

Hino, M. and Burke, M. (2021). The effect of information about climate risk on property values. *Proceedings of the National Academy of Sciences*, 118(17).

Holtermans, R., Niu, D., and Zheng, S. (2024). Quantifying the impacts of climate shocks in commercial real estate markets. *Journal of Regional Science*, 64(4):1099–1121.

Hsiang, S., Kopp, R., Jina, A., Rising, J., Delgado, M., Mohan, S., Rasmussen, D. J., Muir-Wood, R., Wilson, P., Oppenheimer, M., Larsen, K., and Houser, T. (2017). Estimating economic damage from climate change in the United States. *Science*, 356(6345):1362–1369.

Hsiao, A. (2023). Sea Level Rise and Urban Adaptation in Jakarta. *mimeo*.

Hu, B., Zhou, J., Xu, S., Chen, Z., Wang, J., Wang, D., Wang, L., Guo, J., and Meng, W. (2013). Assessment of hazards and economic losses induced by land subsidence in Tianjin Binhai new area from 2011 to 2020 based on scenario analysis. *Natural Hazards*, 66(2):873–886.

Kala, N. (2017). Learning, adaptation and climate uncertainty: Evidence from Indian Agriculture. *MIT Center for energy and environmental policy research working paper*.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *Quarterly Journal of Economics*, 112(2):443–477.

Lee, S. and Zheng, S. (2023). Does Market Price Sea Level Rise Risk? Evidence from High-Tide Flooding. *MIT Center for Real Estate Research Paper No. 23/11*.

Liu, L., Yu, J., Chen, B., and Wang, Y. (2020). Urban subsidence monitoring by sbas-insar technique with multi-platform sar images: a case study of Beijing Plain, China. *European Journal of Remote Sensing*, 53(sup1):141–153.

Liu, Z., Liu, P.-W., Massoud, E., Farr, T. G., Lundgren, P., and Famiglietti, J. S. (2019). Monitoring Groundwater Change in California's Central Valley Using Sentinel-1 and GRACE Observations. *Geosciences*, 9(10):436.

Lixin, Y., Jie, W., Chuanqing, S., Guo, J.-W., Yanxiang, J., and Liu, B. (2010). Land Subsidence Disaster Survey and Its Economic Loss Assessment in Tianjin, China. *Natural Hazards Review*, 11(1):35–41.

López-Morales, C. and Mesa-Jurado, M. A. (2017). Valuation of Hidden Water Ecosystem Services: The Replacement Cost of the Aquifer System in Central Mexico. *Water*, 9(571).

Mahadevan, M., Meeks, R., and Yamano, T. (2023). Reducing information barriers to solar adoption: Experimental evidence from India. *Energy Economics*, 120.

Mahya, M., Kok, S., and Cado van der Lelij, A. (2021). Economic assessment of subsidence in Semarang and Demak, Indonesia. Technical report, Wetlands International.

Nanda, A. and Ross, S. L. (2012). The Impact of Property Condition Disclosure Laws on Housing Prices: Evidence from an Event Study Using Propensity Scores. *The Journal of Real Estate Finance and Economics*, 45(1):88–109.

Nicholls, R. J., Lincke, D., Hinkel, J., Brown, S., Vafeidis, A. T., Meyssignac, B., Hanson, S. E., Merkens, J.-L., and Fang, J. (2021). A global analysis of subsidence, relative sea-level change and coastal flood exposure. *Nature Climate Change*, 11(4):338–342.

Novelo-Casanova, D., Suárez, G., Cabral-Cano, E., Fernández-Torres, E., Fuentes-Mariles, O., Havazli, E., Jaimes, M. A., López-Espinoza, E., Martin-Del Pozzo, A. L., Morales-Barrera, W., Morales-Rodríguez, H., Nieto-Torres, A., Rodríguez-Elizarrarás, S., Solano-Rojas, D., and Velasco-Herrera, V. (2022). The Risk Atlas of Mexico City, Mexico: a tool for decision-making and disaster prevention. *Natural Hazards*, 111(1):411–437.

Ortega, F., Penaranda, F., and Taspinar, S. (2025). When the Water Recedes and Home Prices Don't: Flood Risk Learning and Neighborhood Spillovers. *mimeo*.

Ortega, F. and Taṣpınar, S. (2018). Rising sea levels and sinking property values: Hurricane Sandy and New York's housing market. *Journal of Urban Economics*, 106:81–100.

Ospital, A. (2024). Urban Policy and Spatial Exposure to Environmental Risk. *mimeo*.

Ostriker, A. and Russo, A. (2023). The Effects of Floodplain Regulation on Housing Markets. *mimeo*.

Parsons, T., Wu, P.-C., Wei, M. M., and D'Hondt, S. (2023). The Weight of New York City: Possible Contributions to Subsidence From Anthropogenic Sources. *Earth's Future*, 11(5).

Patel, D. (2024). Learning About a Warming World: Attention and Adaptation in Agriculture. *mimeo*.

Peter, H., Jäggi, A., Fernández-Sánchez, J., and Escobar, D. (2017). Sentinel-1A - First precise orbit determination results. *Advances in Space Research*, 60(5).

Petkov, I. and Ortega, F. (2025). Learning from experience: Flooding and insurance take-up in the flood zone and its periphery. *Journal of Risk and Insurance*, 92(2):212–356.

Petrin, A. and Train, K. (2010). A Control Function Approach to Endogeneity in Consumer Choice Models. *Journal of Marketing Research*, 47(1):3–13.

Rodella, A.-S., Zaveri, E., and Bertone, F., editors (2023). *The Hidden Wealth of Nations: The Economics of Groundwater in Times of Climate Change*. World Bank. License: CC BY 3.0 IGO.

Rodríguez Zamora, C. (2010). Estimation of a semi-parametric hazard model for Mexico's new housing market. *Bank for International Settlements*, Proceedings of the Fifth IFC Conference(34):587–599.

Rollet, V. (2025). Can We Rebuild a City? The Dynamics of Urban Redevelopment. *SSRN 5205801*.

Rosen, S. (1974). Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition. *Journal of Political Economy*, 82(1):34–55.

SACMEX (2012). *El Gran Reto del Agua en la Ciudad de México: Pasado, presente y prospectivas de solución para una de las ciudades más complejas del mundo*. Mexico City.

Secretaría de Gestión Integral de Riesgos y Protección Civil (2018). Causas y Consecuencias de la subsidencia en los cuadrantes Dos (Zona Poniente), Tres y Cuatro de la Zona Cero de la Ciudad de México, Anexo A. Technical report, Secretaría de Gestión Integral de Riesgos y Protección Civil, Ciudad de México.

Sinha, K. (2022). Fracking Disclosure, Collateral Value, and the Mortgage Market. *The Accounting Review*, 97(5):427–454.

Solano-Rojas, D., Cabral-Cano, E., Hernandez-Espriu, A., and Wdowinski, S. (2015). The relationship of InSAR-GPS land subsidence and the groundwater level decrease in wells of the Mexico City Metropolitan Area. *Boletin de la Sociedad Geologica Mexicana*, 67:273–283.

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

Wagner, K. R. H. (2022). Adaptation and Adverse Selection in Markets for Natural Disaster Insurance. *American Economic Journal: Economic Policy*, 14(3):380–421.

Werner, A. D., Zhang, Q., Xue, L., Smerdon, B. D., Li, X., Zhu, X., Yu, L., and Li, L. (2013). An Initial Inventory and Indexation of Groundwater Mega-Depletion Cases. *Water Resources Management*, 27(2):507–533.

Willemsen, W., Kok, S., and Kuik, O. (2020). The effect of land subsidence on real estate values. *International Association of Hydrological Sciences*, 382:703–707.

Wu, Y.-Y. and Madson, A. (2024). Error Sources of Interferometric Synthetic Aperture Radar Satellites. *Remote sensing*, 16(2).

Yoo, J. and Perrings, C. (2017). An externality of groundwater depletion: land subsidence and residential property prices in Phoenix, Arizona. *Journal of Environmental Economics and Policy*, 6(2):121–133.

Yu, C., Li, Z., and Penna, N. T. (2018a). Interferometric synthetic aperture radar atmospheric correction using a GPS-based iterative tropospheric decomposition model. *Remote Sensing of Environment*, 204:109–121.

Yu, C., Li, Z., Penna, N. T., and Crippa, P. (2018b). Generic Atmospheric Correction Model for Interferometric Synthetic Aperture Radar Observations. *Journal of Geophysical Research: Solid Earth*, 123(10):9202–9222.

Yu, C., Penna, N. T., and Li, Z. (2017). Generation of real-time mode high-resolution water vapor fields from GPS observations. *Journal of Geophysical Research: Atmospheres*, 122(3):2008–2025.

Yunjun, Z., Fattahi, H., and Amelung, F. (2019). Small baseline InSAR time series analysis: Unwrapping error correction and noise reduction. *Computers & Geosciences*, 133(104331).

Zappalà, G. (2024). Adapting to climate change accounting for individual beliefs. *Journal of Development Economics*, 169.

Table 1: Survey evidence on sorting and attention to sinking

	Suprised by severity	Ex-post factors used	Ex-ante factors used
Expected 5-year sinking	0.306**	-0.214*	-0.177**
	(0.117)	(0.124)	(0.080)
N	792	798	798
Mean Y (low sinking)	0.256	0.403	0.445
Strata FEs	X	X	X
Borough FEs	X	X	X
Education FEs	X	X	X
Respondent age FEs	X	X	X
Respondent gender FEs	X	X	X

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

Notes: Outcomes in order of appearance are an indicator for reporting that sinking has been worse than expected (with omitted categories "better than expected" and "same as expected"), an indicator for using any ex-post factors (seeing effects in the home or neighborhoods) to infer future probability of having issues due to sinking, and an indicator for using any ex-ante factors (asking an expert, asking around, the news, the risk zone, examining construction quality). "Expected 5-year sinking" is the average sinking rate in meters on the plot multiplied by five. All specifications include survey strata, borough, respondent age, respondent gender, and respondent educational attainment fixed effects. All regressions weighted by sampling weights. Standard errors clustered at the census tract.

A Tables

Table 2: Effect of future sinking on rents

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-0.024	-0.008	-0.006	-0.007
	(0.069)	(0.060)	(0.059)	(0.063)
log HH income	0.349***	0.345***	0.343***	0.339***
	(0.049)	(0.052)	(0.052)	(0.050)
N	702	702	702	702
Age bin FEs	X	X	X	X
Type x Borough FEs	X	X	X	X
N. room FEs	X	X	X	X
N. bedroom FEs	X	X	X	X
Year FEs	X			
Year x Borough FEs		X	X	X
Locality FEs			X	X
N. bedrooms x age trends				X

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

Notes: The outcome in every column is the log rents. All regressions include controls for log household income, an indicator for having a solid roof, and property-level sinking from the time of construction to the year of observation. Standard errors clustered at the census tract.

Table 3: Impact of plot sinking on appraisal values

	AR(1)		Perfect foresight	
	3%	6%	3%	6%
Property-level sinking (m)	-0.06052***	-0.06049***	-0.06075***	-0.06045***
	(0.01968)	(0.01968)	(0.01992)	(0.02001)
NPV future sinking	0.00009	0.00061	-0.00009	0.00012
	(0.00095)	(0.00386)	(0.00076)	(0.00238)
Observations	129,424	129,424	129,424	129,424
θ	-0.002	-0.01	0.001	-0.002
SE	(0.016)	(0.064)	(0.012)	(0.039)
ATE: plot sinking	-0.92%	-0.92%	-0.92%	-0.92%
Mean sinking between observ. (m)	0.151	0.151	0.151	0.151
Property FEs	X	X	X	X
Age FEs	X	X	X	X
Zone x year FEs	X	X	X	X
Borough x age trends	X	X	X	X
N. bedrooms x age trends	X	X	X	X

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

Notes: The outcome in all regressions is the log appraised value of the home. All estimates include property, age, and seismic zone by year fixed effects, and include linear time trends by borough and the number of bedrooms. Estimation limited to properties with more than one appraisal. Standard errors clustered at the census tract. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average sinking between observations in meters.

Table 4: Price response to expected future sinking, new builds

	(1)	(2)	(3)	(4)
AR(1) expectations				
NPV future sinking	-0.006	-0.006	-0.004	-0.005
	(0.005)	(0.005)	(0.005)	(0.004)
Observations	102,768	102,768	102,768	102,768
Mean NPV future sinking	12.008	12.008	12.008	12.008
ATE: Future sinking	-7.18%	-6.62%	-4.59%	-6.39%
Grid cell FEs	X	X	X	X
Year FEs	X			
Year x Seismic Zone FEs		X	X	X
Hedonic controls			X	X
Time trends by hedonic chars.				X
Perfect foresight				
NPV future sinking	-0.008	-0.007	-0.004	-0.004
	(0.005)	(0.005)	(0.004)	(0.004)
Observations	102,768	102,768	102,768	102,768
Mean NPV future sinking	13.172	13.172	13.172	13.172
ATE: Future sinking	-10.11%	-8.82%	-5.37%	-5.68%
Grid cell FEs	X	X	X	X
Year FEs	X			
Year x Seismic Zone FEs		X	X	X
Hedonic controls			X	X
Time trends by hedonic chars.				X

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

Notes: The outcome in all regressions is the log appraised value. Regressions limited to new builds. Standard errors clustered at the census tract. The Average Treatment Effect (ATE) is calculated by multiplying the (unrounded) reported coefficient by the average net present value of future sinking in meters assuming a 6% discount rate. "Ar(1) expectations" calculate the net present value as a function of the rate of sinking; "Perfect foresight" uses the observed sinking post sale through 2020, and then uses a fixed projected rate for subsequent years. Estimates of θ assume $\gamma = -0.060$.

Table 5: Survey evidence: Sinking and structural issues

		Does your home have any issues with...?			
	Useful life remaining	Structural issues	Levelling issues	Cracks: Hallways	Cracks: Walls
	(1)	(2)	(3)	(4)	(5)
Sinking (m)	-4.408***	0.100***	0.053*	0.033	0.078**
	(1.643)	(0.036)	(0.030)	(0.031)	(0.037)
Data source	Repeat appraisals	Survey	Survey	Survey	Survey
Observations	129,475	800	800	800	800
Mean Y	63.310	0.59	0.26	0.20	0.28

Notes: Standard errors clustered at the census tract. Column (1) regresses the appraised years of life remaining on total sinking on the property, controlling for property, age, and zone by year of appraisal fixed effects and time trends by borough, number of bedrooms, and type. These estimates exclude properties with only one appraisal. In Columns (2) through (5), all outcomes are indicators for reporting the stated issue in their home and the regressor is the total sinking on the property in meters from the year of construction (or 2007, whichever is later) through 2024. Columns (2) through (5) include survey strata, borough, respondent age, respondent gender, and respondent educational attainment fixed effects and are weighted by sampling weights.

Table 6: Survey evidence: Sinking and maintenance expenditure

	Any spending	log Spending (MXN)	log Years until next big investment	Any spending on non-sinking related issues
Total sinking on home (m)	0.002	0.443**	-0.188**	-0.004
	(0.044)	(0.175)	(0.089)	(0.034)
N	625	359	616	623
Mean Y (low sinking)	0.578	9,152.879	6.707	0.301

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

Notes: Standard errors clustered at the census tract. Outcomes in order of appearance are: An indicator for any spending on maintenance; log total spending on maintenance over the last 6 months; the log of respondent's estimated time in years to the next large investment in the home; and an indicator for the household mentioning any non-sinking category of response: electric work, remodeling, painting interiors or exteriors, additions, water-proofing, or replacing windows. "Total sinking on home" is the total subsidence on the property in meters from the year of construction (or 2007, whichever is later) to date. All specifications include survey strata, borough, respondent age, respondent gender, and respondent educational attainment fixed effects. All regressions weighted by sampling weights.

Table 7: Neighborhood sinking impacts home values and damages public infrastructure

	log appraised value		Reports issues with flooding		Reports issues with fissures	
	(1)	(2)	(3)	(4)	(5)	(6)
Property-level sinking (m)	-0.046**	-0.055***	-0.026	-0.015	-0.005	-0.158
	(0.020)	(0.020)	(0.121)	(0.116)	(0.170)	(0.132)
Abs. relative altitude change (m)	-0.161***		0.630***		-0.309	
	(0.059)		(0.194)		(0.513)	
SD of sinking within 500m (m)		-0.118**		1.296		5.780**
		(0.048)		(2.995)		(2.265)
Data source	Repeat transactions		Survey		Survey	
Observations	129,427	129,427	800	800	800	800
Mean Y	0.112	0.112	0.620	0.620	0.621	0.621
Mean Sinking measure	0.151	0.151	0.263	0.263	0.263	0.263
Mean Unevenness measure	0.063	0.006	0.028	0.005	0.028	0.005

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

Notes: Standard errors clustered at the census tract. Columns (1) and (2) regress the log appraised value on sinking on the property and the cumulative change in relative altitude (Column (1)) and the standard deviation of sinking within the 500m grid cell (Column (2)). Both columns include property, age, and zone by year of appraisal fixed effects and time trends by borough, number of bedrooms, and type, and exclude properties with only one appraisal. Columns (3) through (6) regress an indicator for the household reporting that their neighborhood has issues with the infrastructure problem in the header on total sinking on the property since it was built and the cumulative relative altitude change since 2007 in Columns (3) and (5), and the average standard deviation of sinking from 2007-2020 in Columns (4) and (6). Columns (3) through (6) include strata, borough, education, respondent age, and gender fixed effects, and regressions are weighted by sample weights.

Table 8: Impact of sinking on the probability of re-development

	(1)	(2)	(3)
Sinking on the property (m)	0.0254***	0.0053*	0.0052*
	(0.0040)	(0.0030)	(0.0029)
Relative altitude (m)	-0.0242*	-0.0182	-0.0182
	(0.0132)	(0.0124)	(0.0124)
log Price			-0.0005
			(0.0017)
N	307,859	307,859	307,859
Mean Y	0.0089	0.0089	0.0089
Adj. R2	0.070	0.079	0.079
Census tract FEs	X	X	X
Borough x year FEs	X	X	X
Hedonic controls		X	X

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

Notes: The outcome in all regressions is an indicator for the appraisal being associated with a new development within 5 years of the appraisal. Specifications with hedonic controls include fixed effects for the type (apartment, single family, etc.) and amenity classification (economical to luxury), square footage, number of bedrooms, number of bathrooms, number of parking spaces, age, plot size, built area, and number of floors. Standard errors clustered at the census tract.

Table 9: Impact of sinking on the probability of new development

	Pr. development on the Pixel		
	(1)	(2)	(3)
Pixel-level sinking (m)	0.00049***	0.00067*	0.00041***
	(0.00016)	(0.00036)	(0.00015)
Observations	1,500,497	1,500,497	1,500,497
Mean Y	0.005	0.005	0.005
Mean sinking (m)	0.34	0.34	0.34
Pixel FEs	X	X	X
Year FEs	X		
Zone x Year FEs		X	
Borough x year FEs			X

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

Notes: The outcome in all regressions is an indicator for construction beginning on a new development in the pixel-year. All specification include pixel fixed effects; (1) includes year fixed effects, (2) includes year by seismic zone fixed effects, and (3) includes borough by year fixed effects. Standard errors clustered at the census tract.

Table 10: Supply elasticity estimates

	First stage	IV	Probit	Probit, AME	Probit, Flexible CF	Probit, no CF
Relative altitude change (m)	-0.006					
	[-0.009, -0.002]					
log Price		0.403	5.935	0.273	5.952	0.010
		[-0.074, 1.398]	[0.101, 24.653]	[0.002, 1.106]	[0.217, 24.719]	[-0.148, 0.215]
Property-level sinking (m)	-0.121	0.091	1.654	0.076	1.660	0.895
	[-0.164, -0.084]	[0.026, 0.255]	[0.705, 4.867]	[0.031, 0.228]	[0.720, 4.887]	[0.347, 1.805]
u			-5.927	-0.272	-6.068	
			[-24.660, 0.016]	[-1.107, 0.000]	[-24.949, -0.304]	
u^2					-0.237	
					[-0.818, 0.122]	
u^3					0.295	
					[-0.046, 0.608]	
Share redeveloped	0.027	0.027	0.027	0.027	0.027	0.027
Observations	101,888	101,888	101,888	101,888	101,888	101,888
Tract FEs	X	X	X	X	X	X
Year FEs	X	X	X	X	X	X
Calendar month FEs	X	X	X	X	X	X
Type FEs	X	X	X	X	X	X
Amenity FEs	X	X	X	X	X	X
Bed, bath, age controls	X	X	X	X	X	X
F stat (KPR)	10.563					
Estimated elasticity		0.403	0.281	0.273	0.281	0.001
95% CI		[-0.074, 1.398]	[0.005, 0.873]	[0.002, 1.106]	[0.014, 0.873]	[-0.017, 0.019]

Notes: The outcome in Column (1) is the log transaction price; the outcome in Columns (2)-(6) is an indicator for the plot being redeveloped. Brackets under each estimate report the 95% bootstrap confidence interval. Standard errors clustered at the census tract in the first stage and linear probability model; confidence intervals calculated from bootstrapped sampling of census tracts in Columns (3)-(6), with 1,000 bootstrap samples drawn of 235 census tracts sampled in each. All estimates control for age, age^2, the number of bedrooms, the number of bathrooms, and fixed effects for the census tract, year, month of transaction, type (apartment, single family, etc.), and amenity classification (economy, semi-luxury, or luxury). Average marginal effects calculated over the sample with each observation assigned the relevant estimated fixed effects.

Table 11: Pigouvian tax on groundwater pumping and water consumption, rates

	Very low income	Low income	Middle income	Unsubsidized
Share of HHs	0.243	0.42	0.113	0.225
Avg. consumption (m^3)	23.35	29.427	27.294	35.098
Avg. marginal rate	\$0.303	\$0.574	\$1.167	\$1.691
Average bi-monthly bill	\$4.175	\$13.154	\$22.106	\$59.34
Tax as % of marginal rate	292.81%	154.44%	75.98%	52.44%
Tax as % of bill	297.56%	119.03%	65.69%	31.47%

Notes: Rates from the 2024 schedule. Average consumption estimated from city block averages reported in 2019. All monetary values converted to 2018 USD using a 19 MXN/USD exchange rate, and deflating 2024 prices to 2018. "Marginal rates" refer to the marginal price of water at average city block consumption per household.

B Figures

Figure 1: Structures affected by subsidence in Mexico City



(a) Building distorted by uneven sinking in downtown Mexico City



(b) Church sinking to one side in Texcoco, Mexico City metro area



(c) Building whose foundation is exposed on one side as pavement sinks unevenly around it, downtown Mexico City



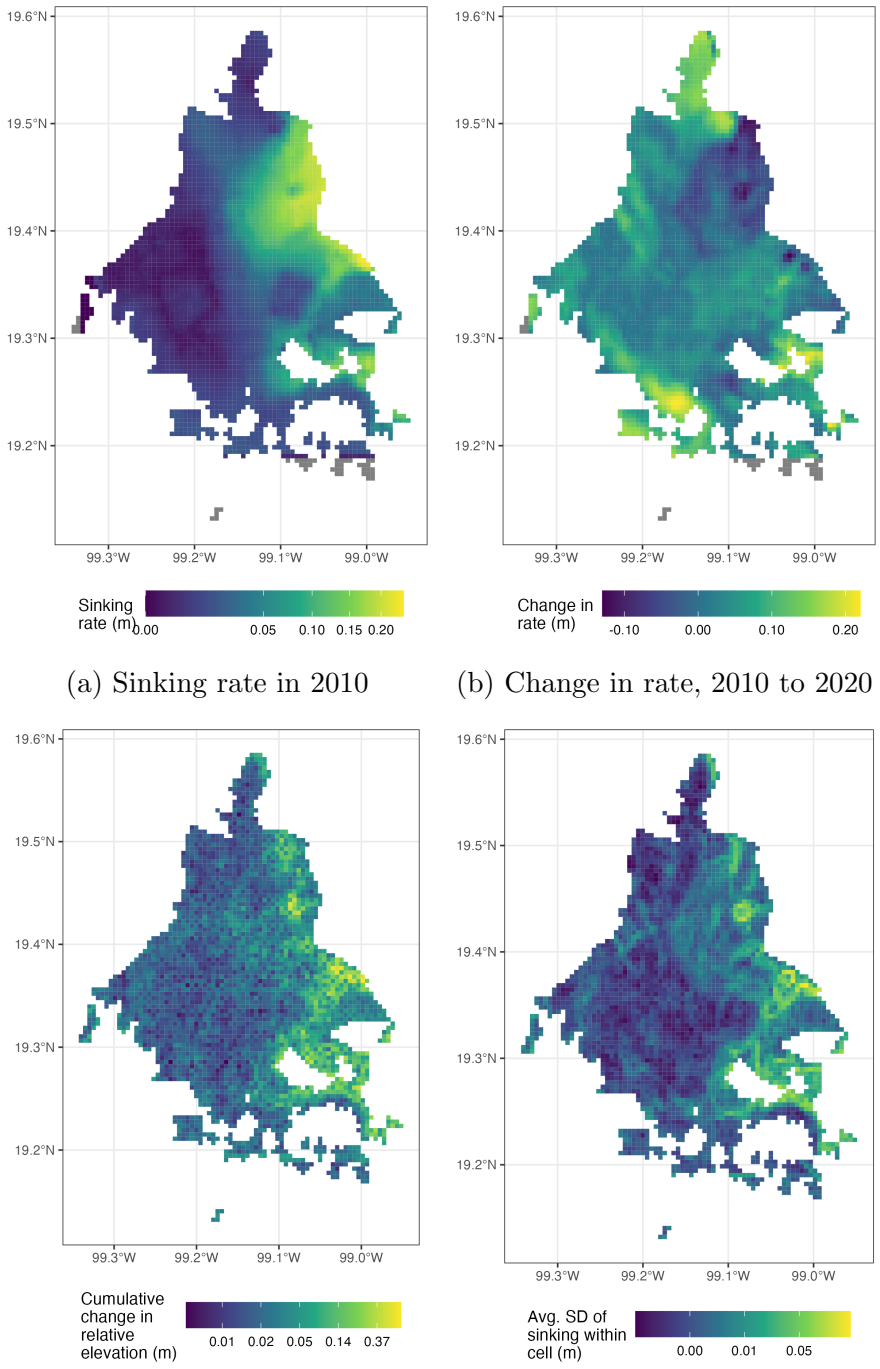
(d) Sinkhole, Coyoacán



(e) Wall tilting into sidewalk, Venustiano Carranza

Notes: Subfigure (a) is from the New York Times, article "Mexico City, Parched and Sinking, Faces a Water Crisis". Sub figure (b) is author's photograph from Texcoco, a city in the greater Mexico City metropolitan area located on the former lake bed. Subfigure (c) is author's photograph from the historic center of Mexico City. Subfigures (d) and (e) are from the Mexico City government's Subsidence Atlas (Secretaría de Gestión Integral de Riesgos y Protección Civil, 2018).

Figure 2: Subsidence rates in Mexico City



(a) Sinking rate in 2010

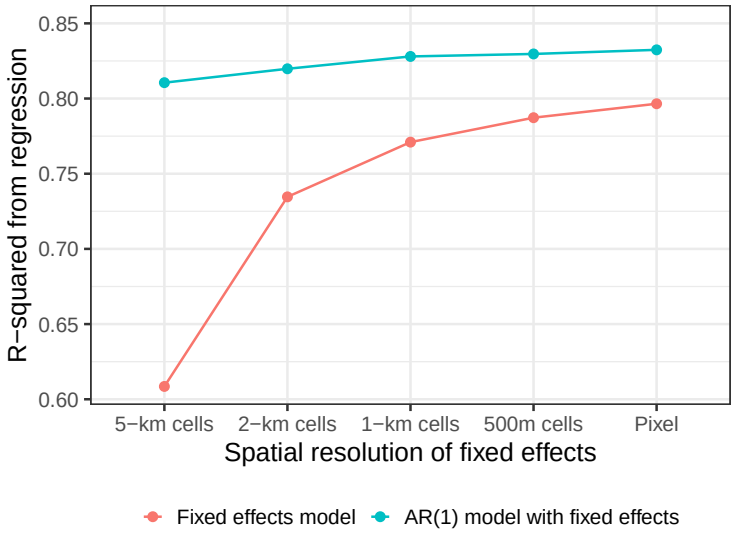
(b) Change in rate, 2010 to 2020

(c) Relative altitude change

(d) Standard deviation of sinking

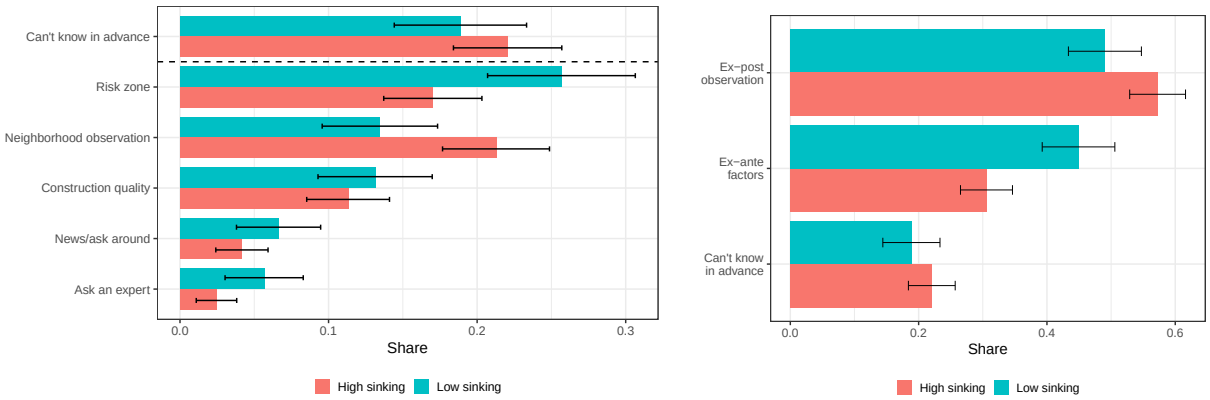
Notes: Map extent limited to grid cells with positive population within the state of Mexico City. Cells in grey do not have data in 2010. Panel (a) plots the sinking rate at the pixel level in 2010. Panel (b) plots the change in the sinking rate from 2010 to 2020. Panel (c) plots the total cumulative relative elevation change from 2007 to 2020. Panel (d) plots the average standard deviation of sinking, which averages the cross sectional standard deviation across years.

Figure 3: Predictability of subsidence over space



Notes: Each point represents the unadjusted R^2 of a regression of the pixel-year level sinking rate on a set of increasingly granular spatial fixed effects whose resolution is noted on the X-axis. The regressions generating the points in pink include only these fixed effects; regressions generating the points in blue additionally include the lagged sinking rate as an explanatory variable.

Figure 4: How can you know if a home will have issues with subsidence in the future?

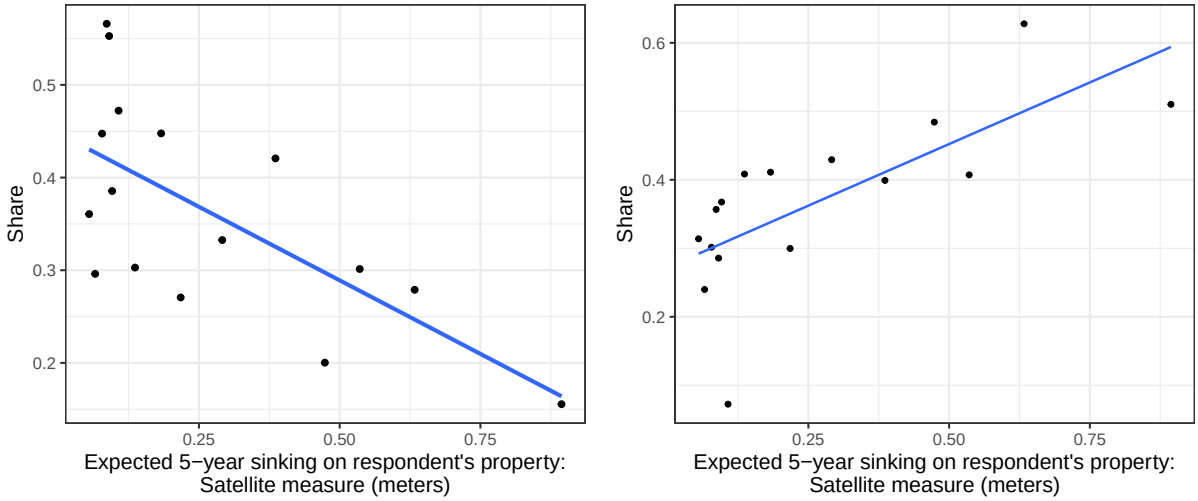


(a) Distribution of responses

(b) Ex-ante vs ex-post ways of knowing

Notes: Each bar plots the average of an indicator for the respondent mentioning a way of knowing that a house will have future issues with sinking in that category, weighted by the population weights and separated by the high-sinking and low-sinking strata. Black bars represent the 95% confidence interval around the mean. "Risk zone" groups any answer that references the location, geology, or risk area of the property. The right figure groups ex-ante and ex-post mechanisms and reports the mean for each. "Any ex-ante available information" includes knowing the risk zone, asking an expert/engineer, seeing issues reported in the news, or knowing the construction quality of the home.

Figure 5: Sinking against attention and surprise

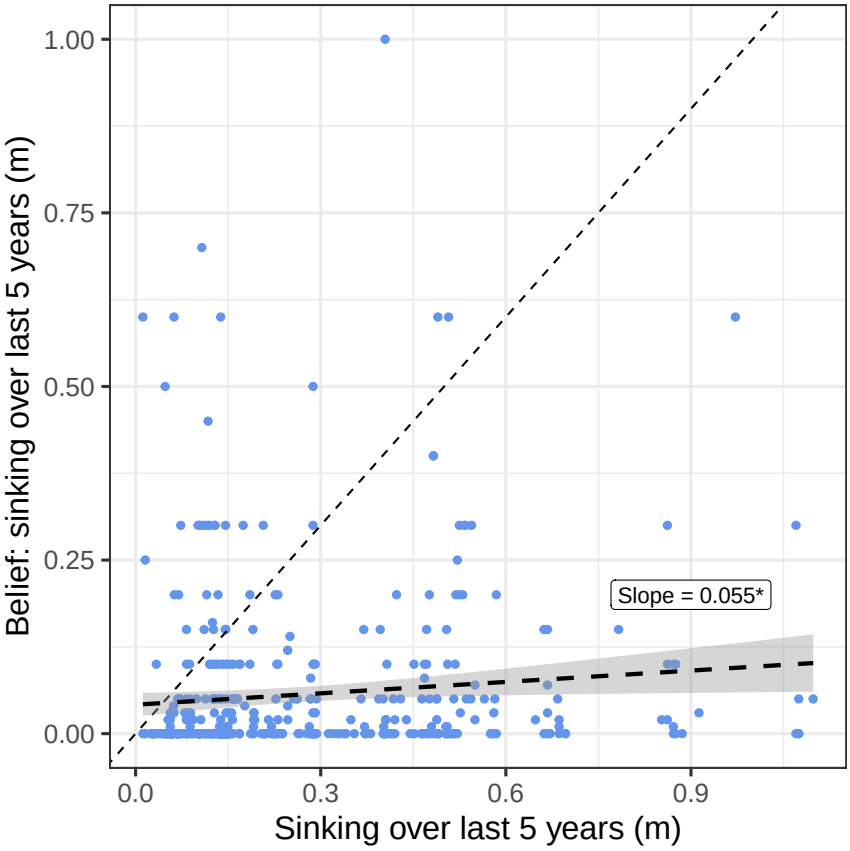


(a) Uses ex-ante available information

(b) Surprised by severity of sinking

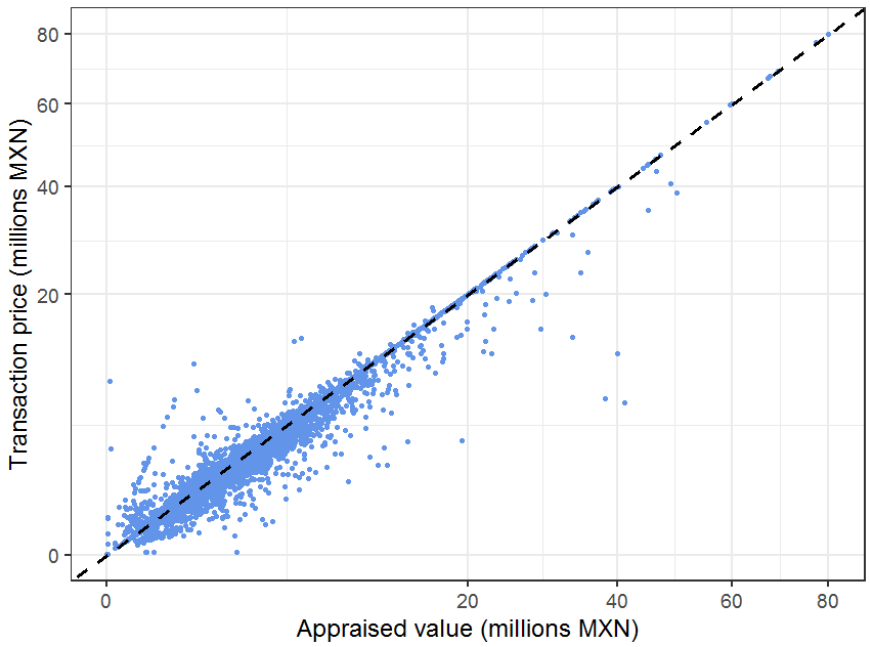
Notes: Each plot calculates the average outcome within 16 quantile bins of expected 5-year plot-level sinking, estimated as $5\overline{s}_j$ where $\overline{s}_j$ is the average sinking rate on the plot. Panel (a) plots the average of an indicator for the respondent reporting they use any ex-ante available information to infer sinking risk (knowing the risk zone, asking an expert/engineer, seeing issues reported in the news, or knowing the construction quality of the home). Panel (b) plots the average of an indicator for the respondent reporting that the issues related to subsidence have been more severe than they anticipated. Linear line of best fit plotted in Panel (a); non-linear relationship in (b) approximated with Loess regression.

Figure 6: Beliefs about past sinking versus satellite measures



Notes: Each point plots the respondent's report of how much their plot has sunk in the last five years against the satellite measure of how much that plot sunk over the same period. The dotted line plots the 45° line. The thick dashed line with prediction interval represents the line of best fit from a linear regression.

Figure 7: Appraisals and transaction prices



Notes: Each point represents an appraisal-transaction price pair from the mortgage data; a random sample of 50,000 pairs was drawn from the full sample of 5,645,851. The dashed line is the 45° line. The slope and adjusted R^2 reported are from a regression of the transaction price on the appraised value with no intercept.

Figure 8: Market clearing

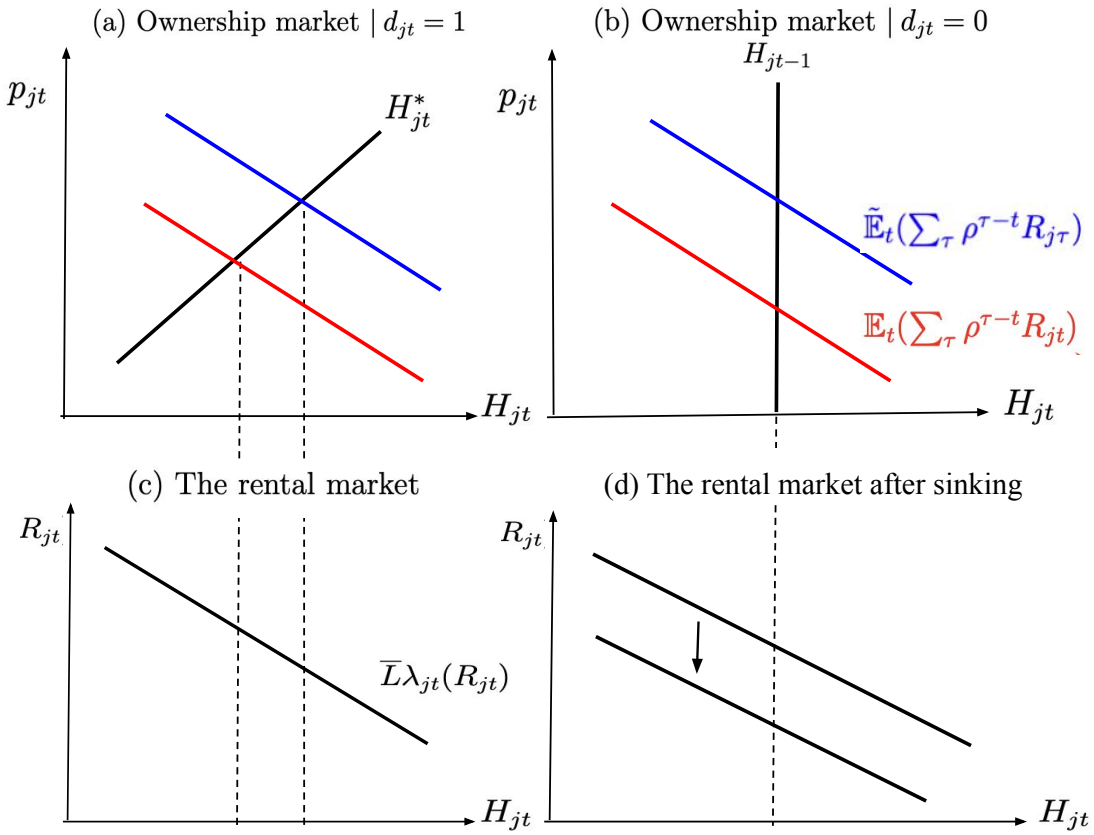
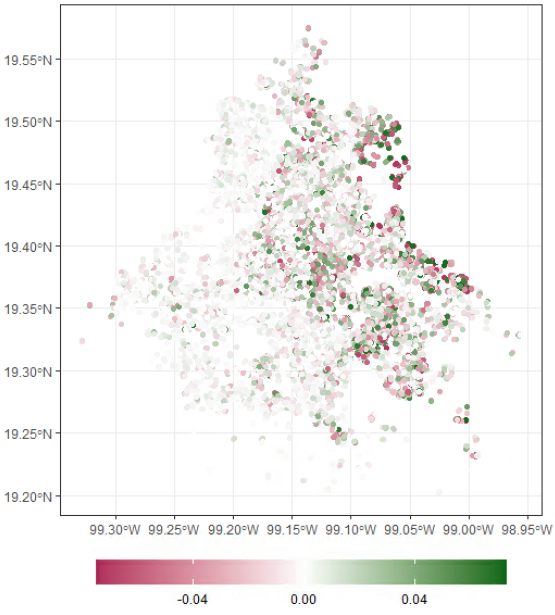
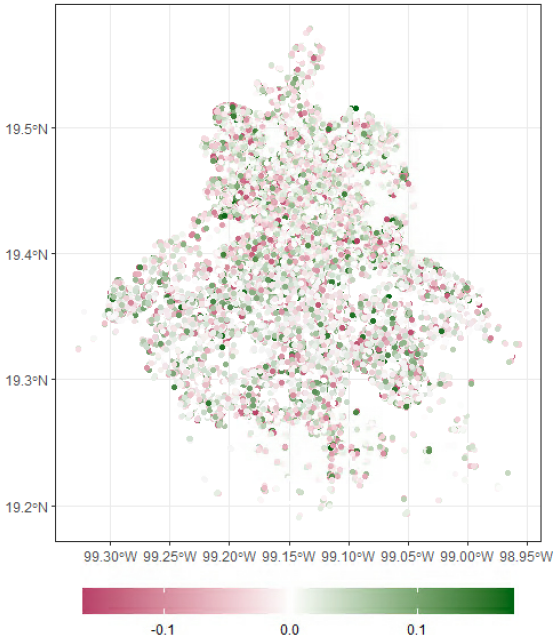


Figure 9: Residualized prices and sinking

(a) Residualized subsidence

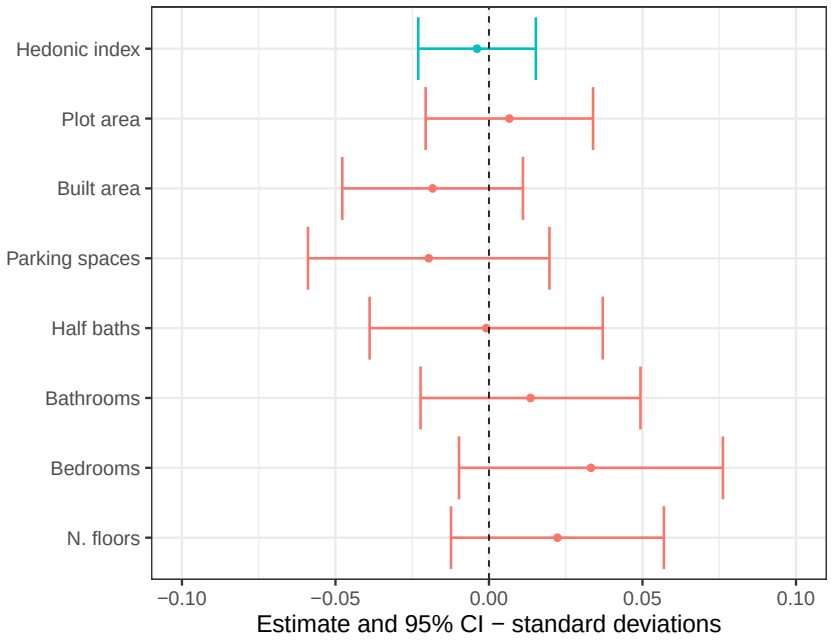


(b) Residualized log appraisal values



Notes: Each point plots the residual from a regression of the outcome on property, age, year of appraisal by zone fixed effects and linear trends by the number of bedrooms and borough. Residuals censored at the 1% and 99% percentiles and pseudo-log transformation applied to color scheme for visual clarity. Points jittered by a factor of 0.002 (maximum displacement of approximately 222 meters) to facilitate visualization of multiple points per property.

Figure 10: Pre-trends test: Subsidence and market composition



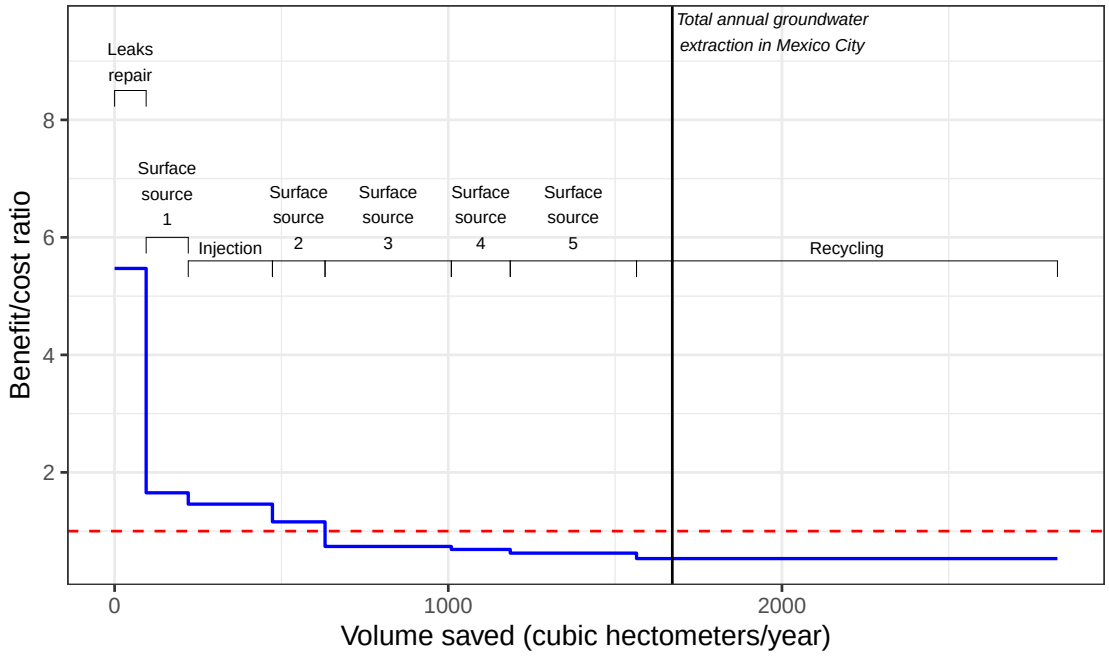
Notes: Plots coefficients and 95% confidence intervals for separate regressions of the outcome on sinking to date on the property in meters. The hedonic index is the simple average of the Z-score of plot area, built area, parking spaces, half baths, bathrooms and bedrooms. All outcomes are standardized by the mean and standard deviation of the variable. Each regression is of the the attribute listed in the y axis regressed on on plot-level sinking, 500 meter grid fixed effects, vintage by year of sale fixed effects, and linear trends by borough and seismic zone. Standard errors clustered at the census tract.

Figure 11: Welfare benefits of counterfactual policies



Notes: Each bar height represents the change in surplus by market participant for the counterfactual minus a baseline equilibrium, expressed in terms of billions of 2018 US dollars. "Stop future sinking" sets $\mathbb{E}(s_{j\tau}) = 0 \, \forall j, \tau > t$ and compares to the observed equilibrium. "Full information" sets $\theta = 1$ and keeps observed sinking, relative to the observed equilibrium. "No sinking | full information" shows the benefits of a counterfactual with $\theta = 1$ and $\mathbb{E}(s_{j\tau}) = 0 \forall j, \tau > t$ relative to a counterfactual with observed sinking but $\theta = 1$. Error bars mark the 5th to the 95th percentile range of estimates from 1,000 samples γ and η from normal distributions centered around their point estimates with standard deviation given by the standard error of the estimate.

Figure 12: Merit order of groundwater recharge policies



Notes: The horizontal axes plots the cumulative water saved by each intervention. The potential surface water sources referenced are taken from López-Morales and Mesa-Jurado (2017) and represent the following basins (in order labeled): Taxhimay, Oriental-Libres, Amacuzac, Temascaltepec, and Tecolutla. 6% discount rate used to calculate present value of costs and benefits. The total benefit of the intervention is in terms of abated subsidence, calculated using a value of $0.89 per cubic meter of abated pumping and bringing the stream of benefits over the life of the intervention to the present value.

C Table Appendix

Table A1: Survey evidence for information frictions: Limiting to respondents who have moved in the last 10 years

	Suprised by severity	Ex-post factors used	Ex-ante factors used
Expected 5-year sinking	0.142	-0.494	-0.340
	(0.290)	(0.326)	(0.271)
N	141	145	145
Mean Y (low sinking)	0.281	0.356	0.405
Strata FEs	X	X	X
Borough FEs	X	X	X
Education FEs	X	X	X
Respondent age FEs	X	X	X
Respondent gender FEs	X	X	X

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

Notes: Data limited to respondents who report moving to their current unit within the last 10 years. Outcomes in order of appearance are an indicator for using any ex-post factors (seeing effects in the home or neighborhoods) to infer future probability of having issues due to sinking, an indicator for using any ex-ante factors (asking an expert, asking around, the news, the risk zone, examining construction quality), and an indicator for reporting that sinking has been more severe than they anticipated. "Expected 5-year sinking" is the average sinking rate times five. All specifications include survey strata, borough, respondent age, respondent gender, and respondent educational attainment fixed effects. All regressions weighted by sampling weights. Standard errors clustered at the census tract.

Table A2: Subsidence and sorting on risk, discounting preferences

	Discounting			Risk aversion		
	(1)	(2)	(3)	(4)	(5)	(6)
Expected 5-year sinking	-0.487			0.300		
	(0.426)			(0.303)		
Sinking in last 5 years (m)		-0.208			0.163	
		(0.447)			(0.305)	
Relative altitude change in last 5 years (m)			1.504			1.253*
			(0.938)			(0.673)
N	391	391	391	353	353	353
Mean Y (low sinking)	0.682	0.682	0.682	0.624	0.624	0.624
Strata FEs	X	X	X	X	X	X
Borough FEs	X	X	X	X	X	X
Education FEs	X	X	X	X	X	X
Respondent age FEs	X	X	X	X	X	X
Respondent gender FEs	X	X	X	X	X	X

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

Notes: Outcomes in order of appearance are the implied discount rate and risk aversion coefficient measured by Falk instruments (higher discount rates indicate higher impatience). Respondents had a random 50% change on being asked the discounting or risk aversion questions, with 400 being surveyed on each each. "Expected 5-year sinking" is the average sinking rate on the property over the full panel multiplied by five. "Sinking in last 5 years (m)" is the total subsidence on the property over the last 5 years. All specifications include survey strata, borough, respondent age, respondent gender, and respondent educational attainment fixed effects. All regressions weighted by sampling weights. Standard errors clustered at the census tract.

Table A3: Repeat sales sample compared to the full appraisal sample

	(1)	(2)	(3)	Mean Y
Single-family home	-0.166***	-0.151***	-0.067***	0.196
	(0.007)	(0.007)	(0.004)	
Apartment	0.171***	0.155***	0.070***	0.794
	(0.007)	(0.007)	(0.004)	
Empty lot	-0.001	-0.000	0.001	0.004
	(0.000)	(0.000)	(0.000)	
Communal property	-0.002***	-0.003***	-0.002***	0.004
	(0.001)	(0.001)	(0.000)	
Medium class	-0.005	-0.014	-0.015**	0.622
	(0.013)	(0.011)	(0.007)	
Plus class	-0.047***	-0.017***	-0.000	0.171
	(0.009)	(0.006)	(0.003)	
Luxury class	-0.020***	-0.013***	-0.002	0.035
	(0.002)	(0.002)	(0.001)	
Age	-5.518***	-5.196***	-2.191***	12.931
	(0.283)	(0.264)	(0.164)	
Plot size (sq. m)	122.636	150.028***	40.917	523.763
	(76.546)	(47.630)	(32.436)	
Construction size (sq. m)	-35.146***	-28.396***	-12.113***	97.059
	(1.344)	(1.353)	(0.645)	
N. floors	-0.237***	-0.199***	-0.205***	1.832
	(0.033)	(0.034)	(0.043)	
N. bedrooms	-0.241***	-0.212***	-0.096***	2.353
	(0.012)	(0.011)	(0.008)	
N. bathrooms	-0.309***	-0.219***	-0.077***	1.644
	(0.019)	(0.015)	(0.008)	
N. parking	-0.210***	-0.136***	-0.041***	1.129
	(0.021)	(0.018)	(0.010)	
Price (millions MXN)	-0.766***	-0.555***	-0.171***	2.224
	(0.049)	(0.044)	(0.015)	
Sinking to date (m)	-0.004	-0.038***	-0.014***	0.181
	(0.011)	(0.009)	(0.004)	
Rel. altitude (m)	0.193	0.130	-0.048	0.338
	(0.180)	(0.155)	(0.071)	
N. observations	322538	322538	322538	322538
Borough x Year FEs		X	X	
Census tract FEs			X	

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

Notes: Each row is the coefficient of a separate regression where the outcome is listed in the row label and the explanatory variable is an indicator for being included in the repeat sales sample. All regressions include an intercept and cluster standard errors at the census tract level. The last column reports the overall mean of the outcome in each row.

Table A4: Appraised values and transaction prices

	Price	Price	% deviation	% deviation
Intercept		0.049***	0.041***	
		(0.008)	(0.005)	
Appraised value	0.967***	0.958***		
	(0.004)	(0.003)		
Sinking intensity (m)			0.006	0.007
			(0.007)	(0.007)
Observations	5,645,851	5,645,851	5,645,851	5,645,851
Mean Y	2.048	2.048	0.044	0.044
Adj. R2	0.528	0.528	0	0.007
Loan controls				X

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

Notes: The outcome in Columns (1) and (2) is the transaction price in millions of pesos. The explanatory variable in (1) and (2) is the appraised value in millions of pesos. The outcome in Columns (3) and (4) is the percent deviation in appraised values from transaction values, $|p_{transac} - p_{appraised}|/p_{transac}$. The explanatory variable in Columns (3) and (4) is average total sinking in the locality from 2007 to 2020; Column (4) includes controls for the income of the borrower, the interest rate assigned to the loan, bank-level fixed effects, and the size of the down payment. Standard errors in all specifications are clustered at the locality.

Table A5: Impact of plot sinking on appraised useful life remaining

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-3.23374***	-4.83280***	-4.10111**	-4.40771***
	(1.14822)	(1.48152)	(1.66468)	(1.64326)
Observations	129,475	129,475	129,475	129,475
Mean sinking between observ. (m)	0.151	0.151	0.151	0.151
Property FEs	X	X	X	X
Year FEs	X	X	X	X
Age FEs	X	X	X	X
Zone x year FEs		X	X	X
Borough x age trends			X	X
N. bedrooms x age trends				X

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

Notes: The outcome in all regressions is the appraised useful years of life remaining. All estimates include property, age, and year of appraisal fixed effects. Estimation limited to properties with more than one appraisal. Standard errors clustered at the census tract. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average sinking between observations in meters.

Table A6: Impact of plot sinking on maintenance expenditures

	(1)	(2)	(3)	(4)
Property-level sinking (m)	0.016	0.023**	0.022**	0.023**
	(0.012)	(0.010)	(0.010)	(0.010)
log HH income	0.062***	0.063***	0.063***	0.062***
	(0.012)	(0.012)	(0.012)	(0.012)
N	5,607	5,607	5,607	5,607
Mean dept. variable	0.117	0.117	0.117	0.117
Age bin FEs	X	X	X	X
Type x Borough FEs	X	X	X	X
N. room FEs	X	X	X	X
N. bedroom FEs	X	X	X	X
Year FEs	X			
Year x Borough FEs		X	X	X
Locality FEs			X	X
N. bedroom x age trends				X

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

Notes: The outcome in all regressions is an indicator for incurring any maintenance spending. Standard errors clustered at the locality. All regressions include controls for log household income and an indicator for having a solid roof.

Table A7: Impact of sinking is stable across specifications

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-0.099***	-0.095***	-0.061***	-0.060***
	(0.014)	(0.015)	(0.020)	(0.020)
Observations	129,431	129,431	129,431	129,427
ATE: plot sinking	-1.49%	-1.44%	-0.92%	-0.92%
Mean sinking between observ. (m)	0.151	0.151	0.151	0.151
Property FEs	X	X	X	X
Year FEs	X	X	X	X
Age FEs	X	X	X	X
Zone x year FEs		X	X	X
Borough x age trends			X	X
N. bedrooms x age trends				X

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

Notes: Standard errors clustered at the census tract. The outcome in all columns is the log appraised value. All estimates include property, age, and year fixed effects. Estimation limited to properties with more than one appraisal. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average subsidence between observations in meters.

Table A8: Impact of sinking on neighboring pixels on appraisal values

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-0.08965***	-0.09683***	-0.06815***	-0.06175***
	(0.01582)	(0.01887)	(0.01977)	(0.02008)
Sinking on neighboring properties (m)	0.00541	-0.00221	0.00052	0.00139
	(0.00364)	(0.00502)	(0.00482)	(0.00482)
Observations	16,438	16,438	16,437	16,437
Mean sinking between observ. (m)	0.034	0.034	0.034	0.034
Start year x end year FEs	X			
Zone x start year x end year FEs		X	X	X
Borough x age trends			X	X
N. bedrooms x age trends				X
Binned floor space x age trends				X

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

Notes: Standard errors clustered at the census tract. The outcome in every column is the change in log appraised value. All estimates difference between subsequent appraisals in the same property. "Start year x End year FEs" are fixed effects by year of appraisal by year of last appraisal, which control for the change in age between appraisals.

Table A9: Repeat appraisals results: Pre-2017 earthquake

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-0.115***	-0.113***	-0.084***	-0.083**
	(0.027)	(0.02847)	(0.032)	(0.032)
Observations	78,453	78,453	78,453	78,453
ATE: plot sinking	-1.28%	-1.26%	-0.94%	-0.92%
Mean sinking between observ. (m)	0.112	0.112	0.112	0.112
Property FEs	X	X	X	X
Year FEs	X	X	X	X
Age FEs	X	X	X	X
Zone x year FEs		X	X	X
Borough x age trends			X	X
N. bedrooms x age trends				X

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

Notes: Standard errors clustered at the census tract. The outcome in every column is the log appraised value. All estimates include property, age, and year fixed effects. Estimation limited to properties with more than one appraisal. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average subsidence between observations in meters. Limited to appraisals for transactions before 2017.

Table A10: Repeat appraisals, Excluding localities with large differences between appraisals and transaction prices

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-0.099***	-0.095***	-0.061***	-0.061***
	(0.014)	(0.015)	(0.020)	(0.020)
Observations	129,228	129,228	129,228	129,224
ATE: plot sinking	-1.49%	-1.44%	-0.92%	-0.92%
Mean sinking between observ. (m)	0.151	0.151	0.151	0.151
Property FEs	X	X	X	X
Year FEs	X	X	X	X
Age FEs	X	X	X	X
Zone x year FEs		X	X	X
Borough x age trends			X	X
N. bedrooms x age trends				X

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

Notes: Estimation limited to localities where the median difference between appraisal values and transaction prices is less than 1%. The outcome in every column is the log appraised value. All estimates include property, age, and year fixed effects. Estimation limited to properties with more than one appraisal. Standard errors clustered at the census tract. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average subsidence between observations in meters.

Table A11: Significance is robust to using Conley standard errors

	(1)	(2)	(3)	(4)
Property-level sinking (m)	-0.099***	-0.095***	-0.061***	-0.060***
	(0.020)	(0.020)	(0.021)	(0.022)
Observations	129,431	129,431	129,431	129,427
ATE: plot sinking	-1.49%	-1.44%	-0.92%	-0.92%
Mean sinking between observ. (m)	0.151	0.151	0.151	0.151
Property FEs	X	X	X	X
Year FEs	X	X	X	X
Age FEs	X	X	X	X
Zone x year FEs		X	X	X
Borough x age trends			X	X
N. bedrooms x age trends				X

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

Notes: Standard errors calculated using 2-kilometer radius for pooled Conley standard error calculation. The outcome in all columns is the log appraised value. All estimates include property, age, and year fixed effects. Estimation limited to properties with more than one appraisal. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average subsidence between observations in meters.

Table A12: Interest rates for mortgages are uncorrelated with tract-level sinking

	(1)	(2)
Intercept	10.135***	
	(0.037)	
Sinking intensity (m)	0.080	0.005
	(0.047)	(0.031)
Observations	5,645,851	5,645,851
Mean Y	10.171	10.171
Adj. R2	0.001	0.217
Loan controls		X

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

Notes: Regression of the interest rate (in percentage points) charged on a mortgage on the total sinking in meters on the census tract from 2007 to 2020. Column (2) includes controls for the appraised value of the home, the borrower's income and the down payment as a share of the value, as well as bank and month fixed effects. Standard errors clustered at the census tract.

Table A13: Effect of sinking on the height of new developments

	(1)	(2)	(3)	(4)
Sinking pre-construction (m)	0.120	0.431	0.121	0.115
	(0.349)	(0.333)	(0.339)	(0.360)
Observations	13,608	13,608	13,608	13,608
ATE: plot sinking	2.55%	9.17%	2.58%	2.45%
Mean sinking (m)	0.213	0.213	0.213	0.213
500m grid FEs	X	X	X	X
Year started x Year finished FEs	X			
Year started x Zone x Year finished FEs		X	X	X
Hedonic controls			X	X
Linear trends by hedonic char.				X
Sinking pre-construction (m)	0.424	0.522	0.061	0.049
	(0.381)	(0.374)	(0.365)	(0.380)
Relative elevation change (m)	-11.616*	-4.428	1.357	1.254
	(6.168)	(4.462)	(4.117)	(4.050)
NPV future sinking (m)	0.017	0.003	0.046	0.047
	(0.041)	(0.039)	(0.034)	(0.034)
Observations	13,608	13,608	13,608	13,608
ATE: plot sinking	9.02%	11.11%	1.29%	1.05%
ATE: Relative altitude change	-14.82%	-5.65%	1.73%	1.6%
ATE: Future sinking	15.34%	2.54%	40.7%	41.82%
Mean sinking (m)	0.213	0.213	0.213	0.213
Mean relative altitude change (m)	0.013	0.013	0.013	0.013
Mean NPV future sinking (m)	8.942	8.942	8.942	8.942
500m grid FEs	X	X	X	X
Year started x Year finished FEs	X			
Year started x Zone x Year finished FEs		X	X	X
Hedonic controls			X	X
Linear trends by hedonic char.				X

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

Notes: The outcome in all specifications is the height (number of floors) of the development. Standard errors clustered at the census tract. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the mean of the corresponding variable.

Table A14: Effect of sinking on the amenities of new developments

Amenity index	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Sinking pre-construction (m)	0.057	0.029	0.084	0.051	0.084	0.046				
	(0.108)	(0.110)	(0.109)	(0.111)	(0.110)	(0.112)				
Relative elevation change (m)			-0.979	-0.909	-0.972	-0.870	-0.621	-0.745	-0.608	-0.714
			(0.751)	(0.739)	(0.764)	(0.759)	(0.736)	(0.738)	(0.744)	(0.755)
NPV future sinking (m)					0.000	0.003			0.001	0.003
					(0.009)	(0.009)			(0.009)	(0.009)
Observations	13,610	13,610	13,610	13,610	13,610	13,610	13,616	13,616	13,616	13,616
ATE: plot sinking	1.22%	0.61%	1.8%	1.08%	1.79%	0.98%				
ATE: Relative altitude change			-1.25%	-1.16%	-1.24%	-1.11%	-0.79%	-0.95%	-0.78%	-0.91%
ATE: Future sinking					0.36%	2.29%			1.14%	2.91%
Mean sinking (m)	0.213	0.213	0.213	0.213	0.213	0.213				
Mean relative altitude change (m)			0.013	0.013	0.013	0.013	0.013	0.013	0.013	0.013
Mean NPV future sinking (m)					8.941	8.941			8.94	8.94
500m grid FEs	X	X	X	X	X	X	X	X	X	X
Year started x Year finished FEs	X		X		X		X		X	
Year started x Zone x Year finished FEs		X		X		X		X		X

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

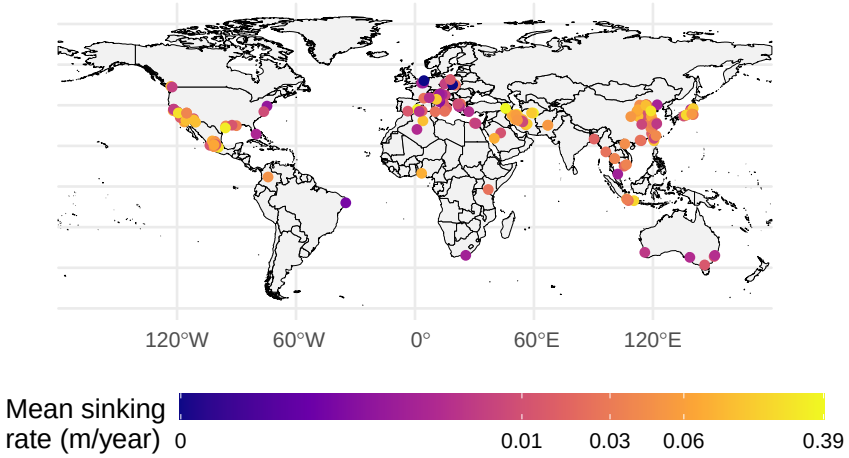
Notes: The outcome in all specifications is the amenity score of the development, which takes the simple average of standardized values for the floor space, construction space, number of bedrooms, number of bathrooms, number of parking spaces, and indicators for the presence of a gym, security desk, pool, and green space. Standard errors clustered at the census tract. The Average Treatment Effect (ATE) is calculated by multiplying the reported coefficient by the average subsidence between observations in meters.

Table A15: Parameters needed for welfare analysis

	Value	**Interpretation**	**How estimated/calibrated**
γ	-0.06	Disutility from subsidence	Repeat transactions estimation
θ	0.08	Extent of information frictions	Value of future sinking in new home price
η	0.28	Supply elasticity	Redevelopment estimation
ρ	1/1.06	Discount factor	Chosen to reflect a 6% social discount rate
μ	1.111	Mean of log cost shocks	Calibrated from distribution of structural residual in data
σ	0.413	Standard deviation of log cost shocks	Calibrated from distribution of structural residual in data
$\tilde{\xi}^S_{jt}$	(vector)	log cost shocks	Exactly matched to data using equilibrium conditions
e_{jt}	(vector)	Unobserved demand shifters	Exactly matched to data using equilibrium conditions
D	2.2	Demolition cost	Calibrated to match mean redevelopment probability of 2.7%

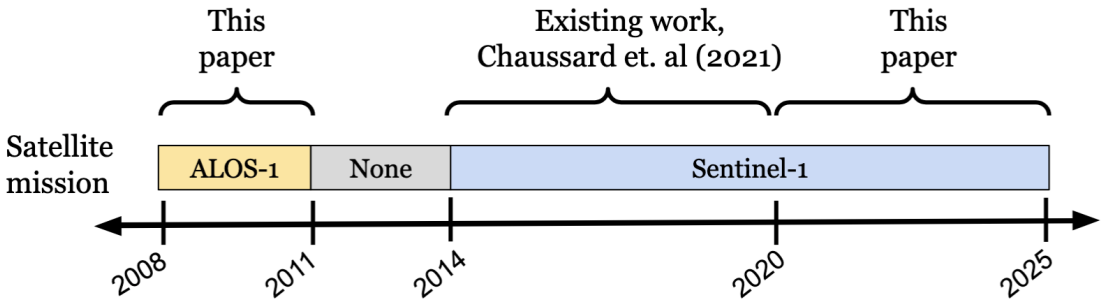
D Figure Appendix

Figure B1: Subsidence rates from the geophysics literature



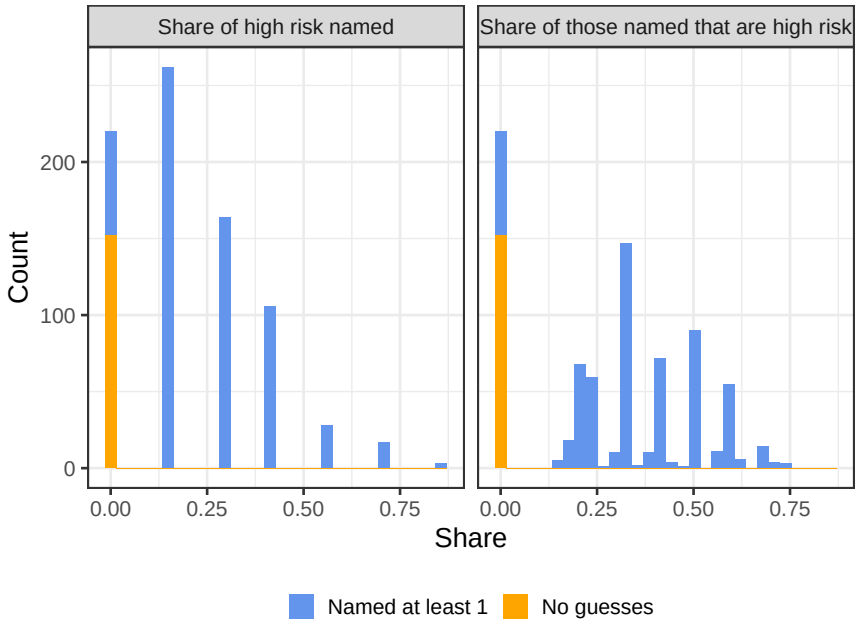
Notes: Figure adapted from data reported in Bagheri-Gavkosh et al. (2021), who survey geophysical studies of land subsidence published before 2021. The median study included in the survey calculated subsidence rates starting in 2001 and ending in 2010. Points are colored by the mean subsidence rate estimated in the study when available, and the maximum rate when that is the only measure reported.

Figure B2: Timeline of satellite data measures



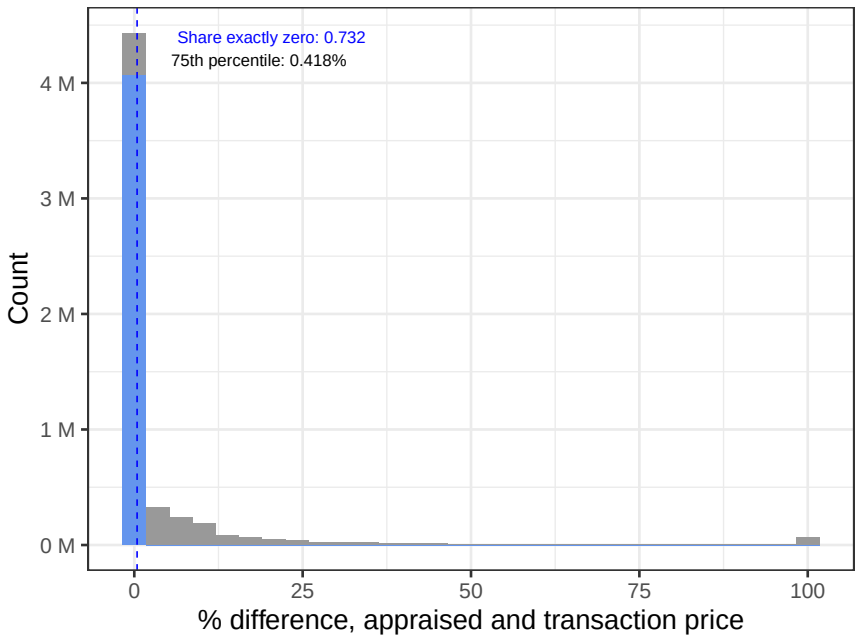
Notes: Chaussard et al. (2021) provided data from 2014-2020. Original interferometry performed from 2008-2011 using ALOS-1 data, and 2021-2024 using Sentinel-1 data.

Figure B3: Evaluating respondent beliefs about the geography of sinking



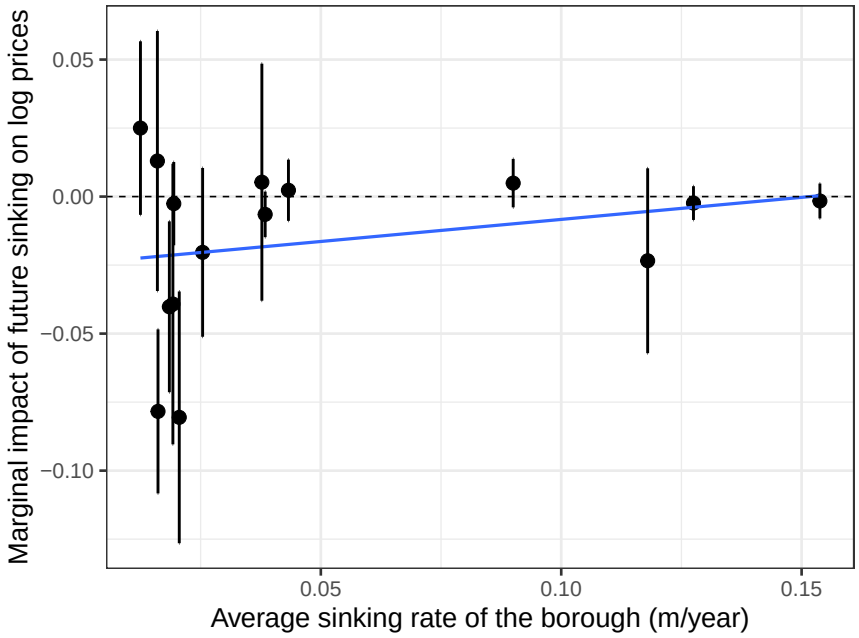
Notes: Each bar plots the number of responses that scored a given "grade" according to two metrics. The right panel scores the share of the seven boroughs considered high sinking risk that the respondent named in response to the question "Which boroughs in Mexico City have issues with subsidence?". The left panel scores the share of named boroughs that are in the severn high risk boroughs. Respondents who could not name any boroughs in response to the question are marked in yellow and are assigned a score of zero.

Figure B4: Distribution of deviations of appraised and transaction values



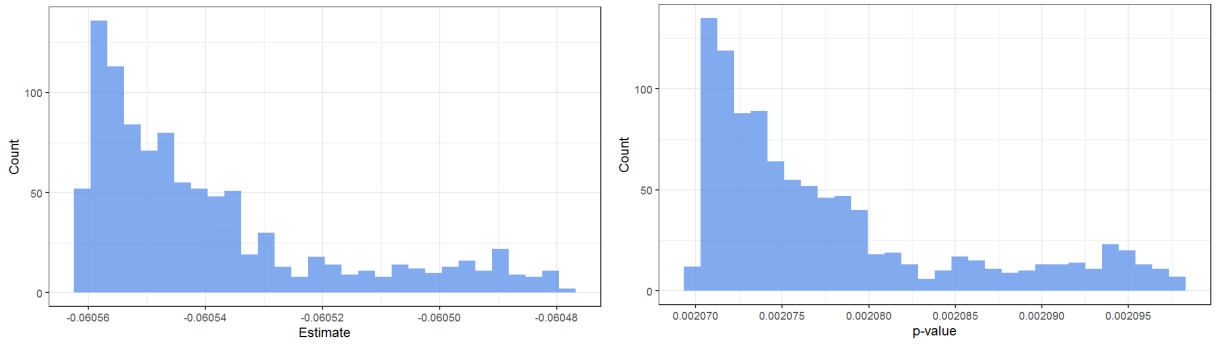
Notes: Height of the bars represents the number of appraisals with the given percentage difference between appraised and transaction values, measured as $100 \times (p^{appr} - p^{transac})/p^{transac}$. Values winzorized at the 99th percentile for visual clarity. Exact zero values colored in blue.

Figure B5: Capitalization of future sinking by borough



Notes: Each point with error bar represents the regression coefficient and 95% confidence interval of the interaction of a borough fixed effect with the net present value of future sinking specified using an AR(1) process and 6% discount rate. Estimation is limited to new builds. All estimates include fixed effects by 500m-grid cell, seismic zone by year, amenity class, number of bathrooms, number of floors, and number of parking spaces. Standard errors clustered at the grid cell. The blue line fits a linear regression through the point estimates.

Figure B6: Monte Carlo estimates simulating measurement error in subsidence

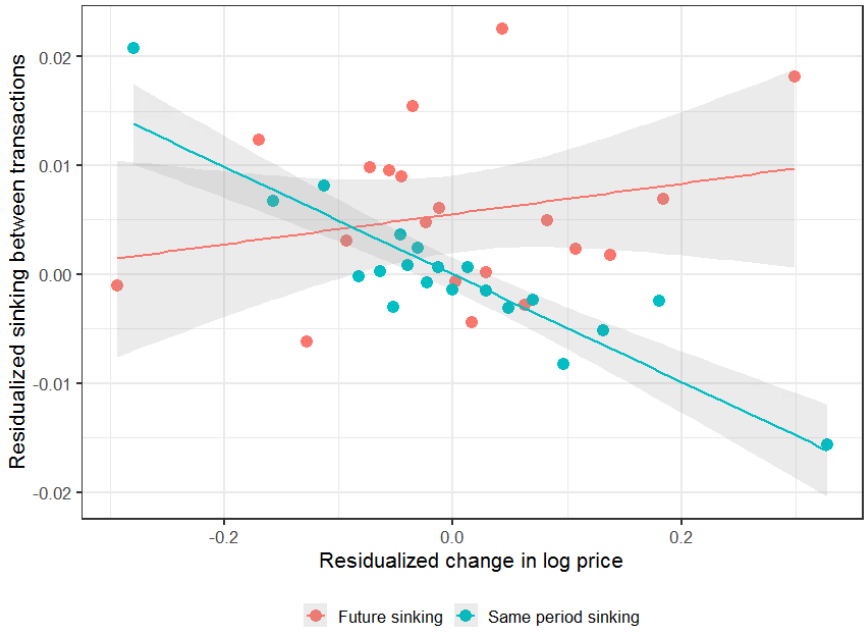


(a) Distribution of Estimates of γ

(b) Distribution of p-values

Notes: Each plot shows the frequency of estimates of γ (Figure (a)) and its associated p-value (Figure (b)) given 1,000 draws of data in which each models total subsidence as $\tilde{S}_{jt} = S_{jt} + \sum_{i=0}^{Age_{jt}} e_{ij}$, where e_{ij} is drawn from a normal distribution with mean zero and standard deviation of 0.01 meters, simulating a 10 millimeter per year white noise measurement error associated with the measured sinking rate. Each point estimate is from a regression of log prices on $\tilde{S}_{jt}$ and property, age, and zone by year of appraisal fixed effects and time trends by borough, number of bedrooms, and type. Estimates exclude properties with only one appraisal. Standard errors clustered at the census tract.

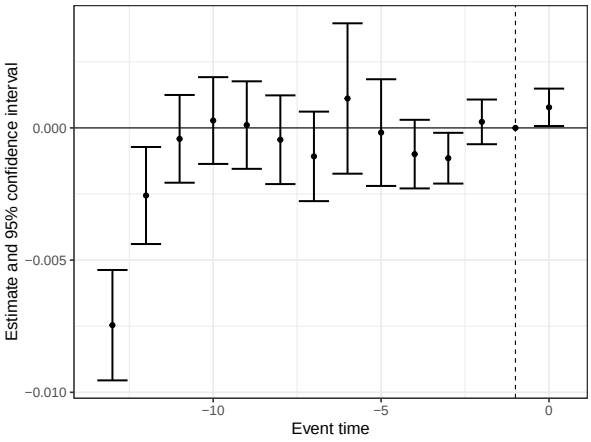
Figure B7: Pre-trends test



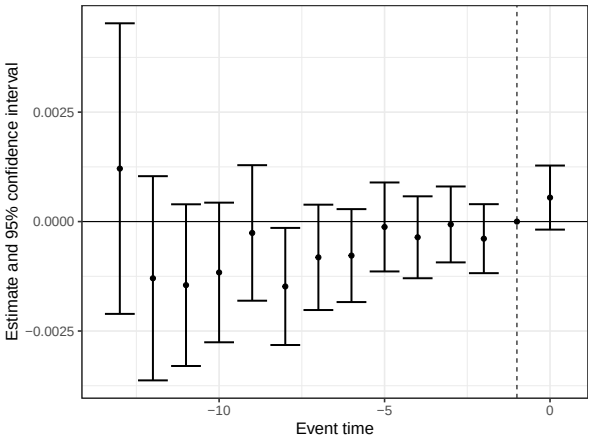
Notes: Points represent mean residualized sinking between transactions for 20 quantile bins of the residualized change in log price between appraisals. Residualization of both variables is done on property, age, and year of appraisal by seismic zone fixed effects. Lines represent the line of best fit through these binned values; grey areas represent 95% prediction intervals. The slopes reported are from a regression of the sinking between appraisals on the change in log price (lagged or contemporaneous), residualized as in the plot. P-values calculated using standard errors clustered at the census tract.

Figure B8: Pre-trends in pixel-level development

(a) Sinking intensity

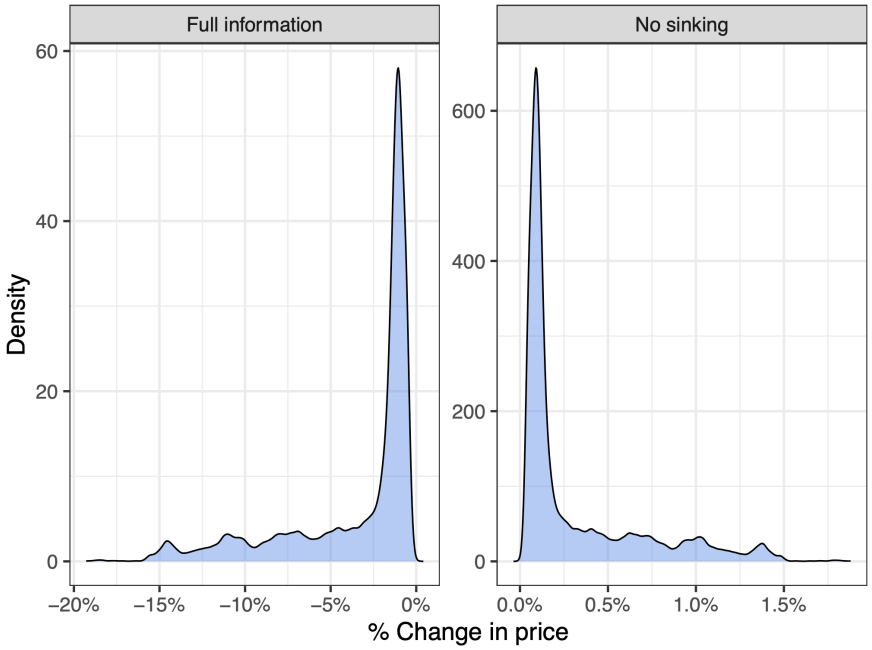


(b) Uneven altitude change intensity



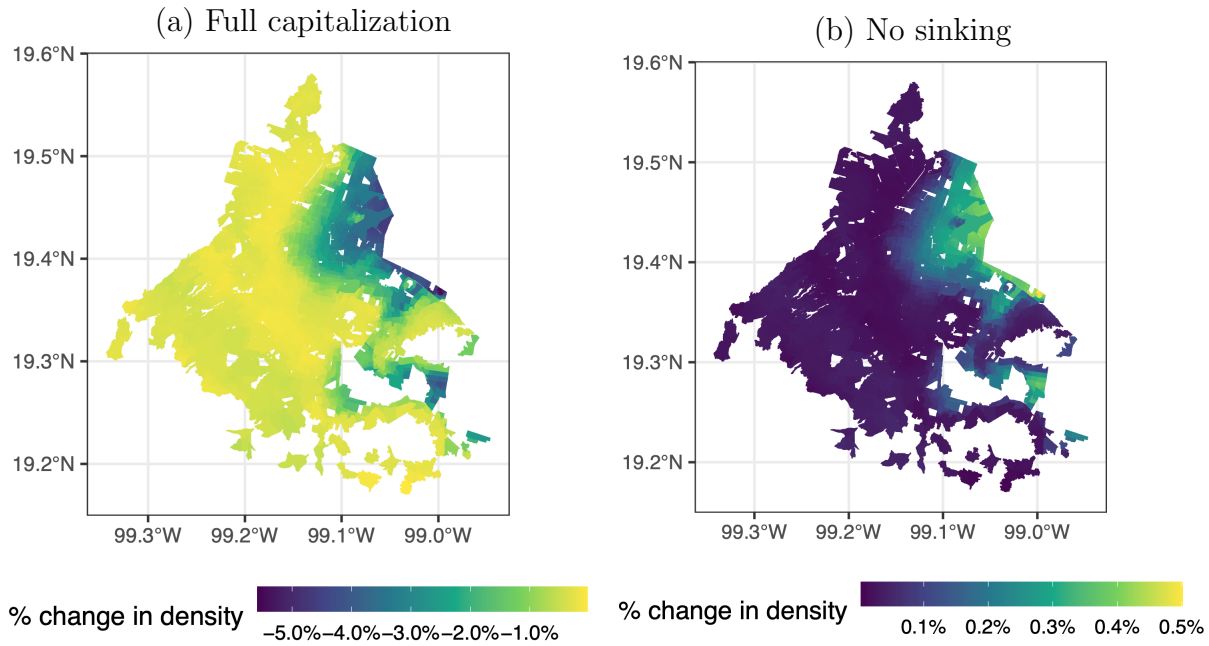
Notes: Depicts point estimates and 95% confidence intervals from a regression of an indicator for any development occurring on the pixel on an indicator for being a high-sinking pixel in 2020 (Panel (a)) or a highly-unevenly sinking pixel as measured by the cumulative relative altitude change (Panel (b)) in different periods. Periods denoted on the horizontal axis are the year relative to the first year that the pixel achieved a sinking intensity above the cut-off. Relative periods <10 are grouped into the -10 period. All regressions include pixel and year fixed effects. Point estimates estimated using Sun and Abraham (2021) staggered event study estimator. Standard errors clustered at the pixel.

Figure B9: Price changes in counterfactuals relative to the observed equilibrium



Notes: Each plot shows the density of the percentage change in price in the counterfactual relative to observed prices. The left panel shows the change between the baseline equilibrium and the equilibrium obtained when sinking remains unchanged but pricing frictions are removed ($\theta = 1$). The right panel shows the change between the baseline equilibrium and the equilibrium obtained when expected future sinking is set to zero.

Figure B10: Change in density under full capitalization relative to the observed equilibrium



Notes: Fill colors per census tract represent the percentage change in the number of units per plot between (a) a counterfactual calculated with no pricing frictions and the observed housing density, and (b) a counterfactual calculated with no future sinking and the observed housing density.

E Data Appendix

E.1 Subsidence data

Subsidence calculated from Sentinel-1 measures (2014-2020) at the pixel level (90 by 90 meter resolution) using the C-band was directly provided by Chaussard et al. (2021). I create an annual pixel-level measure of sinking by mimicking their methodology for creating the average pixel-level analysis reported in their paper by estimating a linear regression of cumulative vertical deformation on time for each pixel-year, the slope of which is an estimate of the sinking in that period. This method is preferred to directly integrating sinking at each time step, which is sensitive to estimates contaminated by error or low coherence (Chaussard et al., 2021). In our setting the linear projection helps remove measurement error that may be related to construction itself, as a change in building height, if detected at all by the satellite, would be smoothed over by the other measurements.

We create annual, pixel-level measures of subsidence for 2007-2011 using ALOS-1 ascending track L-band SAR acquisitions for the Mexico Valley. We process the raw Level-1 data using the ISCE full stack processing framework (Interferometric synthetic aperture radar Scientific Computing Environment) and use MintPy algorithms to correct for tropospheric distortions and unwrapping errors and produce the final corrected time series values (Yunjun et al., 2019) and GACOS to do tropospheric corrections (Yu et al., 2018b,a, 2017). Given the high signal-to-noise ratio in sinking in central Mexico, I do not filter on coherence to maximize the coverage of the data.

Using the same process, we create annual, pixel-level measures of subsidence for 2021-2024 using Sentinel-1 ascending track C-band SAR acquisitions for the Mexico Valley. We process the raw Level-1 data using the ISCE full stack processing framework (Interferometric synthetic aperture radar Scientific Computing Environment) and use MintPy algorithms to correct for tropospheric distortions (using GACOS) and unwrapping errors. I do not filter on coherence.

In order to create a harmonized panel of pixels, I create a regular 100m×100m pixel overlay that spans the two spatial extents of the ALOS and Sentinel data and take an intersection-weighted average of the original pixel subsidence rates. From this grid I interpolate linearly from 2011-2013.

E.1.1 Measurement error

In the direct measures of subsidence produced by SAR interferometry, measurement error arises from four main sources: (i) errors in the characterization of the exact orbit location, which is required to properly compare multiple passes; (ii) errors arising from troposhperic or ionospheric conditions that affect measurement, in particular water vapor; (iii) unwrapping errors, which occur when "unwrapping" the phase length which is modular 2π, and (iv) temporal decorrelation, where measurements jump in the 2π range due to irregular backscatter. We discuss the role for these in our measurements and the magnitude of errors from these sources.

Orbital location errors Errors in characterizing the exact location of the satellite in space at each pass can result in long wavelength phase ramps which essentially represent a

level shift in subsidence that is common to every pixel. These are relatively straightforward to identify in interferograms, and the MintPy algorithm we use to generate the time series implements a linear "de-ramping" of the time series. That said, before de-ramping, the accuracy of orbit information depends on the satellite. The Sentinel-1 orbit information, used in our analysis from 2014-2024, is extremely accurate, and Sentinel's orbit products have an accuracy typically below 5cm in 3D in terms of the exact location of the satellite (Peter et al., 2017). This results in negligible orbit errors in the resulting interferometry process. As an older satellite, the ALOS-1 satellite achieves an accuracty of around 40 centimeters in 3D in terms of the satellite's location in space, so de-ramping is more important for this satellite product.

Tropospheric and ionospheric errors Changes in tropospheric and ionospheric conditions can distort interferograms, which manifests as a phase ramp similar to orbital location errors. We correct for differences in the tropophere by using GACOS hourly temperature, pressure, and water vapor measurements to adjust measurements using MintPy. The magnitude of ionospheric distortions is inversely proportional to the band frequency; in the case of the C-band, which is relatively high frequency, these errors are minimal, on the order of <5 millimeters. For the lower frequency L-band, the ionosphere plays a mroe important role, but its importance is highly dependent on context; ionospheric errors are at their worst during periods of high solar activity, during the day at equitorial and polar locations. The ALOS-1 satellite has a sun-synchronous orbit, so it's ascending orbit passed over Mexico City around 10:30PM local time. The nighttime pass together with the fact that Mexico City is near mid-latitude ranges at 19° latitude means that ionospheric errors are somewhat mitigated, and our de-ramping also decreases ionospheric delay (Chaussard et al., 2021).

Unwrapping errors Unwrapping errors result from too large of a phase difference between passes, and rather than affecting measurements directly it affects the ability to produce a measurement at all in some places. Unwrapping errors can arise from several sources; those most relevant to our setting include (i) too fast of ground deformation; and (ii) steep slopes that mean that measurement is sensitive to the exact line-of-sight of the satellite. Our area of interest does not have particularly steep slopes, as urbanized Mexico City is located within a valley. How fast is "too-fast" of deformation depends on the band, as the phase is wrapped around half the radar's wavelength. For the Sentinel-1 C-band, this means that if displacement exceeds 2.8 centimeters between passes, which occur every 6 days, then displacement would not be able to be estimated. Thus, subsidence would need to exceed a rate of 1.7 meters per year to be undetectable, a rate observed for dramatic events such sudden, discrete events such as volcanic eruptions or ice sheet collapses, but not observed in our sample. The L-band measures we use from the ALOS-1 satellite are able to detect up to 11.8 centimeters between passes, and our average pass frequency is 90 days, so sinking would have to exceed 47.9 centimeters per year to be undetectable, which is almost double the maximum rate observed in our data.

Temporal decorrelation The primary source of temporal decorrelation is changes in the ground surface that cause measurements to behave differently between passes, for example changes in vegetative density or snow cover. These sources of error are minimized in urban settings, where concrete provides a time-consistent radar signature. This is diagnosable using the average coherence in the signal between satellite passes, which when

limiting to Mexico City in our data is over 0.9 on a scale from 0 to 1 almost everywhere.

Summing up By focusing on an urban setting and using state-of-the-art tools for conducting interferometry with tropospheric corrections and linear de-ramping, we minimize the role for measurement error in our analysis. Other studies in urban areas use GPS stations that provide ground measures at discrete points, and find measurement errors on the order of less than millimeters per year (see: 1.5mm/year in Beijing (Liu et al., 2020); 9mm/year in the Antelope Valley, CA (Bawden et al., 2023)). Assuming we achieve similar accuracy to the Beijing study (1.5mm), as this more closely mirrors our urban setting, our measurement error would be an order of magnitude of around 2.1% of the average sinking rate of 7 centimeters/year.

E.1.2 Does development cause subsidence?

While historic buildings in Mexico City such as the Palace of Fine Arts have been known to cause sinking because of their own weight, modern construction code requirements seek to prevent this through soil studies, including underground levels that offset the weight of the building, and using pylons that "anchor" the building into the ground instead of sitting on top of the soil.

We test whether development causes subsidence by estimating the following specification:

$$s_{pt+\tau} = \beta 1\left(\text{Development}_{pt}\right) + \theta s_{pt-1} + \alpha_p + \alpha_{tg(p)} \tag{20}$$

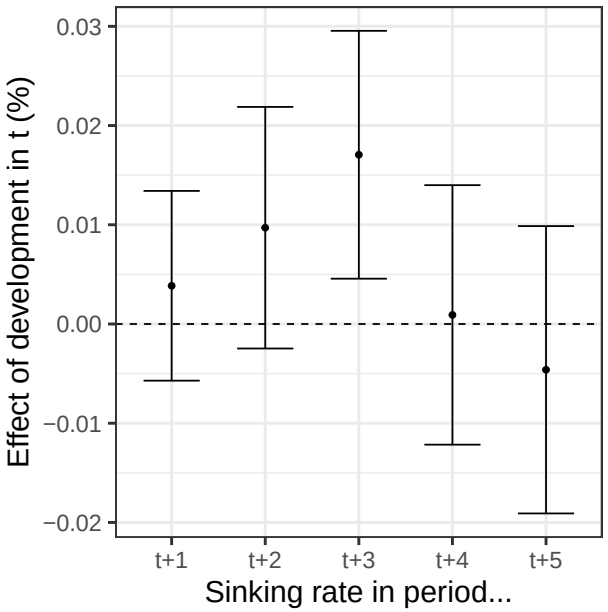
where $s_{pt+\tau}$ is the sinking rate on pixel p in period $t+\tau$ and $1\left(\text{Development}_{pt}\right)$ is an indicator for a new development occurring in pt. s_{pt-1} is the lagged sinking rate, a necessary control given that sinking leads to more development. We also control for pixel fixed effects α_p and seismic zone by year fixed effects $\alpha_{g(p)t}$. To avoid confounding the effect of multiple developments, we limit estimation to the sample of pixels that had either one or no developments over the period. Standard errors are clustered at the pixel.

Figure B11 reports $\widehat{\beta}$ for $\tau \in [1,5]$, transformed to be reported as a percent of the average sinking rate. Most coefficients are statistically insignificant, except for sinking three years after construction ended. However, the value of the coefficient is economically minor; a development in t increases the rate of sinking in $t+3$ by less than 2%. This is generally consistent with studies from Mexico City and elsewhere that find that the role of human development in causing land subsidence is minor (Chaussard et al., 2021; Parsons et al., 2023). The fact that the coefficients are small and largely insignificant also assuages concerns that construction impedes measurement of subsidence by masking as positive elevation gain.

E.2 Survey design and implementation

The survey was designed to be representative of Mexico City housing units, with stratification on two broad zones of high (7 boroughs) and low (9 boroughs) sinking from 2007-2020 defined at the borough level. Within strata, census tracts were sampled with probabilities drawn from the number of housing units in that tract in the 2020 census; 80 tracts were randomly sampled, and 10 households were surveyed per tract. IRB exemption was obtained through the UC Berkeley Office for the Protection of Human Subjects.

Figure B11: Impact of development on future sinking rates



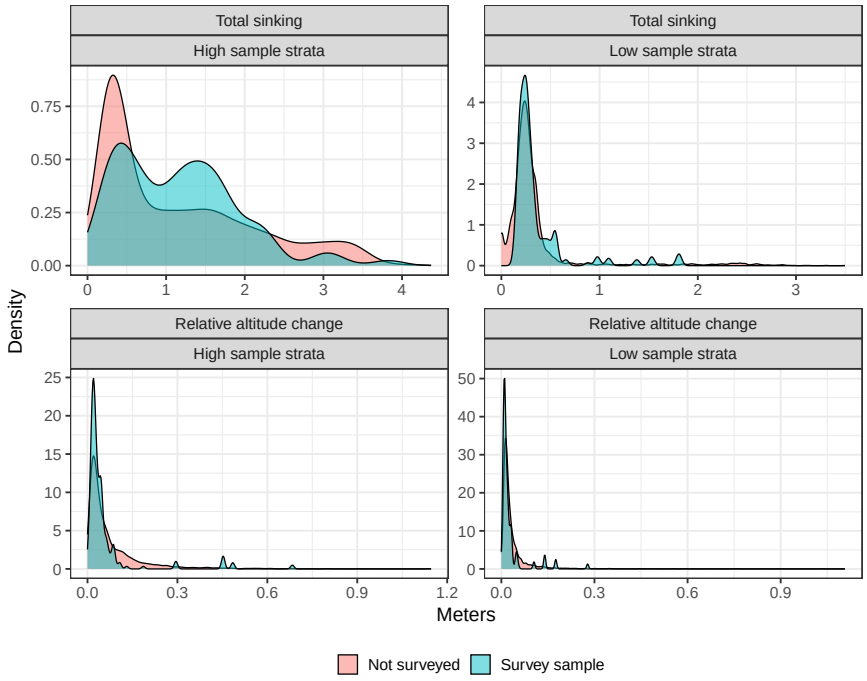
Notes: Plots $\widehat{\beta}/\overline{s}$, where $\overline{s}$ is the average sinking rate over the panel, and 95% confidence interval from estimating Equation (20) for $\tau \in [1, 5]$. Standard errors clustered at the pixel.

The surveyed homes provide good representation of sinking across the city; Figure B12 reports the distribution of average total sinking and total relative altitude loss from 2008-2020 for surveyed and non-surveyed tracts. The distributions are very similar, with some more mass in the surveyed tracts at mid-levels of sinking coming from the distribution of housing units over space within high-risk boroughs.

Table A16 reports comparisons of our resident and unit outcomes with the 2022 Income and Expenditure Survey (ENIGH). Despite the three-year difference and different sampling strategy (housing units versus residents), outcomes are broadly comparable though statistically differentiable between our survey and the ENIGH. Given the sample restriction on homeowners of any tenure and renters with at least 5 years in the unit, our sample is 73% homeowners (compare to 60% in the ENIGH). Looking within homeowners in Table A16, we see that survey respondents are more likely to inhabit single-family homes and live in self-built homes as compared to commercially built homes. Respondents are also significantly older on average. These differences likely highlight selection into being contactable by the survey firm. Reassuringly, our respondents are not meaningfully more educated on average; a 0.58 difference in educational attainment represents an approximate difference of 0.87 years of schooling.[60]

[60]Schooling is coded as a numeric factor representing categories such as "Middle school incomplete", "Middle school complete", "High school incomplete", etc. Given that the modal respondent completed at least middle school and that middle and high school in Mexico each last 3 years, a each unit increase in the education variable represents approximately 1.5 years of schooling.

Figure B12: Distribution of subsidence in surveyed and non-surveyed census tracts



Notes: Each plot is the distribution of average total sinking from 2008-2020 at the census tract level and the total relative altitude change from 2008-2020 on surveyed and non-surveyed tracts, reported separately for the high-sinking strata (N=500) and the low-sinking strata (N=300).

Table A16: Balance of survey outcomes against the 2022 ENIGH

		Full sample		Owners only			
		ENIGH	Our survey	ENIGH	Our survey	Diff, full sample	Diff, owners sample
Unit characteristics							
Single-family home	Mean	0.66	0.78	0.70	0.79	0.12***	0.09***
	SD	(0.437)	(0.412)	(0.409)	(0.401)	(0.019)	(0.022)
	N	2541	800	1595	581	3341	2176
Apartment	Mean	0.31	0.22	0.27	0.21	-0.09***	-0.06***
	SD	(0.422)	(0.412)	(0.392)	(0.401)	(0.019)	(0.022)
	N	2541	800	1595	581	3341	2176
Age of home	Mean	26.82	38.82	30.39	42.28	4.5***	8.27***
	SD	(20.285)	(22.713)	(19.471)	(22.387)	(0.96)	(1.106)
	N	2092	800	1471	581	2892	2052
Self-built	Mean	0.28	0.63	0.47	0.67	0.16***	0.2***
	SD	(0.49)	(0.483)	(0.49)	(0.47)	(0.023)	(0.025)
	N	1597	800	1597	581	2397	2178
Commercially built	Mean	0.26	0.22	0.44	0.23	-0.22***	-0.2***
	SD	(0.465)	(0.409)	(0.465)	(0.42)	(0.021)	(0.023)
	N	1597	800	1597	581	2397	2178
Respondent characteristics							
Female	Mean	0.54	0.57	0.54	0.56	0.04***	0.02***
	SD	(0.499)	(0.496)	(0.499)	(0.498)	(0.019)	(0.022)
	N	8421	800	5449	581	9221	6030
Age (respondent)	Mean	38.45	52.94	41.63	53.73	14.55***	12.11***
	SD	(21.547)	(16.762)	(22.135)	(17.086)	(0.671)	(0.818)
	N	8421	800	5449	581	9221	6030
Education	Mean	4.23	6.12	4.44	6.24	0.71***	0.58***
	SD	(3.541)	(3.081)	(3.563)	(3.118)	(0.123)	(0.148)
	N	6505	800	4221	581	7305	4802

Notes: Means and standard deviations are weighed by sample weights to the population. Differences are conditional on strata fixed effects and use heteroskedasticity-robust standard errors. Stars represent p-values less than 0.10 ($*$), 0.05 ($**$), and less than 0.01 ($***$).

E.3 Appraisal data

E.3.1 Appraised useful life remaining

We investigate whether appraised values for the years of useful life remaining follow systematic rules that would preclude us from exploiting true variation in the home's physical state. If appraisers use fixed rules based on some property characteristic or age to assess the useful life left, then regressions that include property and year of appraisal by year of last appraisal fixed effects would control for all possible variation in this variable.

To investigate this, we define the implied duration of the property's useful lifespan by adding the age at time of appraisal to the appraised years of useful life remaining. Then, we investigate the extent to which the modal value of these implied useful lifespans by the property's borough, type (basic, medium, luxury amenities), vintage, and year of appraisal explains the appraised useful lifespan. We find that there are systematic patterns in the modal appraised useful life, with different boroughs and vintages receiving common appraised useful lifespans of 65, 70, 75, or 80 years. However, there are also substantial deviations from these rules of thumb.

We compute the cell-specific modal useful lifespan for cells defined by the property's location, type, vintage, and year of appraisal, and plot the deviation of the implied useful lifespan from this modal value in Figure B13. While differences are centered around zero, there are appraised useful lifespans that differ substantially from the rule of thumb value. If we define the rule of thumb excluding vintage, we find that newer homes appear to be built with longer useful lives, a finding consistent with technological improvements over time that increase the durability of homes.

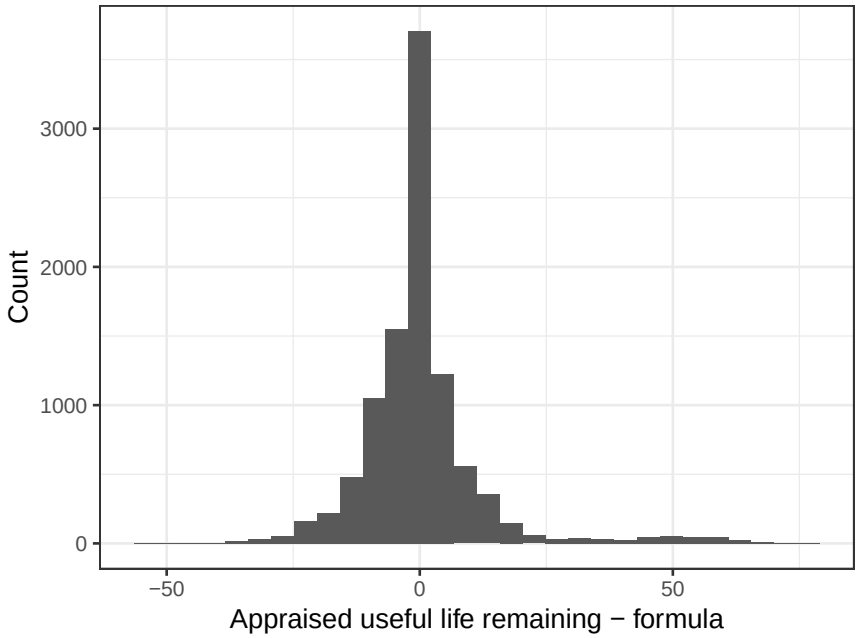
E.4 Housing developer data

We aggregate the quarterly development panel into a dataset of developments by calculating averages of characteristics of the modal unit and its asking price that are weighted by the number of units reported sold in that period. To create a dataset of new units developed, we expand the data by the number of units sold, assigning to each the modal price and characteristics reported in that quarter.

E.4.1 Coverage of the new housing market

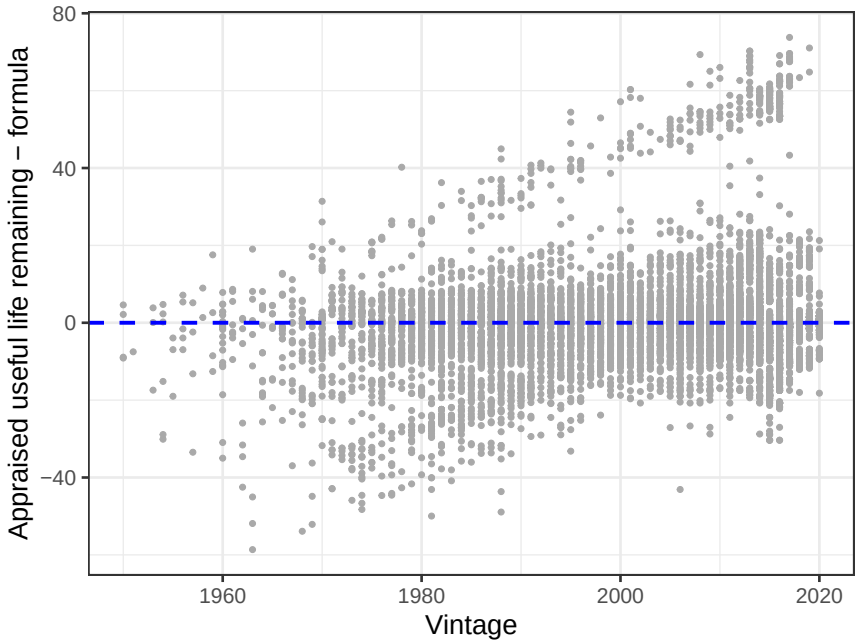
Across Mexico, the Central Bank reports that the DIME should cover about 60% of new homes based on the fact that developers account for 60% of new unit construction. We compare the units reported built in the DIME to the number of appraisals per year to get a better idea of the DIME coverage in Figure B15. Units reported sold by the DIME surpass even the total number of appraisals in some years, and is consistently around double the number of appraisals for new builds (blue line). This is likely because the appraisal data is missing transactions that did not involve a mortgage, a significant share (around 70%). While conclusions about coverage should be tempered by this fact, we see this as suggestive that the DIME covers a significant share of new development in Mexico City.

Figure B13: Deviations of appraised useful lifespan from rule of thumb values



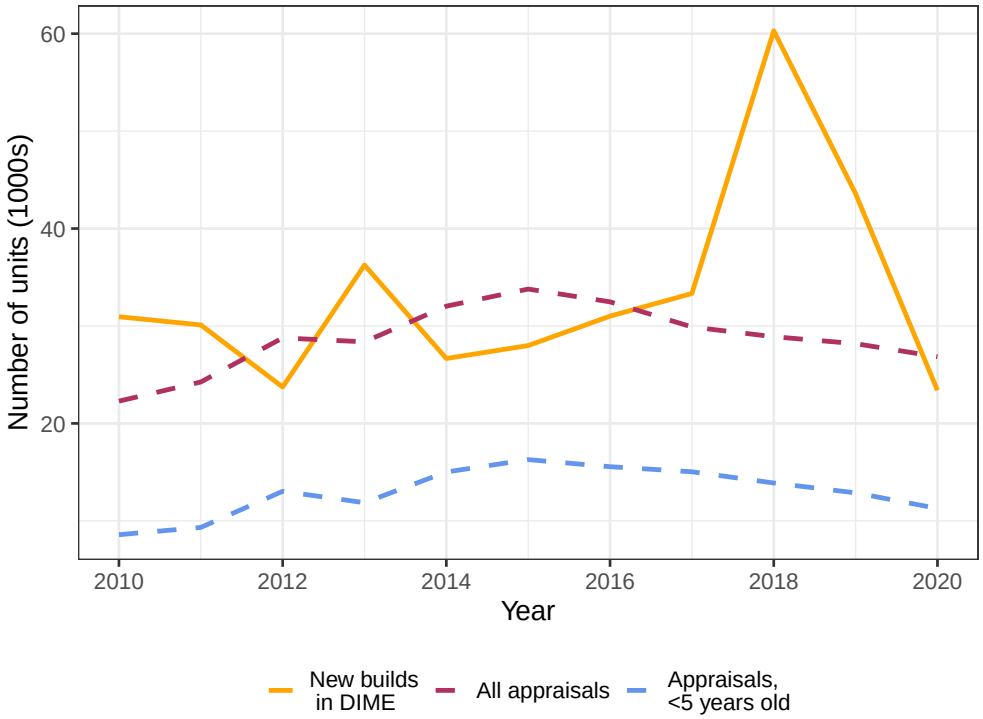
Notes: Histogram of the appraised useful lifespan minus the implied rule of thumb using a 10,000-observation random sample from the data. The rule of thumb is defined as the modal appraised useful lifespan for properties that are located in the same borough, built in the same year, appraised in the same year, and are of the same type.

Figure B14: Deviations of appraised useful lifespan from rule of thumb values



Notes: The rule of thumb is defined as the modal appraised useful lifespan for properties that are located in the same borough, appraised in the same year, and are of the same type, but does not control for vintage.

Figure B15: Total number of newly developed units and appraisals



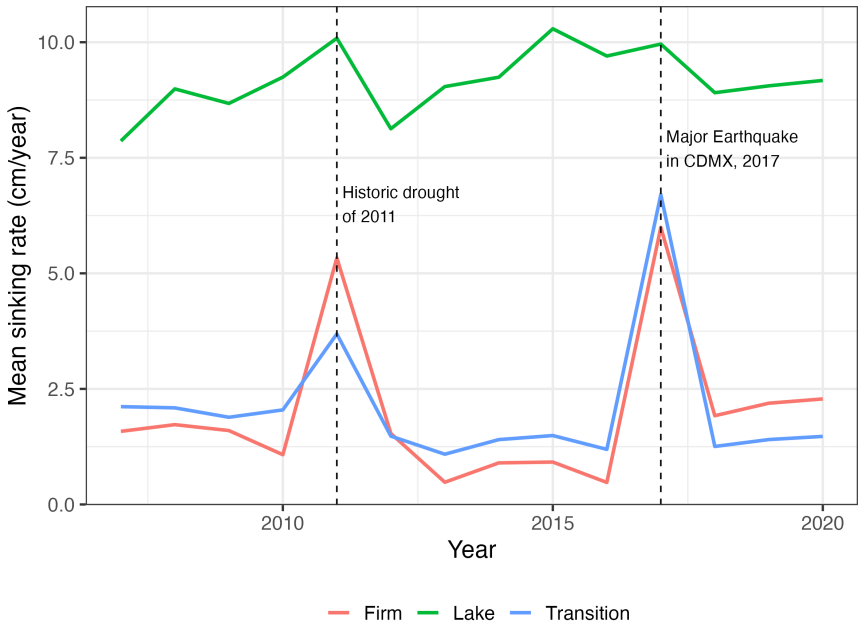
Plots the total number of units per year. The solid orange line plots the total number of units reported sold by the DIME. The dashed lines represent the total number of appraisals (maroon) and the total number of appraisals that are reported less than 5 years since construction (blue).

F Predictability of subsidence over time and space

As seen in Figure 2, subsidence intensity is very serially correlated over space, with the fastest rates of sinking concentrated over the historic lake and the surrounding "transition" zone defined in the construction code. Subsidence is also highly serially correlated; at the pixel level, the lagged subsidence rate has a correlation of 0.85 with the contemporaneous subsidence rate.

However, how sinking has evolved over time varies substantially in Mexico City. First, how sinking rates respond to shocks that affect groundwater levels or geophysical stability depends on zone. Figure B16 plots subsidence rates over time by broad categories of the geophysical environment defined in the construction code. While subsidence rates are consistently high in the lake zone, the more resistant Firm and Transition zones are more responsive to shocks; each experiences a spike in subsidence rates during the drought of 2011 and the large earthquake of 2017. In particular after the earthquake, rates in the Firm zone experienced a level shift up in the sinking rate.

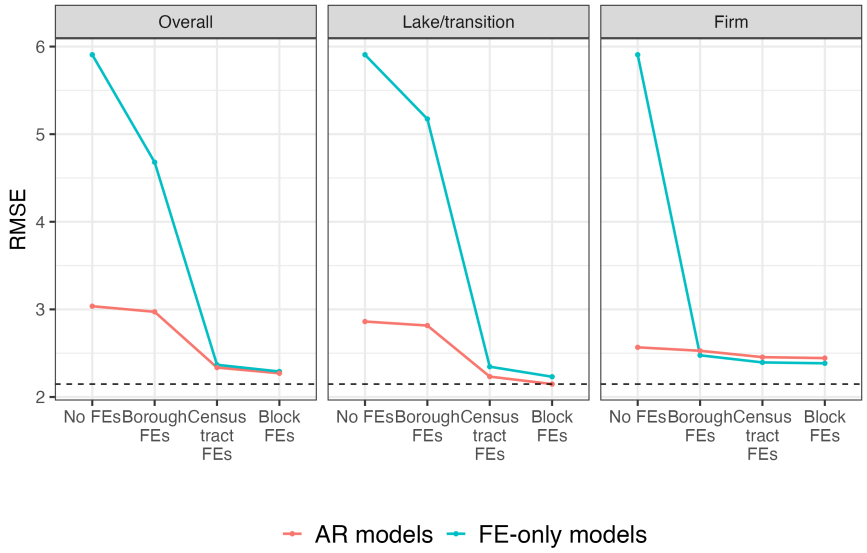
Figure B16: Subsidence rates over time



Notes: Lines represent simple averages of subsidence rates by year and seismic zone as defined by the construction code.

While past sinking is strongly correlated with current sinking (Table A17), it does not provide much additional predictive power conditional on fine enough geographic information. Figure B17 compares the RMSE for models with different levels of granularity of spatial fixed effects that either include or exclude the lagged subsidence rate as a predictor of the contemporaneous rate. This highlights that starting around the census tract level, lagged subsidence contains little additional information above this geographic average.

Figure B17: Predictive models of subsidence rates



Notes: Points plot the RMSE of regressions that include different levels of fixed effects and autoregressive terms. Points and lines in blue include only the fixed effects indicated on the horizontal axis, which range from no fixed effects (simple intercept model) to city block-level fixed effects. Pink points and lines include these same fixed effects, and additionally the lagged sinking rate.

Table A17: AR(1) predictive models of the subsidence rate

	Overall				Lake/transition				Firm			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Intercept	0.698***				0.621***				1.801***			
	(0.004)				(0.006)				(0.006)			
Lagged sinking rate	0.877***	0.797***	0.151***	0.095***	0.922***	0.863***	0.232***	0.152***	0.117***	0.083	0.023*	0.015***
	(0.001)	(0.056)	(0.012)	(0.002)	(0.001)	(0.031)	(0.014)	(0.003)	(0.002)	(0.085)	(0.014)	(0.003)
Observations	821.716	821.716	821.716	821.716	507.539	507.539	507.539	507.539	314.177	314.177	314.177	314.177
RMSE	3.04	2.97	2.34	2.27	2.86	2.82	2.23	2.15	2.57	2.53	2.45	2.44
R squared	0.74	0.75	0.85	0.86	0.81	0.82	0.89	0.89	0.01	0.04	0.10	0.11
Borough FEs		X				X				X		
Census tract FEs			X				X				X	
Block FEs				X				X				X

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

Notes: The unit of observation is a city block-year. Standard errors in parentheses assume iid errors.

G Theory Appendix

G.1 Linearization of expectation terms

First, consider a first-order approximation of $g(S_{j\tau}^{t+1}) = \exp(\gamma S_{j\tau}^{t+1})$ around $S_{j\tau}^{t+1} = \mathbb{E}(S_{j\tau}^{t+1})$:

$$\begin{aligned} g(S_{j\tau}^{t+1}) &= \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1})) + \gamma \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1})) \times (S_{j\tau}^{t+1} - \mathbb{E}(S_{j\tau}^{t+1})) + R(S_{j\tau}^{t+1}) \\ &\Rightarrow \mathbb{E}(g(S_{j\tau}^{t+1})) = \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1})) + \mathbb{E}(R(S_{j\tau}^{t+1})) \end{aligned}$$

The remainder term $R(S_{j\tau}^{t+1})$ is a function of the sum of higher-order moments, each of size:

$$\frac{\gamma^n}{n!}(S_{j\tau}^{t+1} - \mathbb{E}_t(S_{j\tau}^{t+1}))^n$$

These terms will be small if the moments of the distribution of sinking grow at a slower rate than $\frac{\gamma^n}{n!}$. As an example, we can examine the second-order approximation, which include the variance:

$$\begin{aligned} g(S_{j\tau}^{t+1}) \approx & \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1})) + \gamma \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1}))(S_{j\tau}^{t+1} - \mathbb{E}(S_{j\tau}^{t+1})) \\ & + \frac{\gamma^2}{2} \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1}))(S_{j\tau}^{t+1} - \mathbb{E}(S_{j\tau}^{t+1}))^2 \\ & \Rightarrow \mathbb{E}(g(S_{j\tau}^{t+1})) \approx \exp(\gamma \mathbb{E}(S_{j\tau}^{t+1})) + \frac{\gamma^2}{2} \exp(\gamma \mathbb{E}_t(S_{j\tau}^{t+1})) Var(S_{j\tau}^{t+1}) \end{aligned}$$

This is small in the time series of plot-level sinking.

Our second step, after implementing the above approximation, us to linearize the log term around these expectations. Consider a first-order Taylor series approximation of:

$$\begin{aligned} & f(\mathbb{E}_t(\{\tilde{\lambda}_{ju}\}), \mathbb{E}_t(\{e_{ju}\}), \mathbb{E}_t(\{s_{ju}\}_u)) = \\ & \log\left(\exp(-\tilde{\lambda}_{jt} + e_{jt}) + \sum_{\tau=t+1}^{\infty} \rho^{\tau-t} \exp(\gamma \mathbb{E}_t\{S_{j\tau}^{t+1}\} - \mathbb{E}_t(\tilde{\lambda}_{j\tau}) + \mathbb{E}_t(e_{j\tau}))\right) \\ & = \log\left(\exp(-\tilde{\lambda}_{jt} + e_{jt}) + \sum_{\tau=t+1}^{\infty} \rho^{\tau-t} \exp(-\mathbb{E}_t(\tilde{\lambda}_{j\tau}) + \mathbb{E}_t(e_{j\tau}) + \sum_{u=t+1}^{\tau} \gamma \mathbb{E}_t\{s_{ju}\})\right) \end{aligned}$$

around $\tilde{\mathbb{E}}_t(s_{ju}) = 0$, $\mathbb{E}_t(e_{ju}) = e_{jt}$, and $\mathbb{E}_t(\tilde{\lambda}_{ju}) = \lambda_{jt}$ $\forall u > t$. Note this is analogous to assuming that these variables follow a random walk. In the case of e_{jt} this reduces to assuming a distribution for these shocks. In the case of the $\tilde{\lambda}_{jt}$ this approximation comes at a cost in terms of model consistency, as these relative choice probabilities are equilibrium objects that have a distribution that is not likely to be a random walk.[61] However, this approximation reflects the fact that $\tilde{\lambda}_{j\tau}$ is fixed until re-development; because re-development is rare, with a per-period probability of 2.7%, $\mathbb{E}_t(\tilde{\lambda}_{j\tau}) = Pr(d_{jt} = 0)\lambda_{jt} +$

[61] This distribution is difficult to characterize for a given j, as it is a convolution of the distribution of cost shocks and a non-linear function of the random variables in demand that affect prices.

$Pr(d_{jt}=1)\mathbb{E}_t(\tilde{\lambda}_{j\tau}|d_{jt}=1) \approx \lambda_{jt}$.

$$\begin{aligned} f(\{\mathbb{E}_t(\tilde{\lambda}_{ju})\},\{\mathbb{E}_t(e_{ju})\},\{\mathbb{E}_t(s_{ju})\}) \approx & f(\tilde{\lambda}_{jt}, e_{jt}, \vec{0}) \\ & + \sum_{u=t+1}^{\infty} \frac{\partial f}{\partial \tilde{\lambda}_{ju}}\Big|_{\tilde{\lambda}_{jt},e_{jt},\vec{0}} \times (\mathbb{E}_t(\tilde{\lambda}_{ju}) - \lambda_{jt}) \\ & + \sum_{u=t+1}^{\infty} \frac{\partial f}{\partial e_{ju}}\Big|_{\tilde{\lambda}_{jt},e_{jt},\vec{0}} \times (\mathbb{E}_t(e_{ju}) - e_{jt}) \\ & + \sum_{u=t+1}^{\infty} \frac{\partial f}{\partial \mathbb{E}_t(s_{ju})}\Big|_{\tilde{\lambda}_{jt},e_{jt},\vec{0}} \times (\mathbb{E}_t(s_{ju}) - 0) \end{aligned}$$

Note that:

$$\begin{aligned} f(\tilde{\lambda}_{jt}, e_{jt}, \vec{0}) &= \log\left(\frac{\exp(e_{jt}-\tilde{\lambda}_{jt})}{1-\rho}\right) \\ \frac{\partial f}{\partial \tilde{\lambda}_{ju}}\Big|_{\tilde{\lambda}_{jt},e_{jt},\vec{0}} &= -\left(\frac{1-\rho}{\exp(e_{jt}-\tilde{\lambda}_{jt})}\right)\rho^{u-t}\exp(e_{jt}-\tilde{\lambda}_{jt}) = -(1-\rho)\rho^{u-t} \\ \frac{\partial f}{\partial e_{ju}}\Big|_{\tilde{\lambda}_{jt},e_{jt},\vec{0}} &= \left(\frac{1-\rho}{\exp(e_{jt}-\tilde{\lambda}_{jt})}\right)\exp(e_{jt}-\tilde{\lambda}_{jt})\rho^{u-t} = (1-\rho)\rho^{u-t} \\ \frac{\partial f}{\partial \tilde{\mathbb{E}}_t(s_{ju})}\Big|_{\tilde{\lambda}_{jt},e_{jt},\vec{0}} &= \left(\frac{1-\rho}{\exp(e_{jt}-\tilde{\lambda}_{jt})}\right) \times \exp(e_{jt}-\tilde{\lambda}_{jt})\frac{\gamma\rho^{u-t}}{1-\rho} = \gamma\rho^{u-t} \end{aligned}$$

The last line follows from the fact that a given $\mathbb{E}_t(s_{ju})$ term appears in every (discounted) term starting in period u to infinity.

Thus:

$$\begin{aligned} f(\{\mathbb{E}_t(\tilde{\lambda}_{ju})\},\{\mathbb{E}_t(e_{ju})\},\{\tilde{\mathbb{E}}_t(s_{ju})\}) \approx & \; e_{jt} - \tilde{\lambda}_{jt} - \log(1-\rho) \\ & + (1-\rho)\sum_{\tau=t+1}^{\infty}\rho^{\tau-t}\left(\mathbb{E}_t(e_{j\tau}) - e_{jt} - (\mathbb{E}_t(\tilde{\lambda}_{j\tau}) - \tilde{\lambda}_{jt})\right) \\ & + \gamma\sum_{\tau=t+1}^{\infty}\rho^{\tau-t}\tilde{\mathbb{E}}_t(s_{j\tau}) \end{aligned}$$

G.2 Calculating expectations of future sinking

We create two measures of expectations of future sinking. The first assumes that sinking rates follow an AR(1) process at the plot level, such that:

$$s_{jt+1} = \mu_j + \phi s_{jt} + e_{jt+1} \tag{21}$$

where $\mu_j = (1-\phi)\overline{s}_j$ We estimate the parameters of this function, namely, $\{\mu_j\}_j$ and ϕ, and

define the expected value of the sinking rate in each period as:

$$\mathbb{E}_t(s_{jt+\tau}) = \mu_j(1-\phi^{\tau-t}) + \phi^{\tau-t}s_{jt} \tag{22}$$

with the expected value of total sinking in τ being the sum of the expectations of sinking rates from $t+1$ to τ. Solving for the net present value of these terms from $t+1$ to infinity gives:

$$\sum_{\tau=t+1}^{\infty} \rho^{\tau-t}\mathbb{E}_t(s_{j\tau}) = \frac{\rho}{1-\rho}\mu_j + \frac{\rho\phi}{1-\rho\phi}(s_{jt}-\mu_j) \tag{23}$$

For expectations that follow perfect foresight, we use the realized sinking on the plot from $t+1$ through 2024, and use a plot-specific estimate of the rate from 2025-on for sinking after that date. The expectation term under perfect foresight is:

$$\sum_{\tau=t+1}^{\infty} \rho^{\tau-t}\mathbb{E}_t(s_{j\tau}) = \sum_{\tau=t+1}^{2024} \rho^{\tau-t}s_{j\tau} + \frac{\rho^{2025-t}}{1-\rho}\widehat{s}_j \tag{24}$$

G.3 Welfare calculations

G.3.1 Renters

Per-period surplus per renter in the housing unit market is:

$$CS_j^i = R_j(0)\log\left(\sum_k \exp\{-\log(R_k(0)) + Z_k\delta + e_k\}\right) + C$$

In this expression, the marginal utility of money (rent, evaluated in the initial period) is $\frac{\partial V_j}{\partial R_j(0)} = -1/R_j(0)$. We measure the welfare loss to consumers Δ^D as the equivalence variation required to compensate home buyers for the impacts of subsidence relative to the benchmark. Because of the size of the system, it is too computationally intensive to calculate an equivalent variation for each consumer, so we calculate the average equivalent variation as:

Δ_t^D such that:

$$\frac{1}{J}\sum_j CS_{jt}^i = \frac{1}{J}\sum_j R_j(0)\log\left(\sum_j \exp\{-\log(R_j(0)+\Delta_t^D) + Z_j\delta + e_{jt}\}\right) \tag{25}$$

While a building is not redeveloped, rents must adjust so that equilibrium demand on the plot is equal to the (fixed) supply of housing on the plot, so we can express mean utility in each period as a function of utility in the period when the building is new ($S_{jt}=0$) with rental price $R_j^*(0)$.

G.3.2 Property managers

Expected manager surplus is their profits under rational expectations:

$$\Pi_t = \sum_j H_{jt} \left(\mathbb{E}(\sum_{\tau=t}^{\infty} \rho^{\tau-t} R_{j\tau}) - p_{jt} \right) \tag{26}$$

G.3.3 Producer surplus

Per-period producer surplus in the housing unit market is:

$$\begin{aligned} PS_t(\{S\}_j;\theta) =& \mathbb{E}\left(\pi^*_{jt}\right) \\ =& Pr(\exp(\tilde{\xi}_{jt})\tilde{\eta} p_{jt}(0)^{\eta+1} - p_{jt} H_{jt-1} D > 0) \times \\ & \mathbb{E}(\exp(\tilde{\xi}_{jt})\tilde{\eta} p_{jt}(0)^{\eta+1} - p_{jt} H_{jt-1} D | \exp(\tilde{\xi}_{jt})\tilde{\eta} p_{jt}(0)^{\eta+1} - p_{jt} H_{jt-1} D > 0) \end{aligned}$$

with expectations defined over cost shocks. The firm earns positive profits for:

$$\begin{aligned} &\tilde{\xi}_{jt} > -\left(\log(\tilde{\eta}) + (\eta+1)\log(p_{jt}(0)) - \log(p_{jt}) - \log(H_{jt-1}) - \log(D)\right) \iff \\ &\tilde{\xi}_{jt} > -\left(\log(\tilde{\eta}) + \eta \log(p_{jt}) - \gamma(1+\eta) S_{jt} - \log(H_{jt-1}) - \log(D)\right) := -\tilde{\pi}_{jt} \\ &\quad \iff \xi^*_{jt} > \exp(-\tilde{\pi}_{jt}), \qquad \xi^*_{jt} = \exp(\tilde{\xi}_{jt}) \end{aligned}$$

So the expected value from the firm's perspective, who takes prices as given, is:

$$\begin{aligned} \mathbb{E}(\pi^*_{jt}) &= Pr(\xi^*_{jt} > \exp(-\tilde{\pi}_{jt})) \frac{\tilde{\eta} p_{jt}(0)^{\eta+1} \int_{\exp(\tilde{\pi}_{jt})}^{\infty} \xi^*_{jt} f(\xi^*_{jt}) d\xi^*_{jt} - p_{jt} H_{jt-1} D}{Pr(\xi^*_{jt} > \exp(-\tilde{\pi}_{jt}))} \\ &= \tilde{\eta} p_{jt}(0)^{\eta+1} \int_{\exp(-\tilde{\pi}_{jt})}^{\infty} \xi^*_{jt} f(\xi^*_{jt}) d\xi^*_{jt} - p_{jt} H_{jt-1} D Pr\left(\xi^*_{jt} > \exp(-\tilde{\pi})\right) \end{aligned}$$

If $\tilde{\xi}_{jt} \sim N(\mu, \sigma^2)$, then ξ^*_{jt} has a log-normal distribution and the incomplete integral from above (removing indices for clarity) is:

$$\begin{aligned} \int_{\exp(-\tilde{\pi})}^{\infty} \xi^* f(\xi^*) d\xi^* &= \int_{\exp(-\tilde{\pi})}^{\infty} \exp(\tilde{\xi}) f(\xi^*) d\xi^* \\ &= \int_{\exp(-\tilde{\pi})}^{\infty} \frac{\exp(\tilde{\xi})}{\xi^* \sigma \sqrt{2\pi}} \exp\left(-\frac{(\log(\xi^*) - \mu)^2}{2\sigma^2} \right) d\xi^* \end{aligned}$$

Let $\tilde{\xi} = \log(\xi^*) \Rightarrow d\tilde{\xi} = 1/\xi^* d\xi^*$ and $\xi^* = \exp(\tilde{\pi}) \Rightarrow \tilde{\xi} = \tilde{\pi}$. Then:

$$\begin{aligned}
\int_{\exp(-\tilde{\pi})}^{\infty} \xi^* f(\xi^*) d\xi^* &= \int_{-\tilde{\pi}}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp(\tilde{\xi}) \exp\left(-\frac{(\tilde{\xi}-\mu)^2}{2\sigma^2}\right) d\tilde{\xi} \\
&= \int_{-\tilde{\pi}}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2}(\tilde{\xi}^2 - 2\tilde{\xi}\mu + \mu^2 - 2\sigma^2\tilde{\xi})\right) d\tilde{\xi} \\
&= \int_{-\tilde{\pi}}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{1}{2\sigma^2}(\tilde{\xi}^2 - 2\tilde{\xi}(\mu+\sigma^2) + \mu^2 + 2\mu\sigma^2 + \sigma^4 - 2\mu\sigma^2 - \sigma^4)\right) d\tilde{\xi} \\
&= \int_{-\tilde{\pi}}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\tilde{\xi}-(\mu+\sigma^2))^2}{2\sigma^2} + \mu + \sigma^2/2\right) d\tilde{\xi} \\
&= \exp(\mu+\sigma^2/2) \int_{-\tilde{\pi}}^{\infty} \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(\tilde{\xi}-(\mu+\sigma^2))^2}{2\sigma^2}\right) d\tilde{\xi} \\
&= \exp(\mu+\sigma^2/2)\left(1 - \Phi\left(\frac{-\tilde{\pi}-(\mu+\sigma^2)}{\sigma}\right)\right) \\
&= \exp(\mu+\sigma^2/2)\Phi\left(\frac{\mu+\sigma^2+\tilde{\pi}}{\sigma}\right)
\end{aligned}$$

Giving:

$$\mathbb{E}(\pi^*_{jt}) = \tilde{\eta} p_{jt}(0)^{1+\eta} \exp(\mu+\sigma^2/2)\Phi\left(\frac{\mu+\sigma^2+\tilde{\pi}_{jt}}{\sigma}\right) - p_{jt} H_{jt-1} D\Phi\left(\frac{\mu+\tilde{\pi}_{jt}}{\sigma}\right) \tag{27}$$

G.3.4 Expected housing on each plot

The expected housing supply on each plot conditional on conditions in t (the relevant expectation for home prices) is:

$$\begin{aligned}
\mathbb{E}(H_{jt}|H_{jt-1}, p_{jt}) &= Pr(d_{jt}=0)H_{jt-1} + Pr(d_{jt}=1)\tilde{\eta} p_{jt}(0)^{\eta} \mathbb{E}(\xi^*_{jt}|d_{jt}=1) \\
&= Pr(\tilde{\xi}_{jt} < -\tilde{\pi}_{jt}) H_{jt-1} + \tilde{\eta} p_{jt}(0)^{\eta} \exp(\mu+\sigma^2/2)\Phi\left(\frac{\tilde{\pi}_{jt}+\mu+\sigma^2}{\sigma}\right) \\
&= H_{jt-1}\Phi\left(\frac{-\tilde{\pi}_{jt}-\mu}{\sigma}\right) + \mathbb{E}(H_{jt})\Phi\left(\frac{\tilde{\pi}_{jt}+\mu+\sigma^2}{\sigma}\right) \\
&= H_{jt-1}\left(1 - \Phi\left(\frac{\tilde{\pi}_{jt}+\mu}{\sigma}\right)\right) + \mathbb{E}(H_{jt})\Phi\left(\frac{\tilde{\pi}_{jt}+\mu+\sigma^2}{\sigma}\right) \\
&= H_{jt-1} - \Phi\left(\frac{\tilde{\pi}_{jt}+\mu}{\sigma}\right) H_{jt-1} + \mathbb{E}(H_{jt})\Phi\left(\frac{\tilde{\pi}_{jt}+\mu+\sigma^2}{\sigma}\right)
\end{aligned}$$

When $H_{jt-1} \leq \mathbb{E}(H_{jt})$ this is greater than H_{jt-1}: housing supply tends to increase, as development occurs when cost shocks are good. For values of $H_{jt-1} >> \mathbb{E}(H_{jt})$, it is possible that future housing supply can be lower than previous housing supply, conditional on prices.

G.4 Propositions about changes in consumer, producer surplus

Consider the counterfactuals defined as follows:

C1 All future expected sinking is zero: $NPVS_{jt} := \sum_{\tau=t+1}^{\infty} \rho^{\tau-t} \mathbb{E}_t(s_{j\tau}) = 0 \, \forall j, t$.

C2 Mexico City experiences no sinking ever. That is, $NPVS_{jt} = S_{jt} = 0 \, \forall j, t$.

C3 Information frictions are resolved: $\theta \to 1$.

We evaluate each derivative "in the long run", where the long-run is long enough that housing on a plot has a chance to be redeveloped and update to a new equilibrium. In the long-run, given a demand shock (as all of our counterfactuals represent), the new equilibrium represents a movement along the supply curve, and we know the elasticity of this supply curve is η:

$$\frac{\partial \log(H_{jt})}{\partial \log(p_{jt}(0))} = \eta \Rightarrow \frac{\partial \log(H_{jt})}{\partial X_{jt}} = \eta \frac{\partial \log(p_{jt}(0))}{\partial X_{jt}}$$

for some demand shifter X_{jt}.

G.4.1 Change in the probability of redevelopment

The probability that a plot is developed is:

$$Pr(\text{Redev}_{jt}) = \Phi\left(\frac{\mu + \tilde{\pi}_{jt}}{\sigma}\right), \qquad \tilde{\pi}_{jt} = \log(\tilde{\eta}) + \eta \log(p_{jt}(0)) - \gamma S_{jt} - \log(H_{jt}) - \log(D)$$

In counterfactual **C1**:

$$\begin{aligned}\frac{\partial Pr(\text{Redev}_{jt})}{\partial NPVS_{jt}} &\propto \eta \frac{\partial \log(p_{jt}(0))}{\partial NPVS_{jt}} - \frac{\partial \log(H_{jt})}{\partial NPVS_{jt}} \\ &= \eta \frac{\partial \log(p_{jt}(0))}{\partial NPVS_{jt}} - \eta \frac{\partial \log(p_{jt}(0))}{\partial NPVS_{jt}} = 0\end{aligned}$$

In counterfactual **C2**:

$$\frac{\partial Pr(\text{Redev}_{jt})}{\partial S_{jt}} \propto \frac{\partial Pr(\text{Redev}_{jt})}{\partial NPVS_{jt}} \frac{\partial NPVS_{jt}}{\partial S_{jt}} - \gamma = -\gamma > 0$$

In counterfactual **C3**:

$$\frac{\partial Pr(\text{Redev}_{jt})}{\partial \theta} \propto \eta \frac{\partial \log(p_{jt}(0))}{\partial \theta} - \frac{\partial \log(H_{jt})}{\partial \theta} = 0$$

G.4.2 Change in producer surplus

Expected profits are:

$$PS_{jt} = \tilde{\eta} p_{jt}(0)^{\eta+1} \exp(\mu + \sigma^2/2) \Phi\left(\frac{\mu + \sigma^2 + \tilde{\pi}_{jt}}{\sigma}\right) - \exp(\gamma S_{jt}) p_{jt}(0) H_{jt} D \Phi\left(\frac{\mu + \tilde{\pi}_{jt}}{\sigma}\right)$$

Let $\Xi(\tilde{\pi}_{jt}) = \exp(\mu + \sigma^2/2)\Phi\left(\frac{\mu+\sigma^2+\tilde{\pi}_{jt}}{\sigma}\right)$ in the following, and recall that:

$$\frac{\partial \Xi(\tilde{\pi}_{jt})}{\partial NPVS_{jt}} = \frac{\partial \Xi(\tilde{\pi}_{jt})}{\partial \theta} = 0$$

In counterfactual **C1**:

$$\begin{aligned}
\frac{\partial PS_{jt}}{\partial NPVS_{jt}} &= \Xi(\tilde{\pi}_{jt})\tilde{\eta}(\eta+1)p_{jt}(0)^{\eta}\frac{\partial p_{jt}(0)}{\partial NPVS_{jt}} \\
&\quad - \exp(\gamma S_{jt})D\Phi\left(\frac{\mu+\tilde{\pi}_{jt}}{\sigma}\right)\left(\frac{\partial p_{jt}(0)}{\partial NPVS_{jt}}H_{jt} + \frac{\partial H_{jt}}{\partial NPVS_{jt}}p_{jt}(0)\right) \\
&= \Xi(\tilde{\pi}_{jt})\tilde{\eta}(\eta+1)p_{jt}(0)^{\eta}\frac{\partial p_{jt}(0)}{\partial NPVS_{jt}} \\
&\quad - \exp(\gamma S_{jt})D\Phi\left(\frac{\mu+\tilde{\pi}_{jt}}{\sigma}\right)\left(\frac{\partial p_{jt}(0)}{\partial NPVS_{jt}}H_{jt} + \eta\frac{\partial p_{jt}(0)}{\partial NPVS_{jt}}H_{jt}\right) \\
&= \frac{\eta+1}{p_{jt}(0)}\frac{\partial p_{jt}(0)}{\partial NPVS_{jt}}\left(\Xi(\tilde{\pi}_{jt})\tilde{\eta}p_{jt}(0)^{\eta+1} - p_{jt}H_{jt}D\Phi\left(\frac{\mu+\tilde{\pi}_{jt}}{\sigma}\right)\right) \\
&= (\eta+1)\frac{\partial \log(p_{jt}(0))}{\partial NPVS_{jt}}PS_{jt} = (\eta+1)\gamma\theta PS_{jt} \leq 0
\end{aligned}$$

In counterfactual **C2**:
The change in producer surplus is equal to that of **C1**, plus the change due to the change in S_{jt}:

$$\begin{aligned}
\frac{\partial PS_{jt}}{\partial S_{jt}} = &\underbrace{\frac{\partial PS_{jt}}{\partial NPVS_{jt}}\frac{\partial NPVS_{jt}}{\partial S_{jt}}}_{\leq 0}\underbrace{-\gamma p_{jt}H_{jt}D\Phi\left(\frac{\mu+\tilde{\pi}_{jt}}{\sigma}\right)}_{>0} \\
&\underbrace{-\frac{\gamma}{\sigma}\left(\tilde{\eta}p_{jt}(0)^{\eta+1}\exp(\mu+\sigma^2/2)\phi\left(\frac{\mu+\sigma^2+\tilde{\pi}}{\sigma}\right) - p_{jt}H_{jt}D\phi\left(\frac{\mu+\tilde{\pi}}{\sigma}\right)\right)}_{\lesseqgtr 0}
\end{aligned}$$

The sign of this derivative is unknown. However, some patterns emerge:

- As very high or low $\tilde{\pi}_{jt}$, the last term goes to zero, and there is little impact through the changed probability of re-development.
- For values of $\tilde{\pi}_{jt}$ near the middle of the distribution, where cost shocks are more marginal to the decision to develop, things are more complicated for this term.
- Sending this term to zero, we see that benefits are not clear to producers. Subsidence benefits them by lowering the cost of development (term 2), but it also lowers revenues (term 3). As $\theta \to 0$, the costs channel wins.

In **C3**:

$$\begin{aligned}\frac{\partial PS_{jt}}{\partial \theta} =& \frac{\partial \tilde{\pi}_{jt}}{\partial \theta}\left(\tilde{\eta} p_{jt}(0)^{\eta+1}\exp(\mu+\sigma^2/2)\phi\left(\frac{\mu+\sigma^2+\tilde{\pi}}{\sigma}\right) - p_{jt}H_{jt}D\phi\left(\frac{\mu+\tilde{\pi}}{\sigma}\right)\right) \\ &+ \frac{\partial \log(p_{jt}(0))}{\partial \theta}(\eta+1)\,PS_{jt} \\ =& \frac{\partial \log(p_{jt}(0))}{\partial \theta}(\eta+1)\,PS_{jt} \leq 0\end{aligned}$$

Information frictions unambiguously increase producer surplus by increasing revenue.

G.5 Proof of validity of control function

The full structural supply model is:

$$\begin{aligned} d_{jt} &= 1\left(\tilde{V}^S_{jt} \geq 0\right) \\ \tilde{V}^S_{jt} &= \log(\tilde{\eta}) + \eta\log(p_{jt}) - \gamma(1+\eta)S_{jt} - \log(H_{jt-1}) - \log(D) + \log(\xi_{jt}) \\ \log(p_{jt}) &= \gamma S_{jt} + \Pi\text{Rel. altitude change}_{jt} + \tilde{u}_{jt}\end{aligned}$$

The proof follows Blundell and Matzkin (2014), in showing that the residual $\tilde{u}_{jt}$ satisfies control function separability, which follows naturally from the linearity of the first stage:

$$r^2(\tilde{V}^S_{jt}, \log(p_{jt}), Z_{jt}) = \tilde{u}_{jt} = \underbrace{\log(p_{jt}) - \gamma S_{jt} - \Pi\text{Rel. altitude change}_{jt}}_{q(\log(p_{jt}), Z_{jt})}$$

Note that $v(q(\log(p_{jt}), \text{Rel. altitude change}_{jt}), r^1(\log(p_{jt}))) = q(\log(p_{jt}), \text{Rel. altitude change}_{jt})$, so q is strictly increasing in $\overline{p}_{jt}$ and $v()$ is strictly increasing in $q()$. Thus, the structural inverse of the system satisfied control function separability, and a control function can be derived from the first stage to estimate the second stage.

G.6 Computing counterfactuals

Given our sparse data on prices, we seek to estimate counterfactual prices and housing densities using a dataset of plots as a single cross-section. To do this, we consider a steady state in which the joint distribution of the underlying random variables that generate the equilibrium, $(\xi_{jt}, \varepsilon^i_{jt}, e_{jt}, s_{jt})$ is time-invarying and approximately non-stochastic given the large number of plots we consider (985,862). The intution of this steady state is that while the draws of cost and preference shocks are stochastic, given enough plots the distribution of $(p_{jt}(0), H_{jt})$ is constant in each period and can be treated as fixed. In the steady state, we drop the t subscript to emphasize the independence of time.

The equilibrium is most easily expressed in terms of log prices and log relative housing relative to the outside option. Let $\tilde{\lambda}_j = \log(\lambda_j) - \log(\lambda_0)$; in equilibrium, $\tilde{\lambda}_j = \log(H_j) -$

$\log(H_0)$. The equilibrium is characterized by:

$$\tilde{\lambda}_j^D = -\log(R_j(0)) + Z_j\delta + e_j \tag{28}$$

$$\tilde{\lambda}_j^S = \eta \log\left(\frac{\eta}{1+\eta}\right) + \eta \log(p_j(0)) + \tilde{\xi}_j - \log(H_{0t}^S) \tag{29}$$

$$\tilde{\lambda}_j^D \approx -\log(p_{jt}(0)) + Z_j\delta + e_j + \gamma\theta \sum_{\tau=t}^{\infty} \rho^{\tau-t} \mathbb{E}_t(s_{j\tau}) \tag{30}$$

The last equation is derived directly from Equation (16) evaluated at $S_{jt} = 0$.

Given data on the plot level on $(\tilde{\lambda}_j, \log(p_j(0)), Z_j, \sum_{\tau=t}^{\infty} \rho^{\tau-t}\mathbb{E}_t(s_{j\tau}))$ and estimates or calibrated values of the parameters $(\gamma, \theta, \delta, \eta, \rho)$, we first solve the previous equations exactly to back out $(e_j, \tilde{\xi}_j)$. From the inferred values of $\tilde{\xi}_j$, we calibrate μ and σ as the mean and standard deviation of these shocks.

In order to calculate welfare under different counterfactuals, for example setting $\mathbb{E}_t(s_{j\tau}) = 0$ or $\theta = 1$, we alter these parameters and, holding $(Z_j, e_j, \tilde{\xi}_j)$ fixed, solve the system for alternative vectors of prices and relative housing supply that clear the housing market on each plot.

G.7 Proof that the new error function is independent of prices

In order for the control function approach to provide unbiased estimation of the coefficient on prices, prices and the (modified) error term must be independent. We prove that here:

$$\begin{aligned}
\mathbb{E}\left(\tilde{\xi}_{jt}^* | \log(p_{jt})\right) &= \mathbb{E}\left(\mathbb{E}\left(\tilde{\xi}_{jt}^* | \log(p_{jt}), Z_{jt}\right) | \log(p_{jt})\right) \\
&= \mathbb{E}\left(\mathbb{E}\left(\tilde{\xi}_{jt} - \mathbb{E}(\tilde{\xi}_{jt}|\tilde{u}_{jt}) | \log(p_{jt}), Z_{jt}\right) | \log(p_{jt})\right) \\
&= \mathbb{E}\left(\mathbb{E}\left(\tilde{\xi}_{jt} | \log(p_{jt}), Z_{jt}\right) - \mathbb{E}\left(\mathbb{E}(\tilde{\xi}_{jt}|\tilde{u}_{jt}) | \log(p_{jt}), Z_{jt}\right) | \log(p_{jt})\right) \\
&= \mathbb{E}\left(\mathbb{E}\left(\tilde{\xi}_{jt} | \log(p_{jt}), Z_{jt}\right) - \mathbb{E}\left(\mathbb{E}(\tilde{\xi}_{jt}|\tilde{u}_{jt}) | \tilde{u}_{jt}, Z_{jt}\right) | \log(p_{jt})\right) \\
&= \mathbb{E}\left(\mathbb{E}\left(\tilde{\xi}_{jt} | \tilde{u}_{jt}, Z_{jt}\right) - \mathbb{E}\left(\tilde{\xi}_{jt} | \tilde{u}_{jt}\right) | \log(p_{jt})\right) \\
&= \mathbb{E}\left(\mathbb{E}\left(\tilde{\xi}_{jt} | \tilde{u}_{jt}\right) - \mathbb{E}\left(\tilde{\xi}_{jt} | \tilde{u}_{jt}\right) | \log(p_{jt})\right) = \mathbb{E}\left(0 | \log(p_{jt})\right) = 0
\end{aligned}$$

H Cost-benefit Appendix

Our empirical estimates provide a measure of the total cost of subsidence, which expressed in terms of dollars per meter of average elevation loss is:

$$\frac{\text{Total cost}}{\text{meter of elevation loss}} = \frac{\text{Total cost}}{\overline{NPVS}} = \frac{1}{0.931}\text{Total cost} \tag{31}$$

In order to price the externality-generating action, groundwater pumping, we must convert this to dollars per unit of water pumped.

Subsidence is caused by low groundwater tables. We estimate the response of subsidence rates to changes in aquifer storage by taking the average change in subsidence rates over the 13-year period in our panel, and estimating the dosage-response of subsidence rates to changes in aquifer storage:

$$\begin{aligned} \Delta s &= \frac{\Delta s}{\Delta V} \times \Delta V \\ \Delta s &= \underbrace{\frac{0.0361}{-111 \times 13}}_{\substack{\text{Annual loss} \times T = \\ \text{Total loss} \\ \text{over 13 years}}} \Delta V \end{aligned} \tag{32}$$

Here Δs is the change in the annual subsidence rate (change in $m/$ year), and ΔV is the net loss in aquifer storage (millions of cubic meters, or hectometers hm^3). We take the annual net loss in aquifer storage from CONAGUA (2024); details from this report on aquifer demand over this period are described in Table A18.[62]

Table A18: Aquifer demand

	Annual volume (hm^3)
Natural recharge	151.4
Recharge from drinking water leaks	361.4
Total recharge	512.8
Pumping for drinking water	448.5
Other pumping	175.3
Total extraction	623.8
Net volume change per year	-111

Notes: Source: CONAGUA (2024).

Solving this formula for ΔV and assuming no change in recharge implies that aquifer storage would need to increase by a total of 2,495 hm^3 in order to fully abate average

[62] This expression extracts away from dynamics; how long it would take for a higher water table to translate to slower sinking rates is difficult to answer but crucial question for policy. While it may take decades for higher aquifer levels to translate to fully abated sinking, interventions that locally rise water tables such as wastewater injection sites have seen rapid abatement of subsidence.

subsidence rates of 6 cm/year, the equivalent of reducing total pumping by a third for the next 26 years, according to:

$$\Delta V = (\text{Recharge} - (1-r)\text{Pumping}) \times Y$$
$$2495 = (512.8 - (1-r)623.8) \times Y$$

where r is the percentage reduction in pumping, and Y is the number of years this reduction is sustained for.

H.1 Cost-benefit of other recharge policies

Here we discuss the costs and benefits of these policies, a challenging task given the lack of detailed information on the cost structure of some of these interventions. We discuss how we construct the costs and benefits of each of the explored policies, with emphasis on the limitations of these estimates.

The net benefits of a policy that facilitates aquifer recharge is the gains from offsetting subsidence minus the sum of its treatment, procurement and construction costs net what it the current system costs:

$$\Delta C = \Delta W - \Delta\left(\text{Treatment} + \text{Operating} + \text{Construction}\right)$$

Treatment costs are the costs of treating the water, which is a function of the source water quality (aquifer, wastewater, surface water) and its intended use (drinking, injection). *Operating costs* are the costs of bringing the water from its location (aquifer, surface source) into the Mexico City water supply or running the program (in the case of fixing leaks). *Construction* costs are the costs of building the proposed investment. ΔW is the change in welfare, equal to the costs of subsidence offset by the volume of recharge permitted.

H.1.1 Wastewater injection

Table A19: Values and sources for wastewater injection costs

Cost type	Explanation and source
Treatment costs	Cost of treating water to be clean enough to enter the aquifer without activating heavy metals or causing well clogs.
Operating costs	Annual costs of operation, treatment
Construction costs	Building treatment and injection infrastructure.

We take costs of investing in wastewater injection and the potential capacity of an injection site from López-Morales and Mesa-Jurado (2017), who characterize the costs of treating wastewater and injecting in the Chichinautzin Sierra.

H.1.2 Leaks control

Table A20: Leaks repair costs

Cost type	Explanation
Treatment costs	There is no change in treatment costs when fixing leaks because the water supply mix is unaffected.
Operating costs	Annual investment in leaks repair. There is no change in the cost of procuring water because the water supply mix is unaffected.
Construction costs	Some minor investment in identifying leaks and setting up the program

We take reports of the costs and efficiency of efforts to repair leaks in the system from López-Morales and Mesa-Jurado (2017). Repairing leaks may have contradictory impacts on the aquifer; CONAGUA (2024) estimate that most recharge on the aquifer comes from lost water in the Mexico City system, and considering that some of that water is imported from outside the watershed would imply that some of this recharge is additional to what is pumped. We take an optimistic scenario and assume that all leaks that are addressed are done on parts of the system *prior* to entering the city, which would make all water saved additional. However, given that our cost estimates are calculated using efforts conducted *within* the city, it is unclear whether we over- or under-estimate these relative to the true costs of repairing leaks outside the city. While infrastructure outside the city is exposed and therefore less costly to work on than buried lines, these lines are also larger and more economically costly to shut down.

H.1.3 New water sources

Table A21: New water sources costs

Cost type	Explanation
Treatment costs	We assume no change in treatment costs. These are potentially even net negative, as it may be cheaper to treat surface water than groundwater which is heavily contaminated with Arsenic.
Operating costs	Conveyance and storage costs.
Construction costs	Infrastructure investment to transport surface water to the city.

We take reports of the costs and efficiency of efforts to generate new water sources from López-Morales and Mesa-Jurado (2017). They consider five alternative sources of water:

Temascaltepec, a surface water system that connects with the Cutzamala System; Tecolutla (Veracruz); Amacuzac (Morelos); Taximay, a dam in the Estado de México; and Oriental-Libres, a groundwater basin in the state of Puebla. Our final estimate for alternative water sources averages across these. We use their operating costs and initial investment estimates, but assume that construction takes 5 years and use our own discount rate to aggregate costs.

H.1.4 Wastewater recycling

Table A22: Wastewater recycling costs

Cost type	Explanation
Treatment costs	This is difficult to characterize; while grey water is very costly to treat for reintroduction into the water supply, the high Arsenic content of the groundwater currently being pumped is also costly to treat.
Operating costs	Conveyance and storage costs.
Construction costs	Infrastructure investment to transport surface water to the city.

Our primary source for information the costs and benefits of investing in wastewater recycling plants comes from the 2007-2012 Plan for Sustainable Water Management in Mexico City (Gobierno del Distrito Federal, 2007). This document describes an investment in building 6 macro-sites for secondary wastewater treatment. We assume a useful life of 50 years once built following engineering guidance. In the absence of information on operating and treatment costs *relative* to the current cost of treating water from the aquifer, which this would offset, we use a value of $0.34 USD per cubic meter, the same cost as treatment for wastewater injection. If we take the cost of treatment for wastewater injection as the benchmark cost of treating water that comes out of the aquifer, this would imply that treating wastewater is twice as expensive as treating groundwater. To the extent that treating groundwater is more expensive than treating water for injection, this margin is smaller. Also note that construction costs may have evolved substantially between 2007 and now, and better estimates of operating costs are necessary for comprehensive analysis.