

Why Has Construction Productivity Stagnated?

The Role of Land-Use Regulation*

Leonardo D’Amico Edward Glaeser Joseph Gyourko
William Kerr Giacomo A. M. Ponzetto

December 30, 2023

Why does it cost so much to build a home? We formalize and evaluate the hypothesis that land-use regulation reduces the average size of home builders, which limits their ability to reap returns from scale and their incentives to invest in technology. Our model distinguishes between regulation of entry, which acts as a fixed cost and increases equilibrium firm size, and project-level regulation, which reduces project and firm size. If larger firms have stronger incentives to invest in technology, then such investment partially offsets the harm that regulation of entry does to consumers, but reduced investment exacerbates the negative impacts of project-level regulation. We document that the US has higher production costs than comparably wealthy countries, and that these costs are higher in more regulated American cities. Homes built per construction worker remained stagnant between 1900 and 1940, boomed after World War II, and then plummeted after 1970 just as land-use regulations soared. Residential construction firms are small, relative to other industries like manufacturing, and smaller firms are less productive. More regulated metropolitan areas have smaller and less productive firms. Under the assumption that one half of the link between size and productivity is causal, America’s residential construction firms would be 91% percent more productive if their size distribution matched that of manufacturing.

*D’Amico: Harvard University, damico@g.harvard.edu; Glaeser: Harvard University and NBER, eglaeser@harvard.edu; Gyourko: The Wharton School, University of Pennsylvania and NBER, gyourko@wharton.upenn.edu; Kerr: Harvard Business School and NBER, wkerr@hbs.edu; Ponzetto: CREI, UPF, BSE, IPEG and CEPR, gponzetto@crei.cat. We thank Leonardo Bianchi, Davor Djekic, Kaiser Ke, Aman Mahajan, Braydon Neiszner and especially Pietro Dedola for outstanding research assistance. D’Amico and Glaeser thank the generous support of the Harvard Joint Center for Housing Studies. Gyourko thanks the Research Sponsor Program of the Zell/Lurie Real Estate Center at Wharton. Ponzetto thanks the Spanish Ministry of Science and Innovation (grant PID2019-111691RA-I00) and the Spanish Ministry of Economy and Competitiveness through its Severo Ochoa Program for Centers of Excellence in R&D (CEX2019-000915-S). Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004). The appendix to this paper is available [here](#) and on the authors’ websites.

1 Introduction

Most analyses of America’s growing housing affordability problem focus on land values far more than construction costs, because the key scarcity seems to be a shortage of legally developable land (for reviews in economics see Gyourko and Molloy (2015) and Glaeser and Gyourko (2018); Ellickson (2022) provides a recent legal perspective). But US construction costs are also high, and macroeconomists studying sluggish American growth have found that construction stands out as a leading long-term driver of the decline in trend GDP growth over the postwar period, contributing for over a quarter of it: 0.75 percentage points out of 2.8 (Foerster et al. 2022). Goolsbee and Syverson (2023) document that value added per construction worker fell by 40 percent between 1970 and 2020 and convincingly show that this is not an artifact of measurement or estimation error. Manufacturing productivity quadrupled over the same time period. Moreover, there is broader evidence supporting the view that America has a construction productivity problem. Brooks and Liscow (2023) document that the real cost of building a mile of highway tripled between the 1960s and the 1980s. The Transit Costs Database, housed at NYU’s Marron Institute, finds that “transit-infrastructure projects in New York cost 20 times more on a per kilometer basis than in Seoul.”

Why are Americans so good at building tradable products and so bad at producing structures that are fixed in space? In this paper, we formalize and investigate the hypothesis that regulation drives the difference between manufacturing and construction, partially by ensuring that residential construction is built by smaller and less productive firms. Construction sites are both fixed in space and highly visible, while assembly lines can relocate and are typically behind closed doors.¹ Building is regulated *ex ante*, with local laws and state and federal environmental regulations possibly restricting the ability to begin development at all. The approval process can require years of community outreach and catering to the wishes of incumbent residents, and it is often easier for smaller projects.² By contrast, federal regulations of manufacturing

¹Of course, regulation effectively can change the location of new home building from Massachusetts and California to Texas and Georgia.

²In this paper, we will mostly focus on the level of regulatory burden and its interaction with construction productivity. However, regulation can have additional perverse effects as local political economy constraints also affect which types of buildings get built (Hamilton 1975; Fernandez and Rogerson 1996; Calabrese, Epple, and Romano 2007; Krimmel 2021). More generally, regulation can generate policies that

are typically enforced *ex post*, and size can make compliance easier.

In Section 2.1, we present three sets of facts about the construction sector. The first set adds to the time series produced by Goolsbee and Syverson (2023). Using an industry-produced cost series, we document the increasing physical cost of building homes (which excludes the cost of land assembly and purchase) between 1985 and 2022. This series complements Goolsbee and Syverson’s figure showing that the construction price deflator has grown substantially faster than the overall GDP deflator since the 1980s. We then focus on the number of homes produced per worker, and extend Goolsbee and Syverson’s series on this particularly direct measure of productivity back to 1900. From 1900 to the 1930s, home production exhibited no particular trend in growth. From 1935 to 1970, homes produced per construction worker grew at a pace that at times exceeded the growth in the number of cars produced per automobile industry worker or the growth of total manufacturing output per industrial worker. Booming construction productivity during this period suggests that there is nothing intrinsic about housing that limits technological innovation. Since 1970, these three series have diverged sharply. Car and manufacturing output per worker continued to soar but houses built per worker fell dramatically. The post-war pattern for these series corresponds with the post-war patterns in the total factor productivity series shown by Goolsbee and Syverson (2023). This section also highlights that regulation, as proxied by two different time series constructed by Ganong and Shoag (2017) and Jackson (2016), likewise shot up after 1970.³

Our second set of facts compares physical construction costs in the US with costs elsewhere in the world. We use data provided by the infrastructure consultants, Turner & Townsend (2022). This data is particularly oriented towards large buildings. Within the US, building in lightly regulated Houston is substantially cheaper than building in highly regulated San Francisco or New York. Yet, even Houston is not much cheaper than Paris. When we regress construction costs per square meter on costs of inputs at the city level, the US cities are generally significant outliers. US coastal cities are expensive in terms of physical construction costs relative to East Asia and Western Europe.

Our third set of facts concerns the size of residential firms and residential projects. Housing is built by firms that are much smaller than in almost any manufacturing in-

only cater to the incumbents and disregard externalities that could materially affect aggregate outcomes (Glaeser and Shleifer 2005; Glaeser and Ponzetto 2018; D’Amico 2022).

³We are grateful to Jacob Krimmel for sharing the data with us.

dustry, and the typical housing project is also small. Almost 40% of employment in new single-family housing construction is in firms with less than 5 employees—versus 2% for manufacturing. The acreage of the typical land parcel bought in recent years for single-family development is also quite small. Gyourko and Krimmel (2021) report a median size of seven acres in their sample of 3,600 parcels purchased between 2013 and 2018 across 24 Core Based Statistical Areas (CBSAs). That small size reflects plentiful infill development, but really large suburban tracts are also rare now. More than 94 percent of their parcels contain less than 100 acres of land, and there are virtually no parcels with more than 1,000 acres. The great builders of the post-WWII era worked with far larger tracts of land.

Section 3 presents a model that links project-level regulation, firm size and productivity. Firms are led by entrepreneurs who have a limited ability to monitor different projects. These entrepreneurs also choose how much to invest in new technology. The model emphasizes how project-level regulation differs from regulation of entry because of its impact on firm size and investment.

Project-level regulation causes the average project size to shrink which then causes firm size to also shrink because entrepreneurs cannot monitor one hundred projects with two houses each as easily as they can monitor two projects with one hundred houses each. In contrast, regulation of entry limits the number of firms in the market and increases firm size. If technology can be used in all projects without decreasing returns to scope, then bigger firms invest more in technology. The adverse effects of regulation on entry on consumers will be partially offset by the fact that the bigger firms that remain in the industry have lower costs. The adverse effects of project-level regulation on consumers will only be exacerbated by the small scale and limited technological investment of the firms that remain in the industry. The key cross-sectional predictions of the model are that smaller firms are less productive and that in jurisdictions where regulation is more intense, average firm size and firm productivity are both lower.

Section 4 then tests these predictions using data from the Economic Census and microdata from the Longitudinal Business Database (LBD). This Census data set contains establishment and firm level data for all industries, including construction. It also contains direct measures of output, such as the number of homes produced in the construction sector. We test three key implications of the model.

First, we document the strong connection between firm size and productivity in the construction sector. In housing construction, firms with 500+ employees produce six times as many units per employee than firms with less than 20 employees. If only a tenth of the link between size and productivity reflected the causal effect of size, then we estimate that construction would be 18% more productive if the firm size in construction was like that in manufacturing. If half of the link between firm size and productivity is causal, then overall productivity in construction would increase by 90 percent. Thus, if some part of the small firm size in the construction sector reflects regulation, then regulation can help explain why productivity growth in construction is so low.

Our second set of analyses shows that areas with stricter land-use regulation (as measured by the Wharton Residential Land Use Regulation Index (WRLURI)) indeed have smaller establishments (as measured by total revenues per establishment or by the fraction of construction employment in firms with more than 100 employees). This relationship holds across all types of construction firms, but it is especially strong in the subsector involving construction of buildings, where residential regulation is likely to bite the most because it excludes heavy construction and civil engineering. In this subsector, a one standard deviation increase in the WRLURI index (which is approximately the difference between Atlanta and San Francisco) is associated with a 12% reduction in total receipts per establishment and a 4.3 percentage point reduction in the fraction of employment in large firms, which amounts to a more than one-third reduction in the share of employment in these firms.

Third, we show that there is just much less construction activity in areas with stricter land use regulation. The index is associated with lower levels of non-residential as well as residential activity. The impact on non-residential activity is unsurprising, both because land-use regulations also impact non-residential activity and because if land-use regulation deters residential growth, there is likely to be less demand for warehouses and offices due to the smaller number of people in the area.

Section 5 concludes by discussing the path forward for research into why construction is so unproductive.

2 Three Sets of Facts on the Construction Sector

In this Section, we provide evidence that building in the US is very expensive and dominated by many small firms. The construction sector as a whole experienced far slower productivity growth than other industries, and this slump began in the 1970s, which was also when land-use regulation increased. Section 2.1 provides time-series evidence on the evolution of American construction prices, costs, quantities, and regulations. Section 2.2 compares America's construction costs with costs in other countries. Section 2.3 documents the small size of construction firms and construction projects.

2.1 Construction over the Long Haul

2.1.1 Prices


Figure 1 reports the evolution of prices of houses and cars from 1950 to 2022, where we have indexed all series to 100 in 1960. In light blue (top line), we use data from decadal censuses and the American Community Survey (ACS), to report a series of new house values controlling for quality and the overall Consumer Price Index (CPI).⁴ We correct for quality by estimating a household level regression where the logarithm of CPI-adjusted self-reported home values is regressed on physical attributes. We then report the changing value of houses holding physical characteristics constant.⁵ The dark blue (middle) line reports the CPI for new homes from Shiller (2015).⁶ Finally, the green (bottom) line is the CPI for new vehicles collected by the Bureau of Labor Statistics (BLS).

⁴Self-reported value of houses is computed from decadal censuses between 1960 and 2000, and from the ACS 1-yr for 2000 onwards.

⁵The physical characteristics are: the number of rooms in the housing unit, whether the unit has access to a kitchen or to plumbing facilities, and the number of units in the structure. Appendix A.2 reports different specifications for this quality adjustment, the associated coefficients, as well as the path of the price as predicted only by the physical characteristics across the different specifications. Results are similar whether or not we put the physical characteristics on their own, or whether we add state-by-year fixed effects or the price of old houses. To be conservative, we chose as baseline the specification that predicts the larger increase in quality over time, which is the one without controls.

⁶As described in Shiller (2015), the 1953 to 1975 data is the home purchase component of the CPI. The Bureau of Labor Statistics collected data on homes constant in age and square footage. The data from 1975 is the Case-Shiller index, which also intends to keep quality constant.

Figure 1: Prices of Cars and Houses, net of Overall CPI




Note. The figure plots price indexes for homes and cars. Home price indexes are from Census self-valuations (top, light blue) and Shiller (2015) (middle, dark blue). The CPI for vehicles (bottom, dark green) is from the Bureau of Labor Statistics.

Both housing price series were growing at the same pace as the overall CPI and vehicles through the 1960s, before diverging substantially in the 1970s. While new homes now cost twice as much as they did in 1960 in real terms, cars are 60 percent cheaper. This finding mirrors that of Goolsbee and Syverson (2023), who find that the construction output deflator rose much faster than the overall GDP deflator from the 1970s onward.

2.1.2 Costs

Across space, variation in the costs of producing homes only explains a small fraction of the variation in the costs of buying a home (Gyourko and Saiz 2006). Similarly, much of the rise in housing prices reflects the rising cost of land and the increased difficulty of getting a permit (Glaeser and Gyourko 2018). To focus on construction productivity, we take data from R.S. Means, a private provider of construction cost data, on the real cost

Figure 2: Depiction of a Constant “Economy Quality” 1,800 Square Foot House



Note. An “Economy Quality” home, used as the basis for the real cost series in Figure 3. From the 2021 R.S. Means company data book.


of a constant quality 1800 square-foot “economy quality” house.⁷ This cost is meant to include only the physical costs of construction, not land acquisition or other costs. R.S. Means defines “economy quality” as a simple, relatively low-cost, one-story home, as depicted in Figure 2, which comes from their 2021 data book. The quality of this home does change, but relatively infrequently and in relatively easy-to-measure ways.⁸

Figure 3 plots the real cost per square foot of supplying that kind of home from 1985 to 2021. Even after accounting for a slight increase in structure quality over time, we find that from 1985 to 2021 the cost of building this modest quality home increased by 18 percent, or \$17.52 per square foot. This fact supports the view that productivity in the homebuilding sector either stagnated or declined, which contrasts sharply with the

⁷The R.S. Means Company, now owned by Gordian, publishes annual information on construction costs of different types of structures, but not all books are available. We have information from 24 of those 36 years (1985, 1986, 1989, 1991, 1992, 1996, 2002, 2003, and 2006-2021). Those years are marked with dots in Figure 3. The reference for 2021 is *Square Foot Costs with RSMeans Data. 2021, 42nd annual edition (and other years).* (2021).

⁸Quality is transparently reflected in a set of “traits:” whether there is a vapor barrier in the foundation, the number of coats of paint, the type of shingles on the roof, etc. The baseline trait set is constant between 1997 and 2021, and is reported in Appendix Figure A1. Before 1997, however, economy-quality homes had fewer traits. To adjust for this, we reconstruct the cost series before 1997 by keeping the set of traits constant to their 1997-2021 level; which we can do using cost data for each trait, provided by another R.S. Means publication, the *Building Construction Cost Data, 1998, 56th annual edition.* (1998). These costs are inclusive of the materials and the labor cost to apply the trait (e.g. apply an extra coat of paint). Appendix Section A.1 provides more details on the procedure and the traits.

Figure 3: Real Cost of a Constant “Economy Quality” 1,800 Square Foot House



Note. The Figure plots the real cost per square foot of supplying a constant “Economy Quality” home, as depicted in Figure 2. The cost comes from our calculations based on R.S. Means company data and only includes the physical costs of construction, not land acquisition or other costs.

evidence in Figure 1 that firms have gotten better at producing cars and making them more affordable.

2.1.3 Quantities


We now turn to output per employee. To continue the parallel with manufacturing and cars, we consider three indexes: (i) newly-started housing units per construction sector employee, (ii) domestically produced cars⁹ per employee in motor vehicle production, and (iii) manufacturing output per manufacturing employee.

Figure 4 illustrates the evolution of new housing units per construction employee from 1900 to 2023. We create this long series by splicing together data from the Construction Census and other sources. Output per worker is obviously highly cyclical, but the figure shows a flat trend between 1900 and the mid-30s, a steep increase in output per employee between the mid-30s and early 70s, and then a sharp decline.

A significant concern with this series is that it is a series of housing units per

⁹Cars include automobiles and light trucks.

Figure 4: Housing Units Started per Employee in the Construction Sector



Note. The figure plots housing units started per employee in the construction sector, from 1900 to 2023. We take housing units between 1900 and 1959 from the Macrohistory Database, specifically from the sources denominated *US Number of New Private Nonfarm Housing Units Started, One-, Two-, and Three-or-more*. From 1959 onward, the data comes from the Census’s New Residential Construction program, specifically from the *New Privately Owned Housing Units Started* series. Employment data in the construction sector between 1900 and 1945 was obtained from the *Historical Statistics of the United States, 1789 - 1945*, series D62-76. For the 1929-1945 time period, we also consulted a Bureau of Labor Statistics (BLS) historical report, which corroborates our main series. From 1939 onward, employment in the construction industry was taken from the BLS’s Current Employment Statistics (CES).

construction sector employee, not a series of housing units per residential construction sector employee. If residential construction employment changed dramatically as a share of total construction employment, this figure would misrepresent changes in residential construction productivity. Unfortunately, consistent data on employment in residential construction is not available, so we use an imperfect alternative.

We create a proxy for residential construction employment by using data on the number of employees of general contractors engaged in the construction of buildings, which is available from 1935 to today, though irregularly.¹⁰ This approach removes

¹⁰Henceforth, we refer to general contractors engaged in the construction of buildings simply as general contractors. We have data for 1935, 1939, from 1945 to 1952, from 1960 to 1967, 1972, 1977, 1982, 1987 and from 1990 onward. For missing years between 1935 and 1990, we impute the number of general contractors by simply assuming a linear trend for the share of general contractors as a share of total employment and multiplying the predicted share with total employment in construction, which we have for all years. Appendix A.5.3 provides the details.

employees engaged in heavy construction such as highways and bridges, which are generally not built by general contractors or specialty contractors (e.g., the latter are typified by electricians and plumbers, who are mostly employed in maintaining the existing housing stock). Appendix Figure A7 shows that dividing new housing units by the number of employees of general contractors does not change the qualitative patterns shown in Figure 4. Dividing home production by the number of employees of general contractors is also imperfect, because general contractors include those putting up non-residential buildings and exclude specialty contractors that might be engaged in the construction of new housing. In Appendix A.5.2, we construct two other series that address these issues with more recent data and show that the patterns remain identical for a range of possible proxies for employment in residential construction.¹¹

Figure 5 compares our general-contractor-based series of residential construction productivity with general manufacturing and automobile production productivity.¹² Manufacturing productivity is measured with an index of real manufacturing production divided by total manufacturing employment.¹³ Automobile productivity is measured with data from a number of different data sources.¹⁴

This plot uses a base-10 logarithmic scale, and the series are indexed to a value of 100 in 1967 (the value in 1967 corresponds to $\log_{10}(100)$). While US construction


¹¹This is because employment shares within construction have not changed considerably during this period. We use all general contractors as our preferred series because it allows us to go further back in time and because it is the most conservative since, if anything, it suggests a slightly larger productivity gain compared to all other series.

¹²For this exercise, we construct a continuous housing start series that pastes together our data sources, averaging them for the periods where they are overlapping. Appendix subsection A.5.1 spells out the details.

¹³In particular, manufacturing goods production is measured by the *Industrial Production and Capacity Utilization - G.17* (IPCC) index for manufacturing output, released by the Board of Governors of the Federal Reserve System, and covers all years from 1919 onward. Data on manufacturing employment come from two data sources: the *Historical Statistics of the United States, 1789 - 1945* D62-76 table gives information from 1900 to 1945, and the CES reports employment data from 1939 onward.

¹⁴For the numerator, between 1929 and 1975 we use data on all automobiles produced reported in the *US Automobile Production Figures* on Wikipedia. This data covers only automobiles produced in the US from 1900 to 2000 and is collected from the *Consumer Guide* magazine (2000; 2001; 2001; 2004). From 1975 onward we use the IPCC index for automobiles and light truck production, which is available from 1972 to 2023. The advantage of using this index is that it accounts also for light-truck production, which was negligible prior to 1975 but now has the lion's share of the market. For the denominator, we took the data on employment in motor vehicles and part production from the Bureau of Economic Analysis (BEA) *Full-Time and Part-Time Employees by Industry* tables. When the two series on car production overlap (as occurs in 1972–1975), we paste them together, following the same procedure detailed in A.5.1 for homes per employee.

Figure 5: Housing Units per Employee Against Manufacturing Output per Employee and Cars per Employee




Note. The figure plots indexes of production per employee of houses (red line, bottom in 2022), cars (green, middle), and all manufacturing goods (purple, top). Series are indexed to 100 in 1967, and the y-axis uses a base-10 log scale. Cross-shaped markers are used to denote years in which the denominator of the housing series (the number of general building contractors) was estimated through an out-of-sample forecast that assumes a linear trend in the share of general building contractors as a total of all construction employees (see Appendix A.5.3 for details).

firms produce roughly as many houses per employee as they used to almost 90 years ago (e.g. 2022 vs. 1939), manufacturing output per employee grew by ten-fold over the same period, and automobile output per employee rose by 400 percent over the same period. In 1939, we estimate that an individual general contractor produced about 0.96 new homes a year, a similar number to 2022 (0.98). An employee engaged in motor vehicle production in 1939 contributed about 4.82 cars, which increased to around 25 cars in 2020.

2.1.4 Regulation

Figure 6 plots a measure of regulation developed by Ganong and Shoag (2017): the number of land-use legal cases per capita. The figure also shows the logarithm of the index of housing units per employee divided by the index of car production per employee. Regulation rose after 1973, roughly when construction and car productivity started de-

Figure 6: Decoupling and Regulation



Note. The red line (bottom in 2010) plots the log of the ratio between the index of housing units per employee and the index of cars per employee (reported separately in Figure 5). The dark yellow line plots the number of land use cases per capita, an index of land use regulation from Ganong and Shoag (2017). Cross-shaped markers are used to denote years in which the denominator in the housing units per employee series was estimated through an out-of-sample forecast (see Appendix A.5.3 for details).

coupling. Appendix Figure A3 shows that the same conclusion holds if we use a more direct measure of regulation, the fraction of municipalities in California with land-use regulations (Jackson 2016).

2.2 International Comparisons

We now turn to a comparison of construction costs in US cities and elsewhere. We use data from the construction-consulting firm Turner & Townsend (2022) on building costs for different types of buildings in 31 cities around the world. Like the R.S. Means data, these costs include only the hard costs of physical construction, and exclude land and other costs. This data tells us both the total cost per square meter, as well as the cost of inputs, which allows us to try to separate construction productivity from input costs.

Input costs distinguish among labor, material, and plant costs. Labor costs are then further divided among costs for different types of workers (e.g. plumbers and elec-

tricians, carpenters and bricklayers, carpet layers and tilers). Material costs include a host of different materials (e.g. concrete, steel beams, tempered glass panes), with the plant cost defined as the cost of hiring a fifty ton mobile crane and its operator for a day. Following the methodology developed by Langston (2014, pp. 325–26) for Turner & Townsend data, we aggregate the labor and materials subcomponents to create the daily cost of standardized baskets of labor and materials. Langston (2014) selected quantities of each input so that they would approximately reflect the typical budget shares of labor (40%), material (50%), and plant costs (10%) for a building of 20 floors or higher at 2012 Sydney prices.¹⁵

Using these bundles, we focus on the cost of constructing an office building of up to 20 floors in the central business district (CBD), and estimate the following regression.¹⁶

$$\text{Office costs}_c = .171 \times \text{Labor costs}_c + .085 \times \text{Materials costs}_c + .366 \times \text{Plant costs}_c + \varepsilon_c$$

(.03)
(.02)
(.10)

(1)

Unsurprisingly, input costs account for much of the physical construction costs of a 20-floor office building, as a regression of actual vs. predicted costs according to the regression above yields a R^2 of 86%.¹⁷ These three cost bundles were, after all, designed to match the cost of constructing a structure similar to a 20-floor office building.

Table 1 reports the decomposition above for the five US cities in the sample, one “average” city in each continent, and a simple average for the world.¹⁸ The first column reports the raw cost coming from the survey and is followed by the split across the different components as estimated by the regression above, and a residual component

¹⁵See Appendix A.4 for details on the data and on the construction of the bundles.

¹⁶The Turner & Townsend data reports costs for many other types of buildings: high-rise prestige offices in the CBD, shopping malls, high-rise apartments, townhouses, and others. We focus on office buildings of up to 20 floors in the CBD because these seem most similar to the sample of high-rise buildings that Langston (2014) uses to create the input bundles. Furthermore, the restriction that they are located in the CBD helps in taking out differences in where different types of buildings are located, thus making the sample more comparable. Finally, office cost estimates, as opposed to those for residential properties, may be more properly comparable, since the prevalence of different types of housing varies across cities. For example, Turner & Townsend (2022) do not provide cost estimates for the single-family detached houses that are most common in the United States.

¹⁷Appendix Figure A4 reports a scatter plot of actual vs. fitted construction costs.

¹⁸The average city is defined as the closest to the continent average. The sample does not cover all major cities in the continent; for instance, the African continent is merely represented by the cities of Johannesburg and Nairobi. Predicted costs based on inputs explain quite well the total cost, as shown by Appendix Figure A4; and the predicted and actual cost per square meter have a correlation of 92.82%.

Table 1: International Building Costs for a 20-Floor Office Building

	Total Cost	Labor Costs	Material Costs	Plant Costs	Unexplained Residual	WRLURI
	\$/m ²	\$/m ²	\$/m ²	\$/m ²	\$/m ²	
New York	6,994	3,037	1,500	1,645	813	1.05
San Francisco	6,540	3,100	1,219	1,535	686	1.22
Los Angeles	5,602	2,430	1,232	1,287	653	0.65
Chicago	4,642	1,985	1,275	1,316	67	-0.12
Houston	2,949	1,628	1,275	879	-834	-0.13
Paris	3,107	1,492	983	558	74	
Singapore	2,437	669	1,060	602	106	
Johannesburg	1,006	140	793	313	-241	
São Paulo	751	159	708	546	-662	
World Simple Average	3,004	1,182	1,137	706	-21	

Note. The table reports a decomposition of total building costs across labor costs, material costs, plant costs, and an unexplained residual as described in Equation (1). All costs are expressed in dollars per square meter. The last column adds, for the US cities in the sample, a measure of land-use regulatory tightness from Gyourko, Hartley, and Krimmel (2019).

based on the same regression. For US cities, we also include a measure of regulatory strictness: the Wharton Residential Land Use Regulation Index (WRLURI), discussed in Gyourko, Hartley, and Krimmel (2019).¹⁹

The table shows that US cities are the most expensive places to build offices. New York is first, closely followed by San Francisco, and then by Los Angeles. Even Chicago has far higher costs than those found in Paris and Singapore. Note that this higher cost of building is not entirely explained by the cost of inputs, as the unexplained residual is always in the order of 10% to 20% of the total cost. Finally, US cities where building is expensive are also those with the tightest land-use regulations, which suggests a link

¹⁹ The index comes from a survey of over 2,000 primarily suburban jurisdictions, in which the authors ask several questions on the presence of various forms of residential building restrictions that local governments may use (e.g. on the maximum number of permits, zoning laws, density restrictions, etc.). It is standardized, with higher values indicating tighter regulation. See Gyourko, Saiz, and Summers (2008) and Gyourko, Hartley, and Krimmel (2019) for a full discussion of the measure. Here we use the 2018 values of the index (Gyourko, Hartley, and Krimmel 2019) since the date of the Turner & Townsend survey is 2022, while in the rest of the paper we use the 2006 value (Gyourko, Saiz, and Summers 2008). Either version gives the same relative ranking across cities.

between regulation and productivity that we will explore throughout the paper.²⁰

2.3 Size of Construction Firms and Projects

We conclude this section by discussing the size distribution of construction firms and projects today.

2.3.1 Firm Size Distributions

We use two datasets to measure the distribution of firm sizes. First, we take public data from the national aggregates of the 2012 Statistics of US Businesses (SUSB), which allow us to plot firm size distributions across broad construction sub-sectors.²¹ Second, we build a unique sample using the confidential base responses of the 2012 Census of Construction Industries (CCI), which allows us to analyze firm size distributions across detailed subsectors.

2.3.1.1 Evidence from Aggregate Data


Figures 7a and 7b report the shares of employment and total receipts across different employment size classes and industrial sectors, including new single-family housing construction, manufacturing, services, tradables, and nontradables (excluding construc-

²⁰Other analysis reported in the appendix is consistent with this pattern of findings. For example, Appendix Table A2 reports the same exercise for high-rise apartments, showing similar qualitative patterns for total costs, although cost differentials controlling for variation in input costs are not nearly as large for this property type in markets such as New York City and Chicago. That said, the underlying coefficient estimates on apartments are much noisier, with the coefficient on materials being the only statistically significant one. One explanation is that Langston (2014) designed cost bundles that reflect the technology used to build skyscraper-like tall buildings, which may be more representative of office buildings than apartment buildings. Finally, Appendix Figure A6 abstracts away from the Turner & Townsend bundles and simply presents scatterplots of total raw costs against GDP per capita at the country level, for both offices and apartments, which confirm the results presented here. While GDP per capita explains much of the variation in the construction costs, several American cities are notable outliers, especially when it comes to building offices. Building costs for offices in New York and San Francisco are particularly high. Construction costs for apartments are somewhat lower, especially for New York and Chicago.


²¹The SUSB is an annual series of the U.S. Census Bureau which supplies both national and subnational data on sector-specific distributions across establishment size of many indicators, including employment.

Figure 7: Firm Size Distributions

(a) Employment Shares



(b) Receipt Shares



Note. The figure plots the share of total employment (panel a) and receipts (panel b) accounted for by firms in different size bins, across different sectors. Data from the 2012 SUSB.

tion).²² The value of approximately 0.4 for new single-family housing in the 0-4 bracket in Figure 7a indicates that approximately 40% of all employees who build single-family homes work in firms with four or fewer employees. The value of approximately 0.8 for manufacturing in the 500+ category indicates that approximately 80% of total revenues in manufacturing were earned by the largest category of manufacturing establishment.

Figure 7a shows that the average single-family residential construction firm is much smaller than the average firm in the other industries. More than 63% of employees in New Single-Family Housing Construction work in establishments of firms with less than 10 employees. In stark contrast, fewer than 13% of workers in manufacturing, tradables, services, and non-tradables work in such small firms. Establishments of firms with more than 100 employees are a rarity in the single-family residential construction sector, whereas they make up the bulk of employment in the other industries.

Our results on receipts show the same pattern. Approximately 60% of the revenues in the new single family housing construction subsector accrue to firms with less than 10 employees, and less than 13% of such revenues are generated by firms with more than 100 employees. In manufacturing and tradables, about 80% of revenues are generated by firms with more than 500 employees. More than one half of revenues for services and other non-tradables also went to the largest firms. Appendix A.6 shows that these patterns hold also for the number of establishments, number of firms and annual payrolls. Figure A9 in Appendix A.6 shows that multifamily builders are also small. In that sector, the bulk of employment and receipts are in firms between 20 to 99 employees, which is larger than for single-family builders but still considerably smaller than in the other industries.

2.3.1.2 Micro-Evidence from the Census of Construction Industries

We use microdata from the CCI to analyze firms size distribution patterns across more fine-grained sectors.

²²Our list of tradables is: "Agriculture, Forestry and Fishing", "Mining", "Manufacturing", and "Management of Companies and enterprises". Our list of nontradables includes: "Retail trade", "Real Estate and Rent Leasing", "Health Care", "Accommodation and Food Services", and "Other Services."

Data Description.

The 2012 CCI dataset contains establishment-level operating data from more than 100,000 firms. The vast majority of firms in the construction industry have only a single establishment. We restrict the sample to observations that were used in official CCI tabulations, that possess sample weights, and that link to a firm in the Longitudinal Business Database. We require that establishments in the sample have non-zero employment as reported on the CCI form, non-missing breakouts of revenue by sector (described further below), non-missing values for all operating data used in profit/productivity calculations, and be located in a CBSA with a 2006 WRLURI value. These restrictions produce a sample of approximately 107,000 firms and their establishments, which we use throughout our later analyses.²³ All our results are very similar if these sample choices are relaxed.²⁴ A key feature of the CCI data is that the establishment-level report of construction revenues is split between 31 types of activities (e.g., “single-family homes, detached,” “bridges and elevated highways,” “decks, residential types”). This detail allows us to analyze detailed construction sectors. We group the reported variables into the following seven categories: housing, consumer-facing buildings (e.g., restaurants, retail stores), office buildings, warehouses, industrial/manufacturing buildings, other buildings (e.g., dormitories, schools, hospitals), and non-building construction (e.g., bridges, highways, sewage plants).²⁵

We calculate the share of revenues across these seven bins for each establishment and firm. The CCI asks firms to report their revenues for activities that exclude the cost of land and other items installed that are not part of the building structure. In graphical analyses, we collapse warehouses and consumer-facing building, and assign firms to a single specialization bin based upon the majority sources of revenues. We also introduce an eighth category for firms that do not earn most of their revenues from a single sector. These specialization bins are thus mutually exclusive and collectively

²³Here we will present graphical firm size distributions. The next sections present regressions of firm profitability by employment size and sector of operation, and regressions of construction traits at the CBSA level by regulation levels.

²⁴Observations counts through this paper are rounded per Census Bureau disclosure requirements.

²⁵See pages 8-10 of the CCI form. The housing category is codes 316 to 318. Consumer-facing building is codes 324, 326. Office buildings is code 325. Warehouses is code 327. Industrial/manufacturing buildings is codes 321 and 323. Other buildings is codes 319, 328 to 334, and 338. Non-building construction is all codes within category B in the form.

exhaustive.²⁶ These choices mostly follow from disclosure requirements, and we note below robustness to other approaches.

Results

Figures 8a and 8b present firm size distributions (FSD) across construction sectors. Different types of construction are in different shades of blue to purple.²⁷ In green, we show the firm size distributions for non-construction firms based on the 2012 LBD data. The latter sample contains every LBD firm that is not in the CCI sample, and has a modal establishment that is not in the construction sector (NAICS 23).

The horizontal axis of each graph groups firms by employment levels, using the same increments used before: 0-4 employees, 5-9, 10-19, 20-99, 100-499, and 500+. The vertical axis shows the share of sector employment and revenues that are accounted for by each employment level (as above, the distribution sums to 1). Appendix Figure A12 reports payroll and firm counts.²⁸

Differences between construction firms and non-construction firms remain stark. Around 60% or more of the employment, revenue, and payroll of non-construction firms is accounted for by firms with 500+ employees, despite their being a small share of the number of total firms, which confirms that “the typical firm is small, but the typical employee works in a large firm.” The 500+ employee bin for non-construction firms in the LBD often contains 13 to 18 times the activity accounted for by the 0-4 employee bin.

The construction sector is different. Firms specialized in housing construction are the most extreme, with firms of 0-4 employees accounting for the largest share of employees and revenues.²⁹ In most other building construction sectors, the activity in the smallest size bins is comparable to the activity in the 500+ employee bin. For exam-

²⁶In regression analyses at the firm level, we will model these revenue shares as continuous variables.


²⁷Here we focus on housing, consumer-facing buildings, industrial buildings and warehouses, office buildings, and non-building construction. Appendix Figure A13, reports FSD results that also include the other two subsectors: other buildings construction (dormitories, schools, etc.), and firms with no clear specialization (i.e. that do not have more than 50% of their revenues in one type of construction).

²⁸All graphs use sample weights. In the case of multi-establishment firms, we take the minimum weight across the establishments of the firm. Results are very similar under alternative techniques.


²⁹Results on payrolls are reported in Appendix Figure A12. There the 20-99 bin accounts for the highest share (22%), but the 0-4 bin is the second-highest (20%).

Figure 8: Firm Size Distribution

(a) Employment Shares




(b) Revenue Shares



Note. The figure plots the share of total employment (panel a) and revenues (panel b) accounted for by firms in different size bins, across different types of construction firms (in shades of blue to purple) and compared also to all nonconstruction firms (in green). Data for 2012 from the LBD. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

Figure 9: Share of Housing Units Built by Employment Size



Note. The figure plots the share of total housing built accounted for by firms in different size bins. Microdata from the 2012 Census of Construction Industries (CCI). This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

ple, firms with 0-4 employees account for slightly more employment in office building construction than firms with 500+ employees.³⁰

We report firm counts in Appendix Figure A12, where we again observe that most construction firms are small. Most non-construction firms in the LBD are also small, but the construction sector is particularly skewed to smaller firms.

The CCI survey asks firms to report units built (single- and multi-unit residential). Most firms specializing in the housing sector according to their revenues have zero-valued unit counts, presumably reflecting the fact that the firm does not build entire houses but rather components of housing (although in some cases, firms may have chosen not to report the data). Figure 9 plots the distribution of units created across the firm size distribution.

³⁰Investigations with the LBD also confirm the uniqueness of the construction sector. When we look at the FSD by major economic sector in the LBD, the role of small firms is the most significant in the construction sector.

2.3.2 Size of Projects

Small firms go hand in hand with small projects. We conclude this section by providing evidence on the size of construction projects today, and comparing that size with the size of mega-projects in the mid 1950s. We start from what is probably the most famous example of large-scale construction in the US: Levitt and Sons.


In 1947 Levitt and Sons acquired 1,400 acres of Long Island farmland with the idea of efficiently developing thousands of nearly-identical single-family houses. By 1948 the firm was completing more than 35 houses per day or 135 per week. Ultimately, just over 17,000 houses were built and sold for an average price of \$7,990 (\$100,886 in 2023 dollars). By 1950 Levitt was a household name and later in the 1950s they repeated this style of development in Bucks County, PA, building 17,300 homes on 6,000 acres.

To do this, they invented a type of assembly line, but rather than moving the house along a line, they moved construction crews along identical homes—thousands of them. They broke down the construction process into 26 specific component parts and had a team for each of them, used time and motion study techniques, and brought some new processes to homebuilding; they also tried to preassemble as much as possible off-site. Clearly, they became very productive at it, making profits of about \$1,000 per home (roughly \$13,800 today; Rybczynski 2017).³¹

Are there any Levitt and Sons today? To describe how projects look today, we use data from Gyourko and Krimmel (2021), who collected information on land parcels

³¹While the Levitt brothers are the most famous of the early post-WWII homebuilders, they were not unique by any means; rather, they were trend setters. Checkoway (1983) notes other large builders who used Levitt-type production strategies to rapidly construct hundreds or thousands of new homes in the 1950s and 1960s. They existed in a wide range of markets across the country including Baltimore (John Mowbray), Washington, D.C. (Waverly Taylor), Toledo, OH (Don Scholz), Cleveland (Maurice Fishman), Chicago (Irvin Blietz), Kansas City (J.D. Nichols), Phoenix (Del Webb), San Francisco (Carl Gellert and Ellie Stoneson). Checkoway (1983) also notes that a few builders such as Dave Bohannon, Fritz Burns and James Price actually replicated the strategy across multiple markets. Checkoway (1983) concludes that three factors distinguished this new wave of builders: their size, their lower costs and their suburban focus. He argues that these three traits combined to allow a doubling of the number of new housing starts in the 1950s compared to the 1940s (i.e., 15.1 million starts from 1950-59 versus 7.4 million from 1940-49) without engendering rising real costs that could have made the homes unaffordable. Using data from the San Francisco Bay Area, Maisel (1953) was the first to document both lower costs and higher profits for larger builders who could employ their new production techniques across (potentially) thousands of single-family residential parcels located within a single community on expansive tracts of vacant suburban land.

Figure 10: The Size Distribution of Vacant Land Purchases Intended for Single-Family Housing Development (Cumulative Distribution Function for the Share of Parcels Below Given Square Footage Amounts)



Note. The figure plots the cumulative distribution function of parcel sizes. The underlying data are vacant land purchases intended for single-family housing development for 24 CBSAs over the years 2013–2018. The plot is based on 3,640 observations of vacant land parcel purchases. The individual observations were downloaded from proprietary CoStar files and used in Gyourko and Krimmel (2021). See their paper for more details. There are 43,560 square feet in one acre.

purchased for the expressed purpose of single family development across 24 metro areas over the 2013–2018 period. In this period, the largest US single-family residential land parcel purchased was a 1,049 acres site north of Denver: one-sixth of what the Levitts were doing. Figure 10 reports the cumulative distribution across parcel size. The median project is below 10 acres, and the 99th percentile of the parcel size distribution is 314 acres. Projects with more than 500 acres are essentially non-existent. This fact is also true if we restrict our attention to places that have large amounts of empty land around—thus casting doubt on the idea that smaller project sizes are due to the fact that land now is scarcer compared to the 1950s. Atlanta is one example of a place with abundant land, and Appendix Figure A14 shows that Atlanta has a similar project size distribution.

Why are there so few large projects today? A possible explanation rests on the increase in regulatory tightness in construction from the mid-1970s to today. Developing large projects and coordinating construction teams over different projects all working at the same time is becoming more and more difficult. It is hard to coordinate teams across different sites if permits get released asynchronously, so that developers do not know

with certainty when they will be able to work on which plot of land. It is also hard to obtain permits to develop a single type of housing unit on large plots of land given that zoning laws discourage such types of projects, which makes it virtually impossible to adopt the assembly line approach of the Levitts.

We next turn to a model that makes this link formal, deriving the prediction that land use regulation leads to smaller firms and less investment in construction technology. We then test the implications of the model in the cross-section.

3 A Model of Regulated Construction

This section provides a simple theoretical model that highlights how regulation of projects, but not regulation of firm entry, causes firms to be stunted in both size and productivity. Three core assumptions generate this result: (1) contractors have a limited span of control, which makes it difficult to manage many small projects, (2) the benefits of endogenous investment in technological know-how rise with company scale, and (3) the ability of the entrepreneurs that enter the market does not rise dramatically as the number of firms shrinks.

3.1 Market Structure, Technology and Regulation

In a given location, a finite pool of N potential builders decide whether to become real estate developers or to receive their outside option ψ . Expected ability in the construction industry is *ex ante* identical across builders, but their outside options vary and are described by the cumulative distribution function $F(\psi)$. If the builder chooses to become a developer, then they forgo their outside option and must pay a fixed cost Ψ . These fixed costs are meant to capture, in part, the regulatory barriers to entry that are measured by Djankov et al. (2002). We do not take a stand on whether these barriers were designed to entrench incumbents (Stigler 1971) or serve some larger social purpose, but we distinguish these upstream barriers to entry into the industry from downstream barriers to individual projects.

After a builder enters the market, they observe their productivity potential A ,

which is independent of the foregone opportunity ψ and independently drawn from a Pareto distribution with minimum $\underline{A}$ and shape α , hence mean $\alpha \underline{A} / (\alpha - 1)$. After observing their productivity potential, the builder chooses how much to spend on technology: we denote that spending level by K . Investment in technology and productivity potential together determine the physical costs of construction.

Having learned their productivity potential and invested in technology, the builder can develop projects. The building process involves partnerships between risk-neutral developers and risk-neutral owners of building parcels. The developer and parcel owner first propose a project to the project regulator, which could be a local zoning board. The parcel owner commits to a payment to the developer conditional upon the project being approved and built. The regulator then either approves or rejects the project. If the project is approved, then the project gets built, its building units get sold and the developer receives the contracted payment. If the project is rejected, then the parcel has no value to either the developer or the parcel owner.³²

More formally, developers can partner with many parcel owners. There are N_L developable parcels that are identical for developers, but whose owners have heterogeneous reservation values. Reservation values are distributed so that the number of parcels with reservation values less than any amount p equals $N_L p^\sigma$. We will let p_L denote the reservation value for the marginal parcel, which in a competitive equilibrium will also equal the expected return to any parcel owner who participates in partnerships with developers.³³ At the time of contracting, the developer and parcel owner agree upon both the payment conditional upon project approval and the size of the proposed project, b .³⁴

The local regulator then decides whether to approve the project. A project of size

³²Our model assumes that developers are contractors building on land they do not own, which is true in many cases. It is equivalent to an alternative structure in which merchant builders first buy parcels and then propose to develop them, with the parcels becoming worthless to the developer if their project is rejected.

³³Likewise, it would equal the market price of parcels if land were bought by merchant builders.

³⁴These must be simultaneously agreed upon for the partnership to function smoothly, because once the post-development payment to the parcel owner is fixed, then the interests of the developer and the parcel owner diverge.

b is approved with probability:

$$l(b) = \min \left\{ \left(\frac{b}{\underline{b}} \right)^\rho, 1 \right\}. \quad (2)$$

While the fixed cost Ψ measures the regulation of entry, the tightness of project regulation is captured by $\rho \in (0, 1)$. This measure is an index because $\rho = 0$ implies that all projects are approved and $\rho = 1$ implies that it is impossible for the expected number of units in a project to be greater than the minimum number $\underline{b}$.

If the project is approved, b units are developed on the parcel at a cost of

$$m = \frac{\varepsilon}{z} b^{1+\frac{1}{\varepsilon}}. \quad (3)$$

The parameter ε is an inverse measure of the extent to which costs increase with the scale of each project. Project productivity z depends on the builder's span of control, namely the number s of projects that they are supervising. The s -th project a builder undertakes has productivity:

$$z = AK^{\frac{1}{\kappa}} s^{-\frac{1}{\omega}}. \quad (4)$$

Overall productivity combines the developer's idiosyncratic productivity potential (A) with the developer's investment in technology (K), whose returns are governed by the parameter κ . A core assumption in our model is that it is difficult to supervise a large number of projects. The developer's time is limited and, if they are constantly shuttling between small projects, it is harder for them to keep watch over the costs in any one project. We model the limited span of control by assuming that costs are higher for the second project than the first and for the tenth project than for the fifth project; more formally, costs are multiplied by $s^{\frac{1}{\omega}}$ for the s -th project.

After the project is completed, units are sold in a competitive housing market at equilibrium price p_B . Total spending on housing in the location is a constant fraction δ of aggregate income Y in that location. We assume that both developers and absentee landowners live outside the location and do not buy housing, and that aggregate income in the location is fixed. Consequently, the equilibrium price of one unit of physical space, which we will refer to as "a building," equals $p_B = \delta Y / Q_B$, where Q_B is the total number of buildings built in equilibrium.

We assume that $\alpha\kappa / (\alpha + \kappa) > \omega > \varepsilon(1 - \rho)$, which guarantees that a finite number of projects yields a finite number of buildings, and that average contractor size is finite. Since the first inequality also implies that $\kappa > \omega$, this assumption also ensures that the returns to technology adoption rise more slowly than the cost of technology adoption, which guarantees that the technology-choice problem has a unique, interior solution. We also assume that the amount of construction that would be permitted with certainty is negligible, in the sense that it can never be value-maximizing to propose a project that is certain to be permitted: formally, $\underline{b} \leq [1 + \varepsilon(1 - \rho)] p_L / p_B$.

Appendix B provides a closed-form solution for the model and all proofs of the propositions. In equilibrium, all builders choose to undertake projects up to the point where the productivity of the marginal project hits a threshold $\underline{z}$ that is homogeneous across builders. Conditional upon undertaking a project, its optimal proposed size is also determined and increasing in project productivity (z). Contractors also choose optimally their technology investment, and there is a maximum outside opportunity cost that determines whether a potential developer enters the market. Similarly, there is a maximum reservation value for land owners that determines their willingness to partner with a developer. Subsection 3.2 provides comparative statics on productivity potential (A) across builders, holding the equilibrium prices and total quantities constant. Subsections 3.3 and 3.4 provide comparative statics on the place-level variables that shape the equilibrium.

3.2 Contractor Heterogeneity

Contractors with greater potential (A) and technology investment (K) can naturally handle a greater number of projects before stretching their span of control to the point at which their marginal project productivity hits the threshold $\underline{z}$. Intrinsically more productive contractors can grow to a larger scale and thus reap greater benefits from their investment. As a consequence, they are incentivized to invest more in technology.

Proposition 1 formally describes the impact of idiosyncratic productivity on developer behavior.

Proposition 1 *Across builders, the elasticity of technology investment, projects undertaken,*

units built, and revenues with respect to productivity potential (A) is identical and equal to $\kappa\omega / (\kappa - \omega) > 0$. Builders with higher productivity potential also have a higher ratio of revenues to total costs.

The exogenous driver of differences across contractors is the realization of their idiosyncratic potential (A). Contractors with greater potential optimally choose greater technology investment and predictably end up being both bigger and more productive. Our measure of their productivity is the ratio of revenues to total costs (excluding the entrepreneur's opportunity cost). Capital, variable costs and total revenues all have exactly the same elasticity with respect to productivity potential (A), but fixed costs (Ψ) do not change with A . Consequently, fixed costs are smaller relative to revenues in bigger firms. Entrepreneurs with higher values of A lead bigger firms, in which fixed costs are relatively less important and productivity is consequently higher. As long as the share of labor in the input bundle is independent of firm size, then larger firms—with a higher value of A —also have a higher revenues per employee.

The link between size and productivity reflects two-way causality. Innately more productive firms take on more projects and build more in each project, but firms that anticipate taking on more projects also invest more in technology. When we simulate the impact of a shift in the firm size distribution on productivity in the construction sector, we will have to make an assumption informed by this model about how much of the observed empirical link between size and productivity reflects the causal effect of productivity on size.

3.3 The Impact of Project Regulation

We now turn to the market-level impact of regulation on firm size and productivity. Propositions 2 and 3 look at the impact of project-level regulation, which is captured by the permitting parameter ρ . Proposition 4 looks at the regulation of entry, which is captured by the fixed cost parameter Ψ . In these propositions, parameter changes cause the equilibrium prices to change, unlike in Proposition 1 that looked across firms within a given equilibrium.

We begin by emphasizing the equilibrium price and quantity effects of project-

level regulation in Proposition 2.

Proposition 2 *Stricter project regulation (higher ρ) reduces the equilibrium number of active contractors and the number of buildings built, and increases the equilibrium prices of buildings and land parcels.*

More restrictive project regulation increases the cost of proposing large projects and reduces profitability in the construction industry holding prices constant. Consequently, contractors with better outside options abandon the industry, which means that there are fewer contractors in equilibrium. Moreover, since proposing larger projects becomes riskier as ρ increases, any given project becomes smaller. Remaining developers do take on more projects, but the contraction on the extensive margin (the declining number of developers) unambiguously dominates. Variation on the intensive margin (the number of units per developer) is ambiguous, but the total supply of building certainly shifts downwards. As demand is unchanging, this shift leads to a reduction in the quantity of homes built and an increase in the price of homes.

The impact of tighter project regulation on land value is perhaps the most surprising part of Proposition 2. As developers spread their housing units across a larger number of land parcels, their demand for land parcels increases, which means that the equilibrium price of land parcels also increases. This result provides another political-economy explanation for the popularity of zoning: the owners of existing land parcels can become richer if regulatory constraints on project size lead developers to buy up more parcels of land. This case also provides a counter-example to the Henry George Theorem (Arnott and Stiglitz 1979) that land value maximization will generate welfare maximization. In this case, regulation has an effect on market power that transfers rents from both developers and home buyers to landowners.

Proposition 3 turns to the link between project regulation and builder size and productivity.

Proposition 3 *The elasticities of developers' average technology investment and average revenues with respect to the tightness of project regulation (ρ) are equal and negative. A developer's average number (and a fortiori value) of projects are increasing in regulatory tightness, while their average ratio of revenues to total costs is decreasing in regulatory tightness.*

More restrictive regulation hinders contractors' operation by forcing them to undertake inefficiently small projects. They undertake inefficiently many projects, which stretches their span of control. The price per unit increases, but nonetheless they earn lower revenues relative to cost.

Deprived of the ability to operate at scale, contractors cannot reap the full benefits of their technology investment, and they react by investing less. Reduced technology investment naturally entails lower productivity, whether it is caused by lower idiosyncratic potential (Proposition 1) or by tighter regulation (Proposition 3). Consequently, the average ratio of revenues to costs falls with regulatory tightness. Assuming that labor's share of costs is unchanged by either regulation or firm size, tighter regulation reduces average revenues per employee. We will test these implications in the paper's next section.

These intuitive aggregate impacts also underscore the importance of measuring contractor size properly. Contractors build on land owned by the investors or homebuyers who hire them. Thus, land value is not included in their revenues nor in their costs. As Proposition 3 established, tighter regulation reduces average contractor revenues but increases the average value of projects undertaken by each contractor. With unchanging total expenditure on buildings the latter effect must be the stronger, so each contractor must produce a greater average value of buildings while earning lower revenues.

3.4 Regulation of Entry vs. Project Regulation

Land-use regulation is different from the classic kind of regulation discussed in Stigler (1971) or Djankov et al. (2002). Such entry regulations are typically imposed at the firm level, and are presumably best seen as a fixed cost that must be paid by the firm. Here we contrast the impact of project-level regulation—discussed above—which leads to overly many, overly small projects and too little investment in firm-wide technology, with the impact of a firm-level regulation that acts as a barrier to entry into the industry.

Proposition 4 *The elasticities of average technology investment, average revenues, the average number and average value of projects with respect to entry costs (Ψ) are equal and positive. The average number and value of buildings per developer all increase with the entry cost, as does the*

price of a housing unit. The cost of each land parcel and the total number of projects undertaken are independent of entry costs. The number of active builders declines with entry costs. The impact of entry costs on the ratio of revenues to total costs is ambiguous.

Higher fixed costs induce exit of marginal firms, but total spending on housing remains unchanged. Consequently, each firm must have higher revenues. Larger firms build more structures and invest more in technology that can be spread across all of their projects. Regulatory barriers to entry do reduce supply, which intuitively harms consumers, but there is at least a countervailing force: larger incumbents have an incentive to invest more in technology. With project-level regulation that reduces firm size, the added effect from technology is negative too, as smaller firms end up being less efficient.

The overall impact of regulation of entry on the ratio of revenues to costs is ambiguous. Greater investment in technology increases equiproportionally revenues, variable costs, and the cost of technology investment itself. Thus, it increases the ratio of revenues to costs if and only if all these are rising more than proportionally with fixed costs—which may or may not be the case, depending on the magnitude of fixed costs and on the distribution of developers’ outside opportunities.

Perhaps the most surprising part of Proposition 4 is that the number of total projects undertaken, and consequently the land cost per project is unchanged. This results follow from two properties of our model. First, total expenditure on buildings is constant. Second, the equilibrium share of building value that accrues to landowners is independent of fixed costs: it reflects only the tightness of project regulation (ρ) and diseconomies of scale ($1/\varepsilon$) and scope ($1/\omega$). These two results jointly imply that the total value of developed parcels is itself independent of fixed costs. With unchanging supply, that means both price and quantity must be unchanging.

In the next section, we turn to empirical support for the predictions of our model.

4 Cross-Sectional Evidence

We now turn to testing the three main cross-sectional implications of the model: (i) smaller firms are less productive, (ii) tighter project regulation is associated with smaller

firms, and (iii) tighter project regulation is correlated with less productive firms.

4.1 Productivity and Firm Size in Housing Construction


We first estimate housing units per employee and revenues per employee for firms of different sizes. We use the distribution of employment, revenues, and housing built across firms in different size bins (0-4 employees, 5-9, etc.), which were calculated using Census microdata and reported above in Figures 8a-8b and 9. For each set of firms in a given size bin, we compute the ratio between their output, measured in both houses and revenues as a share of total industry output, with their employment, again as a share of total industry employment. This measure, reported in Figure 11, is proportional to units or revenues per employee, and we normalize it to 1 for firms in the 0 to 4 employee category. All productivity measures are consequently relative to productivity in the smallest firms.³⁵

Both revenues and units per employee increase steeply across the size distribution. Firms with 20 to 99 employees produce 55% more units per employee than the smallest firms. Firms with 100 to 499 employees produce more than double the units per employee, while employees in firms with more than 500 employees produce almost six times as much.

We also use the microdata to estimate the cross-sectional relationships between firm size and profits and revenues per employee. We estimate firm profits by subtracting spending on payroll, benefits, normal depreciation charges, rental payments, materials and supplies, contract labor, fuels, electricity, and reported other purchases from firm revenues. This measure is negative for a significant number of companies. We transform it into a z-score, by subtracting the mean and dividing by the standard deviation of calculated profits.

³⁵Note that here we use the employment distribution for firms that are specialized in housing construction, rather than the employment distribution for the universe of construction firms. However, the housing units distribution is coming from all firms in construction. Since the housing-specialized ones are those that are responsible for the majority of house-building, using the employment distribution for such firms gives us a more reliable estimate. For revenues, the measure is fully consistent as we take both the revenues and the employment distribution for firms specialized in housing.

Figure 11: Output and Revenues per Employees



Note. The figure plots housing units (top) and revenues (bottom) per employee for construction firms in different size bins. Microdata from the 2012 Census of Construction Industries (CCI). This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

Table 2 provides descriptive statistics on the sample.³⁶ Table 3 reports the regression results of:

$$y_i = \alpha + \beta \times \log(\text{Empl}_i) + e_i, \quad (5)$$

where y_i is profits, in Column (1), and the log of revenues per employee, in Column (2). Regressions are unweighted and report robust standard errors. Across firms, we see that a 10% increase in firm employment corresponds to around a .014 standard deviation increase in profits and a 1.1% increase in labor productivity.

In Appendix Table A6 we illustrate heterogeneity across different types of construction activity. The link between firm size and productivity is strong and meaningful for housing, but it is even larger for other forms of construction. Intuitively, the forms of construction where smaller firms are more prevalent (e.g. housing) are also those that show the weakest link between firm size and productivity. These patterns are robust across many specification variants.

³⁶The formula using CCI form numbers: $100 - \text{sum}(300, 223, 540, 550, 421, 423, 431-434, 425, 449)$. The results are robust to dropping non-cash expenses like depreciation.

Table 2: Descriptive Statistics from the Firm-level Sample

	Mean (1)	SD (2)
Firm profits, estimated	545.9	5346
Labor productivity	238.7	310.4
Log labor productivity	5.066	0.8543
Employment	32.1	197.4
Log employment	2.325	1.396
<i>Firm revenue composition</i>		
Share housing	42.38	44.36
Share industrial buildings	6.55	18.85
Share consumer-facing buildings	9.38	21.34
Share office buildings	7.38	18.19
Share warehousing	2.06	9.14
Share other buildings	13.29	25.56
Share non-building construction	18.96	36.86

Note. Microdata from the 2012 Census of Construction Industries (CCI). The sample has 107,000 observations (rounded per Census Bureau disclosure requirements). Firm profits are estimated by subtracting from firm revenues the amount the firm spent on payroll, benefits, normal depreciation charges, rental payments, materials and supplies, contract labor, fuels, electricity, and reported other purchases. Values are in thousands of nominal dollars in 2012; values can be negative. Labor productivity is measured as revenues per employee. Firm revenue composition is developed by aggregating establishment-level reporting of their construction revenues split out by 31 types of activities (e.g., “single-family homes, detached,” “bridges and elevated highways,” “decks, residential types”). This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

4.2 Counterfactual Productivity in Construction

We now ask how much construction productivity would increase if construction firms were as large as those typically found in either manufacturing or (other) nontradables. We observe the relationship between size and productivity, but as we discussed in Section 3, this relationship reflects both the impact of size on productivity and the impact of exogenous productivity differences on size. Lacking at present exogenous variation on firm size to identify its effects on productivity, we perform a simple yet transparent exercise, and assume that a fraction ϕ of the observed empirical relationship between establishment size and productivity represents the causal effect of size on productivity.

We let j indicate a firm-size-bin (0-4 employees, 5-9, and so on), and denote with $\{N_{0-4}, N_{5-9}, \dots\}$ the firm size distribution in construction, where N_j indicates the frac-

Table 3: Regressions of Firm Profitability and Labor Productivity on Firm Size

VARIABLES	(1) Profits in unit standard deviations	(2) Log labor productivity
Log employment	0.1375*** (0.0062)	0.1094*** (0.0020)
Observations	107,000	107,000
R-squared	0.0368	0.0319

Note. The table reports results from an OLS regression at the firm level of profitability and productivity against the log of the number of employees. Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Analysis using microdata from the 2012 Census of Construction Industries (CCI). Regressions are unweighted and have 107,000 observations (rounded per Census Bureau disclosure requirements). This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

tion of employment accounted for by firms in bin j . We let $\{N'_{0-4}, N'_{5-9}, \dots\}$ denote the counterfactual firm size distribution, where $N'_{0-4} - N_{0-4} < 0$, for example, means that in the counterfactual we are moving workers out of small firms towards other size bins. Finally, we let $\bar{a}_j$ denote the measure of output per employee for firms in bin j that we estimated in the data.³⁷ If ϕ represents the fraction of observed productivity differences across firms of different sizes that is causal, then the aggregate change in productivity, Δ , from a shift in the firm size distribution is:³⁸

$$\Delta = \sum_j \underbrace{(N'_j - N_j)}_{\text{Reshuffling of workers across firms of } \neq \text{ size}} \times \underbrace{(\phi \cdot \bar{a}_j)}_{\text{Effect of size on productivity}}. \quad (6)$$

Table 4 reports these changes for different values of ϕ . The first three columns report changes in the average units per employee, while the last three focus on average revenues per employee. If we assume that one half of the link between size and productivity is causal, construction firms would produce 91% more units per employee if their size distribution matched that of manufacturing. Even if only 10% of the link between productivity and size is assumed to be causal, this estimate would still be 18.2%. This exercise tells us nothing about why construction firms are so small, but shows that if even a small fraction of the link between firm size and productivity is causal, small firm sizes

³⁷These are the units or revenues per employee across the FSD that we reported in Figure 11

³⁸See Appendix A.8 for a formal derivation of Equation (6).

Table 4: Counterfactual Productivity in Construction under Different Assumptions on the Link between Size and Productivity (ϕ)

Counterfactual Size Distribution	Change in Units per Employee (%)			Change in Revenues per Employee (%)		
	$\phi = 1$	$\phi = 0.5$	$\phi = 0.1$	$\phi = 1$	$\phi = 0.5$	$\phi = 0.1$
Manufacturing	+182%	+91%	+18.2%	+90%	+45%	+9%
Non-tradables	+165%	+82.5%	+16.5%	+81.6%	+40.8%	+8.2%

Note. Values of ϕ indicate different assumptions on how much of the empirical relationship between productivity, defined as units per employee, and size is causal. $\phi = 1$ indicates that we assume that all of the empirical relationship between firm size and productivity is causal, $\phi = 0.1$ assumes that only 10% of it is causal. Analysis using microdata from the 2012 Census of Construction Industries (CCI). This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

may be responsible for a significant part of the underperformance of the construction sector. We now turn to the link between regulation and small firm sizes in construction.

4.3 Regulation, Productivity, and Firm Size

To examine the link between regulation and firm size, we use the 2006 version of the Wharton Residential Land Use Regulation Index (WRLURI) described by Gyourko, Saiz, and Summers (2008). The survey is at the jurisdiction level (typically at the census place level), and we aggregate it at the CBSA level by averaging across jurisdictions.³⁹

The WRLURI index is available for around 2,600 communities, which aggregates up to 550 CBSAs. Since we are missing data entirely for approximately 300 CBSAs, and in many CBSAs we have only a few jurisdictions, we also replicate our analysis using a “projected” WRLURI.⁴⁰ For the CBSAs for which WRLURI is available, we run a regression where WRLURI is predicted based on the (log of) population and population

³⁹Results are robust to weighting jurisdictions using population weights and land area weights. See footnote 19 for a longer description of the index, and Gyourko, Saiz, and Summers (2008) for the full discussion.

⁴⁰Our results are robust to restricting the sample to CBSAs where we have more than 5 jurisdictions responding to the survey. To exploit the full granularity of the data, we also predicted the WRLURI by running our predictive regressions directly at the place level and then aggregating up the predicted place-level values at the CBSA level. Results (not reported) are robust to this more granular specification and often stronger.

density in a given CBSA, the average school years of male and female residents, and fixed effects at the Census Division level.⁴¹ The predicted WRLURI has a 0.6 correlation with the raw index.

We measure firm size with: (i) the log of total receipts per establishment and (ii) the fraction of employment that is in establishments of firms with more than 100 employees.⁴² Both variables come from the Census's 2012 Statistics of US Businesses (SUSB). We analyze firm size for construction as a whole and for its subsectors (three-digit NAICS classification): construction of buildings, heavy and civil engineering construction, and specialty trade contractors.

Using these two measures of firm size and both actual and predicted WRLURI, we estimate:

$$y_i = \alpha + \beta \times \text{WRLURI}_i + X_i' \gamma + \epsilon_i, \quad (7)$$

where y_i is firm size (receipts per establishments or % of employment in large firms), X_i denotes a vector of controls including the log of population, population density and the log of (self-reported) median house values, based on the American Community Survey over the 2008-2012 time period. We run regressions at the CBSA level and we weight by population.

Table 5 reports descriptive statistics for the dependent variables and the WRLURI measures. Table 6 reports results from regressing the log of receipts per establishment on both the raw WRLURI (Panel A) and predicted WRLURI (Panel B). Areas with stricter land use regulation display lower levels of revenues per establishment. The results are stronger for the predicted WRLURI measures than for actual WRLURI measures, which

⁴¹Some CBSAs belong to more than one division. In this case, we attribute a CBSA to the state where the highest share of its population resides. For robustness, we run our prediction exercise in three other ways: (i) using our baseline variables interacted by census division fixed effects; (ii) adding a set of economic variables, the (log of the) self-reported median house values prevailing in the CBSA, residents' self-reported per capita income, the (log of the) median rent, and the (log of the) number of housing units, distinguishing between owner-occupied, renter-occupied and vacant housing units; (iii) interacting baseline and economic variables with Census Division fixed effects. Appendix Table A11 reports the results of these different specifications. In the most saturated specification, our R^2 is 0.574. Results are robust to all different specifications, as shown in Appendix Tables A13 to A20.

⁴²Due to censoring, we cannot retrieve employment counts in large firms in 50 to 100 CBSAs, depending on the particular subsector we focus on. As a consequence, we impute missing employment in large firms using the national mean of employment per firm in each bin, multiplied by the (non-censored) number of firms operating in the CBSA. In appendix section A.10 we detail for how many CBSAs we perform the imputation, and show that alternative imputation methods yield almost identical effects.

Table 5: Descriptive Statistics

	Mean	SD
WRLURI	-0.393	0.713
Projected WRLURI	-0.458	0.423
<i>All construction</i>		
Log of receipts per establishment	14.127	0.503
Frac. of employment in large firms	0.181	0.166
<i>Construction of buildings</i>		
Log of receipts per establishment	14.087	0.684
Frac. of employment in large firms	0.112	0.163
<i>Heavy and Civil Engineering</i>		
Log of receipts per establishment	15.013	0.902
Frac. of employment in large firms	0.324	0.285
<i>Specialty Trade Contractors</i>		
Log of receipts per establishment	13.786	.470
Frac. of employment in large firms	0.133	0.204

Note. Descriptive statistics on WRLURI and SUSB variables.

is compatible with the view that mismeasurement of land-use regulation leads to attenuation bias. Column (1) reports the results when considering all establishments in the entire construction industry. The coefficient implies that a one standard deviation increase in the raw WRLURI measure is associated with a 10.8% decrease in receipts per establishment. Columns (2) to (4) replicate the same analysis for its subdivisions, and show that construction of buildings is particularly sensitive to regulation in panel A, but land-use regulation has a larger impact on firm size in heavy and civil engineering construction in panel B.

Table 7 reports results on the fraction of employment in establishments of large firms. Stricter land-use regulation is associated with a larger share of employment in smaller firms. Across the entire industry, a one standard deviation increase in land use regulation (using the raw index) is associated with a 2.5 percentage point reduction in the fraction of employment in large firms (14% of the mean). In construction of buildings, again for the raw index, a one standard deviation increase in the regulation index is associated with a 4.3 percentage point reduction in employment in large establishments,

Table 6: Log of Receipts per Establishment and Regulation

VARIABLES	(1) All construction	(2) Construction of buildings	(3) Heavy and Civil Engineering	(4) Specialty Trade Contractors
<i>Panel A: original WRLURI</i>				
WRLURI	-0.1616*** (0.0456)	-0.1810*** (0.0525)	-0.1254* (0.0657)	-0.1579*** (0.0489)
ln of population (2008-2012)	0.2452*** (0.0320)	0.3221*** (0.0379)	0.3257*** (0.0403)	0.2148*** (0.0249)
Constant	10.6805*** (1.2267)	9.3842*** (1.2864)	10.9557*** (1.5175)	10.0791*** (1.2303)
Other controls	✓	✓	✓	✓
Observations	545	530	422	543
R-squared	.4224	.4328	.4154	.417
<i>Panel B: projected WRLURI</i>				
WRLURI	-0.2925*** (0.0838)	-0.2899** (0.1135)	-0.3948*** (0.1143)	-0.2188*** (0.0739)
ln of population (2008-2012)	0.2544*** (0.0302)	0.3375*** (0.0356)	0.3562*** (0.0342)	0.2174*** (0.0246)
Constant	9.9374*** (1.2057)	8.6560*** (1.3965)	8.5410*** (1.4053)	10.0033*** (1.2205)
Other controls	✓	✓	✓	✓
Observations	852	808	627	852
R-squared	.4456	.4696	.4568	.4228

Note. The table reports results from a WLS regression at the CBSA level of the log of receipt per establishment against WRLURI, both raw (in Panel A) and projected using demographic characteristics (Panel B). Each CBSA is weighted by population. Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. We use the 2006 version of WRLURI (Gyourko, Saiz, and Summers 2008). Other controls include population density and the log of self-reported home values (from the 2008-2012 ACS5A). Projected WRLURI (Set 1) is obtained projecting WRLURI on the (log of) population and population density in a given CBSA, the average school years of male and female residents, and fixed effects at the Census Division level. These regressors are taken from the 2000's Decadal Census.

Table 7: Fraction of Employment in Large Firms and Regulation

VARIABLES	(1) All construction	(2) Construction of buildings	(3) Heavy and Civil Engineering	(4) Specialty Trade Contractors
<i>Panel A: original WRLURI</i>				
WRLURI	-0.0354*** (0.0128)	-0.0607*** (0.0146)	-0.0396* (0.0217)	-0.0328** (0.0128)
ln of population (2008-2012)	0.0742*** (0.0093)	0.0797*** (0.0092)	0.1096*** (0.0111)	0.0732*** (0.0081)
Constant	-0.5453* (0.3213)	-1.1150*** (0.3390)	-0.8687* (0.4436)	-0.8772** (0.3454)
Other controls	✓	✓	✓	✓
Observations	541	488	381	532
R-squared	.3639	.3251	.3382	.409
<i>Panel B: projected WRLURI</i>				
WRLURI	-0.0618*** (0.0222)	-0.0795*** (0.0268)	-0.0963** (0.0381)	-0.0332 (0.0225)
ln of population (2008-2012)	0.0756*** (0.0086)	0.0801*** (0.0081)	0.1145*** (0.0098)	0.0721*** (0.0077)
Constant	-0.6836** (0.2907)	-1.1601*** (0.2943)	-1.2885*** (0.4695)	-0.7820** (0.3384)
Other controls	✓	✓	✓	✓
Observations	846	730	541	830
R-squared	.3719	.3302	.3673	.4079

Note. The table reports results from a WLS regression at the CBSA level of the fraction of employment in large firms against WRLURI, both raw (in Panel A) and projected using demographic characteristics (Panel B). Each CBSA is weighted by population. Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. We use the 2006 version of WRLURI (Gyourko, Saiz, and Summers 2008). Other controls include population density and the log of self-reported home values (from the 2008-2012 ACS5A). Projected WRLURI is obtained projecting WRLURI on the (log of) population and population density in a given CBSA, the average school years of male and female residents, and fixed effects at the Census Division level. These regressors are taken from the 2000's Decadal Census.

which is 38% of the mean. Appendix Tables A13 to A20 report all coefficients including controls, as well as robustness to different projections of WRLURI.

4.4 Micro-evidence Using the Census of Construction Industries

We now turn to our results using the Census of Construction Industries, using the same sample of firms described in Section 2.3.1.2.⁴³ We measure the local construction sector's revenues per capita (total, and split by housing vs. other construction), housing units built per capita, revenues and payroll per construction employee and per firm, and the average size of construction establishments compared to the average national size of such establishments.

While we only report the coefficient on the WRLURI index, our model is identical to the one described in the previous section. We always control for log CBSA population, log housing value, and density. Observations are again weighted by CBSA population. Variables are winsorized at their 1% and 99% values. Our CBSA count is modestly higher than the external data given the internal records and the use of the projected WRLURI 2006 value. Our results are robust to using different predictive regressions and using the raw index.

Table 8 roughly corresponds to Propositions 2 and 3 in the model. Panel (a) corresponds more closely to Proposition 3. Regressions (1), (2) and (5) report the link between regulation and firm size. A one standard deviation increase in the WRLURI index (0.42) is linked to a 12.8% and 10.3% decline in revenues and payroll per firm. Firms shrink by 6 percentage points relative to the average construction firm as WRLURI increases by one standard deviation. Columns (3) and (4) turn to the link between regulation and firm productivity, measured by revenues and payroll per employee. The coefficients in those specifications imply that a one standard deviation increase in WRLURI corresponds to a 5.4% decrease in revenues per employee and a 3% decrease in payroll per employee. These results confirm that more restrictive project regulation is associated with lower labor productivity in construction firms, whether measured in average or

⁴³To accurately assign CBSA location, these measures are calculated using establishment-level data. Weights are included in the CBSA collapse.

Table 8: Firm Size, Construction Output per Capita, and Regulation

(a) Firm Size					
VARIABLES	(1) Log revenues per constr. firm	(2) Log payroll per constr. firm	(3) Log revenues per constr. employee	(4) Log payroll per constr. employee	(5) Firm size relative to ave. national size
Projected WRLURI	-0.3235*** (0.094)	-0.2565*** (0.0867)	-0.1301*** (0.0367)	-0.0729*** (0.0282)	-0.1426* (0.0857)
Controls	✓	✓	✓	✓	✓
Mean	7.400	5.307	5.785	3.696	0.8652
SD	0.5359	0.2521	0.5243	0.1898	0.2944
Observations	650	650	650	650	650
R-squared	0.412	0.373	0.433	0.491	0.283

(b) Construction Activity Per Capita				
VARIABLES	(1) Log of housing units built per capita	(2) Log of total CCI revenues per capita	(3) Log of housing CCI revenues per capita	(4) Log of non-housing CCI revenues per capita
Projected WRLURI	-0.5299* (0.2721)	-0.3718*** (0.0849)	-0.2577*** (0.0901)	-0.4010*** (0.0972)
Controls	✓	✓	✓	✓
Mean	-7.361	1.115	-0.320	0.761
SD	1.475	0.584	0.656	0.701
Observation	650	650	650	650
R-squared	0.209	0.335	0.393	0.280

Note. The table reports results from a WLS regression at the CBSA level of different firm size (panel a) and construction activity measures (panel b) against WRLURI projected using demographic characteristics. Each CBSA is weighted by population. Robust standard errors in parentheses, *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The analysis uses the 2006 version of WRLURI (Gyourko, Saiz, and Summers 2008) and microdata from 2012 Census of Construction Industries (CCI). CBSA regressions control for log population, log housing values, and population density. Variables are winsorized at their 1% and 99% values. Sample has 650 observations (rounded per Census Bureau disclosure requirements). Projected WRLURI is obtained by projecting WRLURI on the (log of) population and population density in a given CBSA, the average school years of male and female residents, and fixed effects at the Census Division level. These regressors are taken from the 2000's Decadal Census. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2396 (CBDRB-FY24-P2396-R11004).

marginal terms.⁴⁴

Panel (b) shows us the correlation between land use and the size of the construction sector, which corresponds to the predictions of Proposition 2. More regulated places build fewer homes, and total construction-related revenues are lower. A one SD increase in WRLURI is associated with a 20.2% decrease in housing units built per capita and a 14.6% decrease in revenues per capita. Although the index mostly captures residential regulation, this decrease in non-housing revenues is unsurprising both because land-use regulation affects non-housing activity as well, and because lower housing supply hampers population growth, which in turn hampers growth of other activities. Non-housing-related revenues may decline by more than housing-related revenues, perhaps because the amount of housing built per capita is less flexible than the amounts of other forms of construction per capita.

5 Conclusion

In this paper, we have formally presented the hypothesis that project-level regulation, as opposed to regulation of entry, reduces firm size and the incentive to invest in technological innovation. We presented a series of facts that are compatible with that hypothesis, including documenting the small size of construction firms, especially in more regulated areas, and the lower productivity of smaller firms. We also developed a back-of-the-envelope calculation suggesting that firm size alone could explain a significant fraction of the low productivity seen in American residential construction. We believe, however, that more work is required to build confidence that low productivity in construction reflects small firm sizes which in turn reflects project-level regulation.

There are two links in this argument, and the case for either link needs to be made stronger with more data. We found a broad cross-sectional relationship between establishment size and regulation. This empirical connection could be strengthened with a panel analysis of regulatory changes. We also see value in work that documents more

⁴⁴In Section 3 we assumed for simplicity that the cost of inputs is independent of industry conditions. However, generalizing to a finite elasticity of labor supply at the city-industry level, a decline in construction wages would follow from the contraction in construction employment that our theory predicts as a consequence of tighter project regulation.

compellingly the link between regulation and small project size, and that identifies the link between project size and probability of approval. Finally, future work can further disentangle the complex relationship between firm size and productivity. One promising avenue is focusing on inputs like investment in technology and innovation. There may also be structural approaches that could be applied to this setting.

Forty years ago, Mancur Olson's (1984) *The Rise and Decline of Nations* described a process through which insiders enact rules that protect themselves from change. Those rules in turn stymie innovation and maintain the status quo. If the hypothesis described in this paper is correct, then it is a variant of the Olson view. Project-level regulations have been put in place that reduce innovation, not by barring it, but by limiting project and firm size. The small scale of the firms, and the fact that they could not grow dramatically even if they made a breakthrough, then limits innovative activity.

References

- Arnott, Richard J. and Joseph E. Stiglitz (1979). "Aggregate Land Rents, Expenditure on Public Goods, and Optimal City Size". In: *Quarterly Journal of Economics* 93.4, pp. 471–500.
- Brooks, Leah and Zachary Liscow (2023). "Infrastructure Costs". In: *American Economic Journal: Applied Economics* 15.2, pp. 131–158.
- Building Construction Cost Data, 1998, 56th annual edition.* (1998). Kingston, MA: R.S. Means Company.
- Calabrese, Stephen, Dennis Epple, and Richard Romano (2007). "On the Political Economy of Zoning". In: *Journal of Public Economics* 91.1, pp. 25–49.
- Checkoway, Barry (1983). "Large Builders, Federal Housing Programmes and Postwar Suburbanization". In: *Readings in Urban Analysis, Perspectives on Urban Form and Structure*. Ed. by Robert W. Lake. Chap. 1.
- Consumer Guide, Auto Editors of (2001a). *Cars of the Fabulous '50s: A Decade of High Style and Good Times*. Publications International.
- (2001b). *Cars of the Sizzling '60s: a Decade of Great Rides and Good Vibrations*. Publications International.

- Consumer Guide, Auto Editors of (2004). *Cars of the Classic '30s: A Decade of Elegant Design*. Publications International.
- D'Amico, Leonardo (2022). "Place Based Policies with Local Voting: Lessons From the EU Cohesion Policy". In: *SSRN Electronic Journal*.
- Djankov, Simeon, Rafael La Porta, Florencio Lopez de Silanes, and Andrei Shleifer (2002). "The Regulation of Entry". In: *Quarterly Journal of Economics* 117.1, pp. 1–37.
- Ellickson, Robert C. (2022). *America's Frozen Neighborhoods, The Abuse of Zoning*. New Haven and London: Yale University Press.
- Fernandez, Raquel and Richard Rogerson (1996). "Income Distribution, Communities, and the Quality of Public Education". In: *Quarterly Journal of Economics* 111.1, pp. 135–164.
- Flammang, James M. and the Auto Editors of Consumer Guide (2000). *Cars of the Sensational '70s: A Decade of Changing Tastes and New Directions*. Publications International.
- Foerster, Andrew T., Andreas Hornstein, Pierre-Daniel G. Sarte, and Mark W. Watson (2022). "Aggregate Implications of Changing Sectoral Trends". In: *Journal of Political Economy* 130.12, pp. 3286–3333.
- Ganong, Peter and Daniel Shoag (2017). "Why Has Regional Income Convergence in the US Declined?" In: *Journal of Urban Economics* 102, pp. 76–90.
- Glaeser, Edward and Joseph Gyourko (2018). "The Economic Implications of Housing Supply". In: *Journal of Economic Perspectives* 32.1, pp. 3–30.
- Glaeser, Edward and Andrei Shleifer (2005). "The Curley Effect: The Economics of Shaping the Electorate". In: *Journal of Law, Economics, and Organization* 21.1, pp. 1–19.
- Glaeser, Edward L. and Giacomo A. M. Ponzetto (2018). "The Political Economy of Transportation Investment". In: *Economics of Transportation* 13, pp. 4–26.
- Goolsbee, Austan D. and Chad Syverson (2023). *The Strange and Awful Path of Productivity in the U.S. Construction Sector*. Working Paper 2023-04. University of Chicago, Becker Friedman Institute for Economics.
- Gyourko, Joe and Jacob Krimmel (2021). "The Impact of Local Residential Land Use Restrictions on Land Values across and within Single Family Housing Markets". In: *Journal of Urban Economics* 126.
- Gyourko, Joseph, Jontahan Hartley, and Jacob Krimmel (2019). *The Local Residential Land Use Regulatory Environment Across U.S. Housing Markets: Evidence from a New Wharton Index*. Working Paper 26573. National Bureau of Economic Research.

- Gyourko, Joseph and Raven Molloy (2015). "Regulation and Housing Supply". In: *Handbook of Regional and Urban Economics*. Ed. by J. Vernon Henderson Gilles Duranton and William C. Strange. Vol. 5. Elsevier. Chap. 19, pp. 1289–1337.
- Gyourko, Joseph and Albert Saiz (2006). "Construction Costs and the Supply of Housing Structure". In: *Journal of Regional Science* 46.4, pp. 661–680.
- Gyourko, Joseph, Albert Saiz, and Anita Summers (2008). "A New Measure of the Local Regulatory Environment for Housing Markets: The Wharton Residential Land Use Regulatory Index". In: *Urban Studies* 45.3, pp. 693–729.
- Hamilton, Bruce W. (1975). "Zoning and Property Taxation in a System of Local Governments". In: *Urban Studies* 12.2, pp. 205–211.
- Jackson, Kristoffer (2016). "Do Land Use Regulations Stifle Residential Development? Evidence from California Cities". In: *Journal of Urban Economics* 91, pp. 45–56.
- Krimmel, Jacob (2021). *Reclaiming Local Control: School Finance Reforms and Housing Supply Restrictions*. Working Paper. Federal Reserve Board.
- Langston, Craig (2014). "Construction Efficiency: A Tale of Two Developed Countries". In: *Engineering, Construction and Architectural Management* 21.3, pp. 320–335.
- Maisel, Sherman J. (1953). *Housebuilding in Transition*. Berkeley, CA: University of California Press.
- Olson, Mancur (1984). *The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities*. Reprint edition. New Haven, Conn.: Yale Univ Press.
- Rybczynski, Witold (2017). "The Pioneering "Levittowner"". In: *Wharton Real Estate Review*, pp. 73–80.
- Shiller, Robert J. (2015). *Irrational Exuberance*. Princeton University Press.
- Square Foot Costs with RSMeans Data. 2021, 42nd annual edition (and other years)*. (2021). Gordian.
- Stigler, George J. (1971). "The Theory of Economic Regulation". In: *Bell Journal of Economics and Management Science* 2.1, pp. 3–21.
- Turner & Townsend (2022). *International Construction Market Survey*.
- U.S. Bureau of Labor Statistics (1952). *Construction During Five Decades: Historical Statistics, 1907-1952*. Bulletin of the United States Bureau of Labor Statistics.