Regulation and Frontier Housing Supply*

Dan Ben-Moshe† and David Genesove‡

September 15, 2024

Abstract

Regulation is a major driver of housing supply, yet often difficult to observe directly. This paper estimates *frontier cost*, the non-land cost of producing housing absent regulation, and *regulatory tax*, which quantifies regulation in money terms. Working within an urban environment of multi-floor, multi-family housing and using only apartment prices and building heights, we show that the frontier is identified from the support of supply and demand shocks without recourse to instrumental variables. In an application to new Israeli residential construction, and accounting for random housing quality, the estimated mean regulatory tax is 48% of housing prices, with significant variation across locations. Higher regulation is associated with proximity to city center, higher density, and higher prices. We construct a lower bound for the regulatory tax that allows quality to differ systematically over location and time, by assuming (weak) complementarity between quality and demand. The bound is most useful after prices have increased, so that at the end of our sample period, with prices at their highest, we bound the regulatory tax between 40% (using a 2km radius) and 53%.

Keywords: Housing, regulation, stochastic frontier, real estate

*We thank Joe Tracy, Anthony Murphy and Peleg Samuels for helpful comments, and Yotam Peterfreund for excellent research assistance. We thank seminar participants at the Conference on Low-Income Housing Supply and Housing Affordability 2022, UEA 2020, ESAM 2021, AFES 2021, ESEM 2024, Aarhus University, Bar Ilan University, Hebrew University of Jerusalem (Mount Scopus and Rehovot), Ben Gurion University, Tilburg University, Western Galilee College.

†Department of Economics. Ben-Gurion University of the Negev. Email: dbmster@gmail.com

‡Department of Economics. The Hebrew University of Jerusalem. Email: david.genesove@mail.huji.ac.il

1 Introduction

Housing economics attributes a major role to regulation in determining housing prices and residential development (e.g., Glaeser and Ward, 2009; Gyourko and Saiz, 2006; Molloy, 2020). However, the diverse forms of regulation and its often inconsistent enforcement can make direct observation and quantification difficult (e.g., Cheung et al., 2009; Gyourko and Molloy, 2015). Our solution is to estimate first the frontier cost, defined as the non-land cost of producing housing in the absence of regulation, and then the regulatory tax, defined as the money-equivalent extent of regulation. Working with predominantly densely populated urban environments characterized by multi-floor, multi-family housing, we use data on apartment prices per square meter and building heights for our analysis.

Assuming homogeneous housing, we show that the lowest observed price identifies frontier average cost (AC) below minimum efficient scale (MES) and frontier marginal cost (MC) above MES. Our approach replaces standard identification assumptions of exogenous variation with an assumption on the support of demand and supply shocks, without concern for simultaneity. Supply shocks are taken as the differential costs induced by the regulatory environment. For our purposes, regulations principally reflect restrictions on building height. We then account for random housing quality differences using stochastic frontier analysis (SFA) and for systematic differences over location and time with a bounds analysis that relies on (weak) complementarity between quality and demand.

Figure 1 provides intuition for identification of frontier costs for homogeneous housing. Each plotted point represents an observed equilibrium price and height at the intersection of a supply curve that is shifted up by regulatory constraints and a demand curve. The red curve, tracing the locus of equilibria in unregulated markets as demand increases, is frontier marginal cost above MES (i.e., the firm's inverse supply in the absence of regulation). The blue curve, tracing out the locus of equilibria with break-even demand as regulation is relaxed, is frontier average cost below MES. For illustrative purposes these curves are drawn as continuous. As the figure suggests, identification of frontier cost, by minimum price at each height, depends on the support of demand and supply shocks, requiring sufficient variation of demand in unregulated markets in the region with diseconomies of scale (i.e., above MES) and sufficient variation in both demand

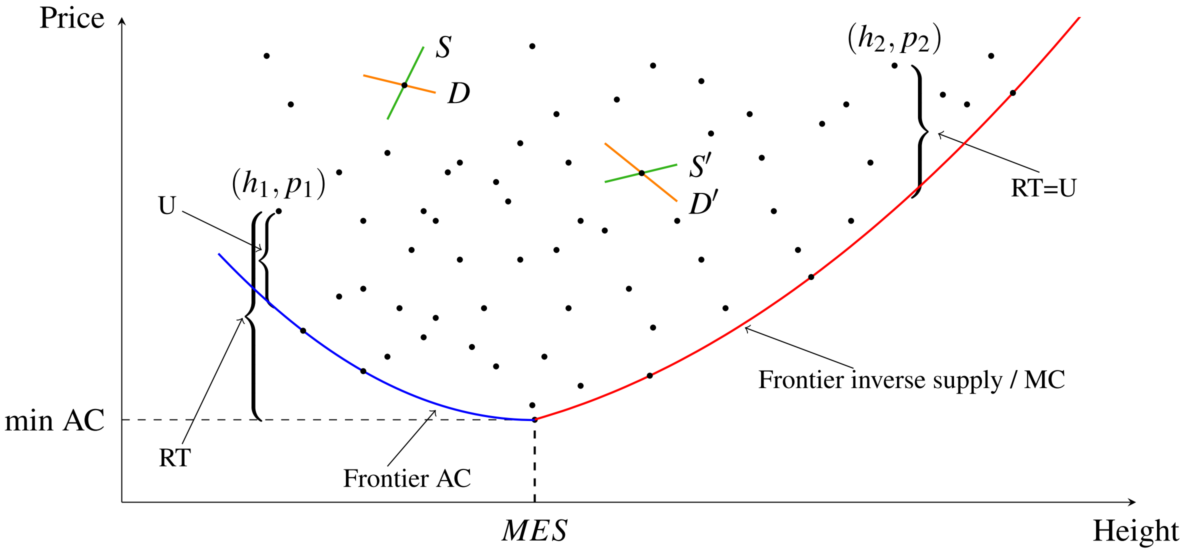

Figure 1: Each point represents an equilibrium price and height. At heights with decreasing economies of scale, the red curve represents the firm's frontier inverse supply. At heights with increasing economies of scale, the blue curve represents the firm's frontier average cost. The regulatory tax is RT. The deviation from the frontier is U.

and regulation in the region with economies of scale (i.e., below MES).

The regulatory tax quantifies the impact of regulation in money-equivalent form, representing both the shadow cost of actually enforced restrictions on construction, as well as delays and additional expenses required to circumvent these restrictions.[1] In an unregulated environment but with the same demand, this tax would induce firms to choose a given building height (i.e., number of floors). Implicitly assuming diseconomies of scale, Glaeser et al. (2005) define the regulatory tax at a given price and height as the price less frontier marginal cost (see Figure 1). Because of the discreteness of building height, as number of floors, there is a range of prices on the supply frontier at any given height (see Figure 2). To address this issue, we amend the definition of regulatory tax to be the maximum of zero and price minus the frontier cost of building an additional floor.

The regulatory tax definition needs modifying for heights below MES, where no tax in an unregulated environment would induce firms to build. To account for such observed buildings, we conceptualize the relevant land areas as covering multiple plots. Then, when demand at minimum average cost falls short of MES, equilibrium absent regulation will consist of some plots developed to MES and others left undeveloped, with average height over all plots equal to quantity demanded. We thus define regulatory tax in the region with economies of scale equal to price less the frontier minimum average cost (see Figures 1 and 3).

[1] The precise definition of regulatory tax is provided in Section 2.3.

In the ideal scenario of Figure 1 the frontier is identified by the minimum observed price at each building height, while regulation is identified based on the deviation between a given price and the frontier. However, this identification is complicated by unobserved quality, such as additional appliances, flooring quality, underground parking, quality, or stage, of construction or exterior aesthetic enhancements. We consider both random quality differences and quality that differs systematically over location and time.

We treat random quality differences as part of measurement error. These errors, which obscure the frontier, are addressed using SFA methods (e.g., Greene, 2008; Kumbhakar et al., 2022). In contrast to standard SFA, which typically relies on the skewness of deviations from the frontier and symmetry of errors for identification, our main approach exploits the spatial structure of the data. Assuming that constant-quality price varies smoothly over space allows us to posit equal regulatory taxes for nearby buildings of the same height. This enables us to separate regulation from measurement error by leveraging within-building, between-building, and between-bloc variation in prices, where a bloc is a land registry geographical division averaging about 150 apartments in multi-housing-unit buildings. Additionally, unlike SFA, we allow the frontier and distributional parameters to vary arbitrarily with height. This approach is feasible because the data contain hundreds to tens of thousands of observations at each height, enabling us to perform estimation separately at each height.

Quality may, alternatively, be systematically related to locational amenities if consumers prefer higher quality housing in areas with more desirable amenities. In this case, the frontier represents the non-land costs of producing housing with minimal, rather than average, quality. Yet, without further structure, it is not possible to distinguish the effects of regulation from those of quality above the minimum.[2] To address this issue, we assume additionally that within a limited distance structural quality and amenities are (weak) complements—higher-quality housing is built in locations with at least as desirable amenities. This allows us to bound the regulatory tax by comparing frontier costs and prices for nearby buildings. Similarly, preferences for quality may change over time. Weak complementary of quality with temporal demand shocks allows us to use construction-year effects in hedonic price regressions for existing homes to identify quality

[2]The analogous difficulty for the SFA literature would be distinguishing product quality from firm inefficiency. This issue has not received much attention in the SFA literature, although it is an important issue in the productivity literature (and more generally) since Klette and Griliches (1996).

differences across time.

Our empirical application uses newly constructed residential buildings in Israel from 1998 to 2017 relying on variation in prices across both space and time. This market is particularly suitable for our study due to the extensive variation in enforced regulation. Even neighboring buildings may face different levels of enforcement depending on builders' success in securing permits, which they must obtain from at least two different levels of local planning committees, each with considerable discretion (see Czamanski and Roth, 2011; Rubin and Felsenstein, 2019).

Our study yields six main findings. First, the estimated frontier decreases at low heights, indicating economies of scale, while a mean regression increases steeply. Second, estimates of the frontier elasticity of substitution of land for capital (defined as all non-land inputs in construction) is less than 0.5 at heights just above MES and at high heights but exceeds unity at medium heights, where marginal costs are flat. This suggests that building upwards is easy at medium heights but hard at low and high heights. Third, the mean regulatory tax estimates are about 48% of market price, which aligns with the findings of Glaeser et al. (2005) for residential buildings in Manhattan, and Cheshire and Hilber (2008) for UK office buildings, both of which rely on commercially available cost estimates. This suggests that suppliers would build taller buildings in unregulated markets, despite the difficulty in building upwards. Fourth, the estimated regulatory tax as a percentage of price has a standard deviation of about 16%, indicating significant variation. Fifth, areas that are higher priced, denser, and closer to city centers have higher regulatory tax. Finally, allowing quality to systematically differ over location and time, we bound mean regulatory tax. In 2017, when prices were at their peak in our sample—so that the lower bound is especially informative—we estimate a lower bound of 40% (using a 2km radius) and an upper bound of 53%.

Estimation of the (mean) housing production function has enjoyed a recent renaissance (e.g., Albouy and Ehrlich, 2018; Brueckner et al., 2017; Cai et al., 2017; Combes et al., 2021; Epple et al., 2010). However, most of this research deals with single family housing, with only a few papers addressing building height. Ahlfeldt and McMillen (2018) measure the land price elasticity of height, but disclaim any variation in regulatory conditions in their coverage area. Henderson et al. (2017) focus on uncertain property rights rather than regulation, and take a structural approach. Tan et al. (2020) infer the bindingness of observed height restrictions from their effect on the land-price to housing-price relationship, an approach, unlike ours, requiring data on land prices.

A significant challenge in using housing data, as in many other economic applications, is the difficulty of directly measuring costs and regulations, which are often not fully observable. Hence, quantitative assessment of housing regulation typically infers regulatory effects from the partial correlation of housing market outcomes with observed measures of regulatory strictures, such as the Wharton Index of Gyourko et al. (2008) or the new Wharton index of Gyourko et al. (2021). Early studies were concerned with the capitalization of regulation into mean housing prices (e.g., Katz and Rosen, 1987; Pollakowski and Wachter, 1990). More recent work has focused on the effect of regulation on housing market response to demand shocks by considering housing price variability (Paciorek, 2013), market supply elasticity (Saiz, 2010), or income pass-through to prices (Hilber and Vermeulen, 2016).

In contrast, Glaeser et al. (2005) and Cheshire and Hilber (2008) directly measure the regulatory tax by comparing housing prices to an external assessment of construction costs. Our analysis shares the goal of measuring the regulatory tax, but avoids relying on external, commercial cost assessments, which are often not available. Even when they are available, our approach of inferring costs from prices and quantities aligns with the Industrial Organization literature's skepticism regarding the use of industry cost estimates for economic analysis. Another approach computes the regulatory tax by the excess of the intensive value of land, inferred from housing prices, over the extensive value of land, observed from land transactions (Gyourko and Krimmel, 2021). However, this method is likely appropriate only for single-family homes.

Measuring housing costs and regulation is important for several policy issues. Building upwards can mitigate urban sprawl by increasing density, offering an alternative to outward expansion (e.g., Brueckner and Helsley, 2011; Fu and Somerville, 2001; Nechyba and Walsh, 2004). Variation in housing regulation across locations may reduce productivity by causing spatial mismatches between labor and capital (Hsieh and Moretti, 2019), although we are agnostic about the welfare consequences of the regulation we measure here. Additionally, housing deregulation is an important policy tool for checking growing inequality of wealth, particularly if due to increasing land scarcity (e.g., Rognlie, 2016). Understanding the effect of regulation on housing is crucial for designing effective policies to address these and other related policy issues.

The remainder of the paper is organized as follows. Section 2 focuses on identification. Section 3 describes the estimators. Section 4 reviews the data. Section 5 presents the empirical results.

2 Identification

This section presents a demand and supply framework for identifying frontier costs when observing only equilibrium prices and quantities - which, as we will discuss, are essentially heights in our context. Section 2.1 analyzes frontier supply and Section 2.2 frontier average costs at low heights with economies of scale. Section 2.3 defines the regulatory tax and Section 2.4 bounds the regulatory tax when allowing for quality to be demand-dependent, over location or time. Section 2.5 adjusts prices when households' willingness to pay depend on the apartment's floor and the building's height and Section 2.6 explains how to use SFA techniques to incorporate unobserved random quality as measurement error. Section 2.7 addresses potential critiques by discussing the limitations and assumptions involved in our identification strategy.

2.1 Frontier Supply (Marginal Cost at Heights Above MES)

This section provides conditions under which frontier supply is identified by the joint distribution of equilibrium prices and quantities, in an idealized environment of perfectly competitive markets for a single good produced by equally efficient firms. Since competitive firms supply only at quantities where there are no economies of scale, this discussion concerns such quantities only. The identifying conditions place no restrictions on the joint distribution of the unobserved and observed variables, other than their support. Simultaneity will not be a concern.

Consider multi-floor housing built on parcels of one unit of land each. For simplicity, at most one building can be built on each parcel, with the building covering the entire parcel. Buildings consist of homogeneous housing. Define one unit of housing as a 1-floor building on one unit of land. Then the quantity of housing in one building is its number of floors. We observe the price per unit of housing, $p \in (0,\infty)$, and the number of floors, which we refer to as height, $h \in \{1,2,\ldots\}$, for each newly constructed building.

Consider parcel-level supply (analogous to firm supply in basic theory), which includes any regulatory restrictions. Since the quantity of housing is the number of floors, a supply curve can take nonnegative integer values only, and so is fully characterized by the jump discontinuities at $p_1, p_2, \ldots$, where p_h is the minimum price at which profit maximizing suppliers would build h units of housing under the given regulations. In other words, p_h is the marginal cost of the h-th floor. A strict maximum height restriction at h floors would take the form of $p_{h+j} = \infty$ for $j > 0$.

More generally, builders may be able to overcome restrictions by sufficient expenditure on legal efforts or lobbying; these additional costs explain the vertical gap between non-frontier (regulated) and frontier (unregulated) supply.[3] We derive conditions under which the frontier marginal cost of building the h-th floor p_h^f is identified by the minimum price at height h.

Next consider, for conceptual purposes only, an *area* with a collection of unit land parcels. Consumers consider housing services provided on any parcel as identical to those provided on any other parcel in a given area.[4] Inverse demand for housing in the area, which is assumed continuous, is therefore a function of the total housing consumed in the area. Define parcel-level demand as market demand for the area divided by the total number of parcels in the area.

Figure 2 shows parcel-level supply and demand curves. The red curve is the inverse frontier supply curve, the object of interest, while the green curve is some inverse non-frontier supply curve. The blue curve is inverse demand for a low demand shock, while the orange curve is inverse demand for a high demand shock (violet will be considered later).

Equilibria are at the intersections of inverse demand and inverse supply curves. The figure shows the unique equilibrium for each combination of demand - low (D_L) or high (D_H) - and supply - unregulated (S_U) or regulated (S_R). The equilibrium with no regulation and low demand is E_1. At this equilibrium, price lies between the frontier marginal cost of constructing a 3-floor building, p_3^f, and a 4-floor building, p_4^f, and so only 3-floor buildings are built.

The equilibrium with no regulation and high demand is E_2. At this equilibrium, price equals p_4^f with suppliers indifferent between building 3-floor and 4-floor buildings and the market clears at the fraction of 3-floor buildings built.

The two remaining points show equilibria under supply with regulation. The equilibrium with regulation and high demand is E_3. Absent regulation, and at the associated equilibrium price p_3, suppliers would build 4-floor buildings. Regulation costs lead suppliers to build only 3-floor buildings. Similarly, at E_4, with low demand, 2-floor buildings are built, although suppliers prefer to build an additional floor.

[3] Cheshire and Dericks (2020) argue that, in London, builders may overcome restrictions by employing 'trophy' architects. The regulatory tax would then include the distortion resulting from the excess outlay on the architect compared to the added value to the buyer. Payments or favors to officials are the more likely tool for overcoming restrictions in our application.

[4] In using area as a conceptual device, one need not imagine a contiguous expanse. See Piazzesi et al. (2020) for evidence of buyers searching over noncontiguous areas.

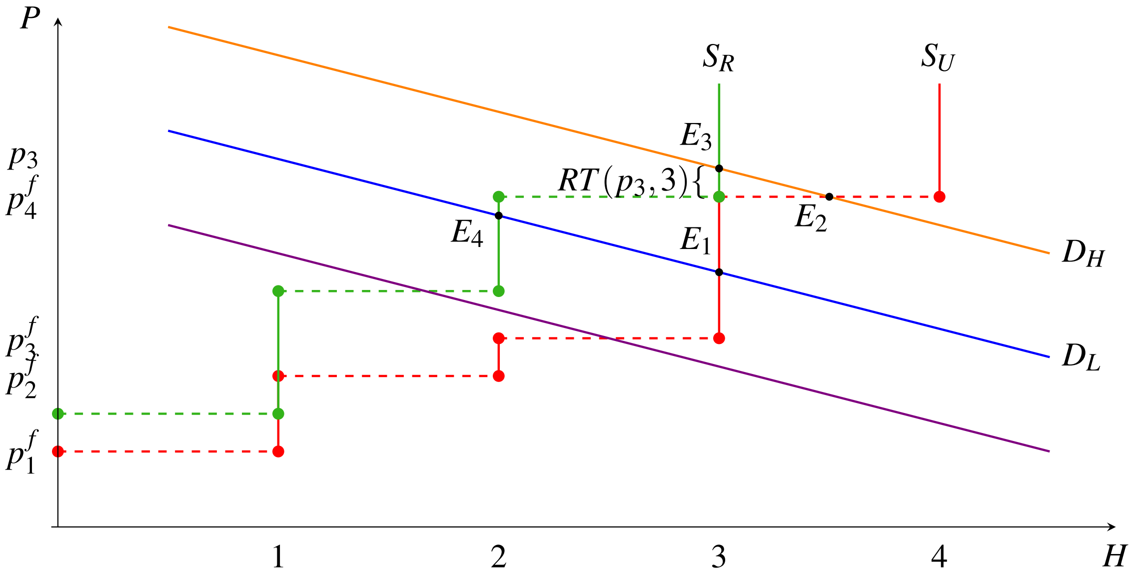

Figure 2: Parcel-level inverse supply and demand curves.

Our empirical analysis conditions on building height. Consider 3-floor buildings, which are built at E_1 (where suppliers want, and are permitted, to build 3-floor buildings), E_2 (where suppliers are indifferent between three and four floors, and some build three floors), and E_3 (where suppliers want to build four floors but permitted only three). The lowest price among these three equilibria is at E_1, which is greater than the minimal price p_3^f required to induce unregulated suppliers to build 3-floor buildings.

Hence, if the pictured high and low demand curves were the extent of demand variation then p_3^f would not be identified. Identification requires a positive probability of frontier supply and a demand curve cutting it at p_3^f. The violet demand curve in Figure 2 is just one such curve that would allow identification. Note that E_2, where the high demand curve intersects the unregulated supply curve, identifies the minimal price to build 4-floor buildings p_4^f. Identification of the frontier supply curve as a whole, then, requires sufficient variation in demand in unregulated markets.

Formally, inverse demand $P^d(h,\varepsilon)$, with random demand shock ε, is assumed continuous in height $h \geq 0$. Inverse supply is defined by the correspondence $P^s(h,W) = \{p \mid p_h^W \leq p \leq p_{h+1}^W\}$, with random supply shock W and $h \in \mathbb{N}$. The frontier inverse supply is defined by $P^s(h,f) = \{p \mid p_h^f \leq p \leq p_{h+1}^f\}$, with $p_h^f = \min_{w \in \text{Support}(W)} p_h^w$, for each h. An equilibrium (P,h,α) is a price $P \geq 0$, height $h \in \mathbb{N}$, and fraction $0 \leq \alpha < 1$, such that the market clears: $P = P^d(\alpha(h-1)+(1-\alpha)h, e) \in$

$P^s(h,w)$, for some $(e,w) \in \text{Support}(\varepsilon, W)$. Now define

$P(h) = \{P : (P,h,0)$ or $(P,h+1,\alpha), 0 \leq \alpha < 1$, is an equilibrium, for some $(e,w) \in \text{Support}(\varepsilon, W)\}$.

If there exists e with $(e,f) \in \text{Support}(\varepsilon, W)$ and $0 \leq \alpha < 1$ such that $P^d(\alpha(h-1)+(1-\alpha)h, e) = p_h^f$, then p_h^f is identified by $\min\{P(h)\}$. In other words, we are assuming sufficient realizations of frontier supply, and demand intersecting it at the frontier price. Note that issues of simultaneity do not arise here. This identification result suggests the sample minimum price at height h as a natural estimator for p_h^f.

2.1.1 Spatial Dependence

The above discussion considers each building in isolation. Real estate markets, however, are characterized by spatial dependencies. The inverse demand curves above can be reinterpreted as residual demands for each building, given prices of other new and existing buildings. However, the probability statements need further consideration under spatial dependencies.

Spatial dependence here can be either local or global. Locally, price may be sensitive to local density, directly through density's effect on utility and indirectly through price response to local supply. To account for this, consider indexing both demand and supply shocks by location, and assume weak dependence so that price dependence diminishes, and approaches independence, as the distance between locations increases (see, e.g., Conley, 2010). Similarly to the nonspatial framework, for each height we assume a positive probability of a location for which the collective shocks within a neighborhood, outside of which prices are essentially independent, lead to realizations of the frontier supply, and demand intersecting it at the frontier price for that height.

Globally, spatial dependence may arise from cross-location arbitrage in an at least partially closed market, where prices arise from aggregate supply and local demand. Then, the price at any location will be determined by an aggregate of market-wide shocks—regulation and demand across all locations, such as the boundary rent curve of Fujita (1989)—and location-specific shocks—regulation and quality at the location. Identification will then be ensured if there is a positive probability of no regulation at any given value of locational quality. The fully encompassing greenbelt example discussed in Section 2.7.1 would be a case in which identification would fail for some heights, but not others.

2.2 Frontier Average Cost at Heights Below MES

In perfectly competitive unregulated markets, firms never construct buildings at heights where there are economies of scale as building at heights at or above MES would always be more profitable. However, under regulation, suppliers might build at heights below the frontier's MES. Minimum price at such heights could not correspond to frontier supply. Rather, the minimum price identifies frontier average cost, under conditions shown below.

Figure 3 shows the textbook example of a U-shaped frontier average cost curve, along with its associated marginal cost curve. For simplicity, we present continuous curves. The frontier supply function maps prices below minimum AC to height equal zero (i.e., the land is left undeveloped) and maps prices above the minimum AC to the inverse MC (the red curve in Figure 3). At price equal to minimum AC, suppliers are indifferent between leaving the land undeveloped and building at MES. Thus an equilibrium where the parcel-level housing quantity demanded at minimum AC falls short of MES involves price equal to minimum AC, with some parcels left undeveloped and the remainder developed to height MES, with their shares such that the market clears. An equilibrium where the quantity demanded at minimum AC exceeds MES entails an above minimum AC price and construction on every parcel at a common height above MES.

Inferring frontier costs at heights below MES thus requires the realization of non-frontier supply. Equilibrium E_h must be generated by some such supply curve intersecting with a demand curve (neither is shown). However, lower prices at the same height h could also be observed, given appropriate demand and regulated supply shocks. The lowest possible observable price is $p'_h = AC(h)$, which would be generated by the joint realization of a demand and non-frontier supply that intersect at E'_h.[5] No lower price is possible at h; otherwise, firms would suffer losses.

Hence, whereas minimum price, conditional on height, converges to MC at heights for which AC is increasing, it converges to AC where AC is decreasing. Minimum price thus identifies the maximum of frontier AC and MC, denoted as $G(h) = \max\{AC(h), MC(h)\}$, which in Figure 3 is the blue curve $\min\{P(h)\} = AC(h)$ and the red curve $\min\{P(h)\} = MC(h)$. Whereas identification at heights of increasing AC requires variation in demand in unregulated markets, identification at heights of decreasing AC requires variation in both demand and regulation.

[5]Recall that firms are perfectly competitive and that the demand that passes through E_h or E'_h are market demands scaled down to the parcel, and so firm, level.

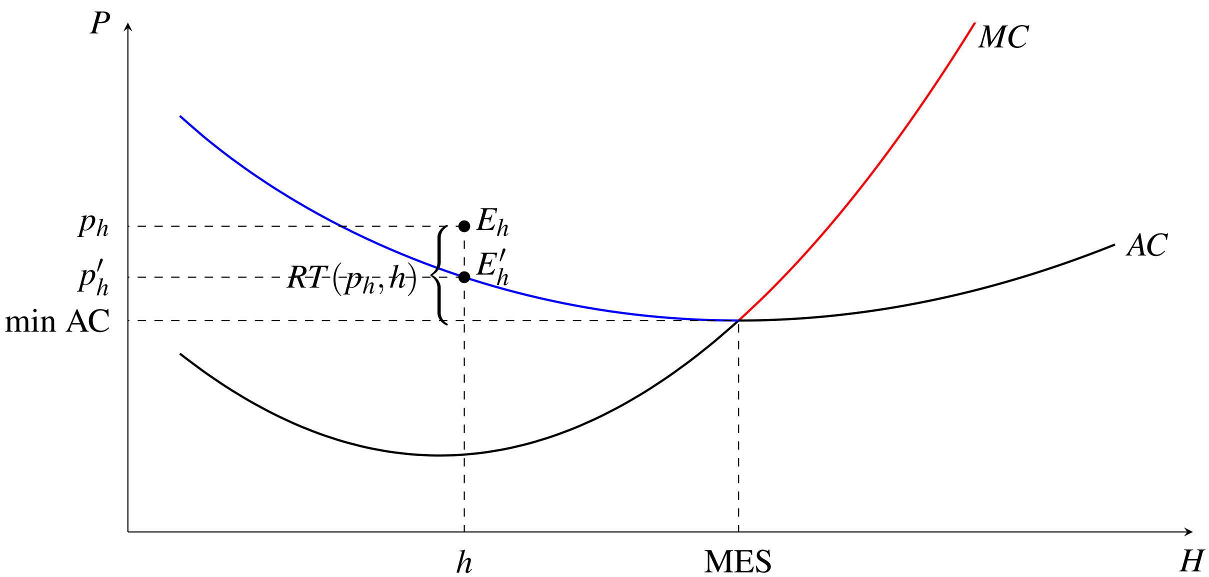

Figure 3: Frontier AC and MC curves.

Assuming a U-shaped frontier average cost curve is an important simplification. In principle, the cost structure might differ. First, average costs might be declining for some region at high heights. However, the maximum extent of the rate of decline decreases with height, since total costs are weakly increasing $(AC(h)-AC(h-1))/AC(h-1) \geq -1/h$. Second, there may be regions where marginal frontier costs exceed average costs yet are decreasing, where firms would ordinarily not operate, but might under regulation. This would be especially difficult to handle as the minimum observable price would actually exceed frontier marginal costs. Furthermore, incorporating such irregular cost structures would involve multiple local turning points, as opposed to the single one at MES that we have here. For these reasons, we impose the condition of a U-shaped average cost curve.

2.3 Regulatory Tax

Define the regulatory tax (RT) as

$$RT(P,h) = \begin{cases} P - AC(MES), & h < MES, \\ \max\{0, P - MC(h+1)\}, & h \geq MES, \end{cases} \tag{1}$$

where $MES = \operatorname{argmin}_{h \in \mathbb{N}}\{AC(h)\}$. This is the minimum tax in an unregulated environment for which one would observe a height h with price P. Define the regulatory tax rate (RTR) as $RT(P,h)/P$.

Below MES, the only possible equilibrium price in an unregulated market is minimum average cost $AC(MES)$. In such an equilibrium, parcel-level height demanded is h and firms are indifferent between not building at all and building to MES. Some parcels are left undeveloped and others built to MES, with the share such that demand equals supply. Hence, at E_h in Figure 3, the regulatory tax is $RT(p_h, h) = p_h - AC(MES)$, which would raise average costs so that h-floor buildings would be built absent other regulation.

Above MES, for an unregulated competitive firm to choose height h, we must have $MC(h) \leq p \leq MC(h+1)$. Thus when price is below the marginal cost of adding another floor, the regulatory tax is zero and when price exceeds the marginal cost of adding another floor, the regulatory tax is equal to the difference. Hence, at E_3 in Figure 2, the regulatory tax is $RT(p_3, 3) = \max\{0, p_3 - MC(4)\} = p_3 - p_4^f$, which would raise marginal costs so that 3-floor buildings would be built absent other regulation.

2.4 Bounds for Systematic Quality

Unregulated suppliers will build better when building higher if households with greater locational amenity preferences also prefer higher quality housing, or if households prefer higher quality when purchasing in better locations, or in periods with greater demand for housing. This will result in quality differing systematically over location and time. Consequently, the frontier will represent the non-land costs of producing minimal-quality, rather than average-quality, housing, but the difference between the price and frontier will be the sum of regulatory effects and the excess of quality above the minimum-quality frontier, requiring some method to separate the two. In this section, we bound the regulatory tax.

To begin, assume total costs are $C(h) + zh$, where $C(h)$ is the frontier-quality cost of building to height h and zh is the extra cost of building at quality $z \geq 0$; with this specification, additional quality adds the same amount to marginal as to average cost, and profit-maximizing quality is independent of height and thus of regulation. Now, for any building i, its price P_i is the sum of frontier cost $G(h_i) = \max\{MC(h_i), AC(h_i)\}$, marginal cost due to quality z_i, and deviation $U_i \geq 0$,

$$P_i = G(h_i) + z_i + U_i. \tag{2}$$

The deviation captures the regulatory effects, and when cost curves are continuous, and suppliers build above MES, as in Figures 1 and 3, then the deviation is exactly the regulatory tax.

Since $z_i \geq 0$, an upper bound for the regulatory tax is obtained when $z_i = 0$,

$$U_i \leq P_i - G(h_i),$$

$$RT_i := RT(P_i, h_i) \leq \begin{cases} P_i - AC(MES), & h_i < MES, \\ \max\{0, P_i - MC(h_i + 1)\}, & h_i \geq MES, \end{cases}$$

This bound, the same as (1), assigns the entire difference between price and frontier to regulatory restrictions, dismissing any contribution from quality.

Next, for a lower bound on the deviation for focal building i, consider a comparison building j. Taking the difference between equation (2) for buildings i and j, rearranging, and using the nonnegativity of the deviation for building j, $U_j \geq 0$, yields a bound for the focal building's deviation:

$$U_i \geq \underbrace{(P_i - P_j)}_{(i)} - \underbrace{(G(h_i) - G(h_j))}_{(ii)} - \underbrace{(z_i - z_j)}_{(iii)}. \tag{3}$$

Of these three components, we now focus on the quality differential (iii), as it is not observed, and must be inferred through additional structure. To that end, decompose $z_i - z_j = (z_i - z(a_j, t_i)) + (z(a_j, t_i) - z_j)$. Here, $z(a_j, t_i)$ represents the quality that would arise at the comparison building's location but at the focal building's transaction period. The spatial component, $(z_i - z(a_j, t_i))$, represents the quality difference due to different locations, at the focal building's transaction period. The temporal component, $(z(a_j, t_i) - z_j)$, represents the quality difference due to different transaction periods, at the comparison building's location.

To bound the spatial component, write the price of housing with amenities a, transaction time t, and quality z as $P(a, z, t)$. We assume local (weak) complementarity between amenities and quality, i.e., the returns to quality are nondecreasing with amenities: $P_{za} \geq 0$.[6] This still allows for different trade-offs between amenities and quality in different geographic areas; indeed, imposing global complementarity between amenities and quality would contradict a constant quality frontier.

A profit-maximizing, price-taking supplier, unconstrained in choice of quality, will choose quality $z(a, t)$ to satisfy the first order condition

$$P_z(a, z(a, t), t) = 1. \tag{4}$$

[6] We use the standard notation f_x to denote the partial derivative $\partial f / \partial x$.

For the spatial component, fix time t. Totally differentiating the first order condition (4) and price $P(a,z,t)$ implies,[7]

$$dz = \frac{1}{1-(P_{zz}P_a/P_{za})} \times dP \equiv \kappa_S(a,z) \times dP. \tag{5}$$

Weak complementarity $P_{az} \geq 0$, the second order condition $P_{zz} \leq 0$, and $P_a > 0$ (by definition) imply $0 \leq \kappa_S(a,z) \leq 1$. Thus if a building's locational amenity is smooth in location, we can conclude that $z(a_j,t_i) - z_i \approx \kappa_{Si} \times (T_{ij}P_j - P_i)$ for all comparison buildings j sufficiently close to focal building i, and for some $\kappa_{Si} \in [0,1]$, where $T_{ij}P_j$ is defined as building j's price deflated to building i's transaction period using a housing price index. Hence, for each building we choose $\kappa_{Si} \in [0,1]$ to minimize $\kappa_{Si} \times (T_{ij}P_j - P_i)$.

Next, we bound the temporal component. Assume, for newly constructed housing, the standard hedonic price specification $P(a,z,t) = \exp(\gamma(t))P^0(a,z)$, so that $\gamma(t)$ are time fixed effects in the log-linear specification. Importantly, this builds in complementary, as $P_{\gamma(t)z} = P_z \geq 0$. Fix amenity a. Totally differentiating the log first order condition for quality (4) and log price, we obtain[8]

$$dz = \frac{\delta(a,z)}{1+\delta(a,z)} \times dP \equiv \kappa_T \times dP, \tag{6}$$

using the first order condition $P_z = 1$, and where $\delta \equiv -P_z^2/(PP_{zz}) \geq 0$ is an inverse measure of the convexity of P as a function of z (a constant for P isoelastic in z). This allows us to write $z(a_j,t_i) - z_j \approx \kappa_T \times (T_{ij}P_j - P_j)$ for all comparison buildings j sufficiently close in time to focal building i.

In contrast to the coefficients κ_{Si} for the spatial component, κ_T can be estimated. Generalizing our price specification above to accommodate existing homes, and noting that the choice of quality for housing constructed at time t can be written as $z(a,\gamma(t))$, let the log price of housing constructed in period s and sold in period t be $\ln P = \gamma(t) + \ln P^0(a,z(a,\gamma(s)))$. Then a linear approximation of the price around the quality of new construction at an arbitrary time period 0, $z(a,\gamma(0))$, is[9]

$$\ln P \approx \gamma(t) + \ln P^0(a,z(a,\gamma(0))) + \delta(a,z(a,\gamma(0))) \cdot \gamma(s). \tag{7}$$

This motivates estimating δ by the proportionality coefficient in a restricted log price regression

[7] We solve $dP = P_a da + P_z dz$ and $0 = P_{za} da + P_{zz} dz$ for unknown dz and da.

[8] We solve $d\ln P = d\gamma(t) + (P_z/P)dz$ and $0 = d\gamma(t) + (P_{zz}/P_z)dz$ for unknown dz and $d\gamma(t)$.

[9] This follows from $\frac{\partial P^0}{\partial\gamma(s)} = \frac{P_z}{P} \cdot \frac{\partial z}{\partial\gamma(s)} = \frac{P_z}{P} \cdot (-\frac{P_z}{P_{zz}}) \equiv \delta$.

that conditions on the dates of transaction ('period effect') and construction ('cohort effect'), with the cohort effect constrained to be proportional to the period effect, and with parcel fixed effects for $\ln P^0(a, z(a, \gamma(0)))$.[10]

Returning to inequality (3), inserting the approximations for the spatial and temporal components of the quality differentiation, accounting for discrete height and nonnegativity of the focal building's own deviation, and noting that the inequality holds for all local buildings, we choose the largest bound for the set $\Omega_i(d)$ of buildings j within a radius d from building i. The lower bound is now obtained by a minimax,

$$U_i \gtrapprox \min_{\kappa_{Si}\in[0,1]} \max_{j\in\Omega_i(d)} \{[G(h_j) - G(h_i)] - [(P_j - P_i) - \kappa_T(P_j - T_{ij}P_j) - \kappa_{Si}(T_{ij}P_j - P_i)]\}, \tag{8}$$

$$RT_i \gtrapprox \min_{\kappa_{Si}\in[0,1]} \max_{j\in\Omega_i(d)} \max\{0, [G(h_j) - G(h_i+1)] - [(P_j - P_i) - \kappa_T(P_j - T_{ij}P_j) - \kappa_{Si}(T_{ij}P_j - P_i)]\}. \tag{9}$$

Thus, the regulatory tax is bounded from below by the difference between the frontier-quality construction costs of any sufficiently close building j and those of the focal building, minus the difference in their quality-adjusted prices. Choosing the radius d involves a tradeoff: a larger d results in higher lower bounds but reduces the accuracy of the spatial component in the quality approximation. Therefore, we consider how the lower bound changes with respect to d.

2.5 Adjusting Prices for Consumer Preferences of Apartment Floor and Building Height

We account for consumers valuing apartment floor or building height by "efficiency unit" modeling of housing services, with log price

$$\ln(\text{price}) = \ln p + \ln m(f, h), \tag{10}$$

where m is an unknown function representing the premium that all households are assumed willing to pay for an fth-floor apartment in an h-floor building, and p is the price net of this, reflecting the value of the building's location. Hence, per unit of land the quantity of housing in an h-floor building is the sum of the premiums, $q(h) = \sum_{f=1}^{h} m(f, h)$.

Although building height maps one-to-one to the quantity of housing (and in our data they are

[10]In principle, δ can vary across locations. However, allowing δ to vary by city in the empirical analysis does not change our results. That issue, along with depreciation and the relationship of the proportionality restriction to the well known period-cohort-age problem are discussed further in Appendix A.1.

very close, with $0.05 < (q(h) - h)/h < 0.1$), they are not identical. Since the discrete levels of quantity will not be integers, it will usually be convenient to express cost as a function of height. Yet, with price stated per unit quantity, we make this relationship explicit. Let $h(q)$ denote the inverse of $q(h)$.[11] Then $\widetilde{C}(q) = C(h(q))$, where $\widetilde{C}(q)$ is the frontier cost of building quantity q and $C(h)$ the frontier cost of building to height h.

Break-even market price for an h-floor building is

$$AC(h) = \frac{C(h)}{\sum_{f=1}^{h} m(f,h)} = \frac{\widetilde{C}(q(h))}{q(h)}.$$

This is the lowest possible observed adjusted price in a region with economies of scale.

For diseconomies of scale, the lowest possible observed adjusted price at any given height equals the marginal cost savings from building the next lowest feasible quantity,

$$MC(h) = \frac{C(h) - C(h-1)}{\sum_{f=1}^{h} m(f,h) - \sum_{f=1}^{h-1} m(f,h-1)} = \frac{\widetilde{C}(q(h)) - \widetilde{C}(q(h-1))}{q(h) - q(h-1)}.$$

2.6 Measurement Error and Random Quality

Our empirical analysis also allows for random quality differences at both the building and apartment levels.[12] While it is not necessary to classify systematic and random types of quality differences, evident sources of random quality differences include varying stages of construction completion at time of transaction (as noted in Combes et al., 2021) and small capital goods, such as appliances, that are available at the same cost to both suppliers and buyers, and are part of the apartment price. We treat such quality differences as measurement errors and apply techniques from SFA.[13] Literal measurement errors—such as transcription mistakes or misreports of apartment price or floor area—are also considered part of these errors.

Separating the convolution of deviations and errors without additional information can be achieved by restricting their distributions (e.g., Florens et al., 2020; Schwarz and Van Bellegem, 2010). In practice, SFA often identifies deviations by imposing skewness on deviations and symmetry in measurement errors. We prefer not to rely on shape restrictions such as symmetry

[11] This inverse exists as long as $m(f,h) > 0$, for all $1 \leq f \leq h$, which is the case empirically.

[12] For a demand and cost specification that includes both systematic quality z_s and random quality z_r, specify price as $P(a, z_s) + z_r$ and cost as $C_0(h) + h \times (z_s + z_r)$, so that firms will be indifferent over choices of random quality.

[13] Unlike SFA, which assumes unregulated markets with deviations representing firm inefficiency, our approach assumes equally efficient firms, with deviations representing regulation. Allowing for a region of increasing returns to scale, as in Section 2.2, also differs from SFA.

and instead leverage the hierarchical structure of the data (Kotlarski, 1967),[14] assuming constant-quality prices vary smoothly over space. This then implies equal deviations for sufficiently nearby buildings *of the same height*. Variances are now identified using variation in prices within buildings, across buildings, and across blocs. Additionally, estimation conducted separately at each height allows these variances, the frontier, and the distributional parameters to vary arbitrarily with height.

2.7 Further Discussion of Identification

Identification of the frontier only requires observable prices and quantities (i.e., heights), with the distributions of deviations from the frontier allowed to depend on height, obviating the usual need for exogenous variation. Also, no parametric or separable conditions need be imposed on the structure of demand or (regulated or unregulated) supply. Other characteristics of the environment become critical, though.

2.7.1 Support

We have assumed a positive probability of observing unregulated markets at heights for which there are diseconomies of scale, and regulated markets at heights for which there are economies scale, in place of the standard exogeneity assumptions for identification. The frontier is not identified if these markets are not realized. Of course, there can be no hope of uncovering costs in the absence of regulation that is always imposed, such as nationwide safety regulations. Thus “unregulated" should really be interpreted as “minimally regulated", and it is the “minimally regulated" frontier that is our estimation objective. The problem arises rather when minimal regulation is realized at certain heights, but not at others. However, that scenario might be detectable if one ends up estimating a nonsensical cost function. For example, consider transaction prices from a period of stable prices in a city well characterized by the demand conditions of the monocentric city model, where willingness to pay decrease from the city center. A greenbelt, where construction is forbidden, that surrounds the city would leave no way to identify marginal costs for heights that would have otherwise been built there. In this case, identification failure for this part of the frontier supply would be apparent from the gap in the distribution of prices, unconditional on height, whereas multiple time periods with changing overall prices will introduce more unregulated height away from the greenbelt, thus restoring identification.

[14]Our estimates using the hierarchical structure indicate relatively symmetric deviations.

2.7.2 Perfect Competition and Equally Efficient Firms

We have assumed perfect competition and equally efficient firms, with the same costs over firms, space, and time. To account for cost changes over time, we adjust prices using the Israeli Central Bureau of Statistics' residential construction input-prices index.[15] Non-land cost differences over space are small according to industry participants.[16] This is corroborated by similar frontier estimates on samples that remove areas known to face greater technical challenges (see Figure 8).

A perfectly competitive residential construction industry with identical cost firms is the standard assumption in the empirical housing production function literature (e.g., Albouy and Ehrlich, 2018; Brueckner et al., 2017; Cai et al., 2017; Combes et al., 2021; Epple et al., 2010). It also accords with conditions in our application: the Israeli construction industry is structurally competitive, with a ten-firm concentration ratio of 0.15 only (Ministry of Finance, Chief Economist Branch, 2017). Local concentration is also likely to be low: the larger firms operate throughout the country, and the locality-level Herfindahl concentration index of auctioned-off building rights for housing units on government-owned land is 0.025, equivalent to forty equally sized firms (see Appendix C.3).[17]

In addition, builders face competition from existing homes, both renovated and unrenovated. The annual construction flow is two percent nationally, so, as elsewhere, the stock of existing homes is much larger than the quantity of newly constructed housing. Also, we are working within a generally urban environment characterized by high density and apartment buildings. The typical new building in our sample is located in an area with an existing population of 5,260 people, or about 1,600 apartments, within a 500 meter radius.

Although the structural conditions in the market suggest that markups should not be a major problem here, it is instructive to consider what implications markups have for our regulatory tax estimates. Importantly, a markup per se need not bias them. A constant absolute markup among all buildings would simply shift up the measured frontier, which would not absorb that markup.

[15] Estimates without adjusting for construction cost changes are similar (see Figure 15).

[16] Industry participants point out two variations, which are small relative to price differences: the cost of protecting the underground portion of very tall buildings from water encroachment in Tel Aviv and potentially lower labor costs in the Beer Sheva district. These interviews were conducted for Genesove et al. (2020).

[17] The country is about the size of New Jersey, with about half its area a semi-arid, lightly populated desert. The transaction data we use lack information on the builder identity.

The regulatory tax would still reflect regulations. The regulatory tax rate will still measure the tax as a fraction of price, although now not as a fraction of marginal cost, but rather as a fraction of the sum of marginal cost and the constant markup.

Only a varying markup will bias the measured regulatory tax. In that case, the measured regulatory tax will exceed the actual regulatory tax by an amount equal to the excess of the markup over its minimum. Likewise, if firms differ in their efficiency, the measured regulatory tax will reflect firm inefficiencies as well.[18] Since both markups and inefficiencies should decrease with market size (Sutton, 1991; Syverson, 2004), the relationship between the measured regulatory tax and measures of market size should indicate how important a bias this is. In fact, Section 5.7 shows the estimated regulatory tax increasing strongly in population density. As an added check the appendix conditions also on the number of buildings constructed in the vicinity of the apartment over the period of the sample. We find only a weak economic relationship with inconstant sign between construction and the regulatory tax, suggesting little bias.

2.7.3 No Subsidization and Price Expectations

Below-cost prices would undermine our frontier estimates. Below-cost prices can be due either to government subsidization of construction costs, forced building beyond profit-maximizing heights or expectation mistakes. Although there have been periods of government subsidization of construction costs, notably in response to the mass immigration from the ex-Soviet Union of the early 1990s (Genesove, 2021), these were absent during our period of analysis.

If builders expect a higher apartment price than what materializes, price may not cover cost. We do not think this is a major concern, however. Building specific expectation mistakes can be included in measurement error: under rational expectations, the observed price is a random deviation from the expected price, which is the relevant price for determining the cost frontier. As modeled, however, measurement error fails to cover market-wide misperceptions. This should not be an issue, however, as parsimonious models forecast prices over the sample period fairly well. A yearly AR(1) specification with a trend and structural break in trend at 2009 yields a

[18]The frontier would now be the cost curve of the most efficient firm with the lowest markup in the least regulated market. This approach is in the spirit of Sutton (1991), who in estimating the lower envelope of concentration ratios across normalized market sizes assumes a positive probability of maximally competitive conditions. Note that the spatial component of the lower bound for the regulatory tax in Section 2.4 can accommodate minimum markups that are weakly complementary with spatial amenities in the same manner as housing quality.

root mean squared error of 0.018.[19] Also, we do not see large variation in mean price differences across transactions within buildings that take place the year before, the year of or the year after construction, as we would expect to see if substantial surprises were common. Finally, when repeating our estimation on the pre-2008 period only, a period with relatively stable prices, we get similar results (see Figure 8).

3 Estimation

3.1 The Model

Consider the log prices of apartments in buildings of height h,

$$y_{kij} = g + u_k + w_{ki} + v_{kij}, \qquad k = 1, \ldots, K,\ i = 1, \ldots, n_k,\ j = 1, \ldots, J_{ki}, \tag{11}$$

where y_{kij} is the observed log price per square meter of apartment j in building i in bloc k, g is the frontier, u_i is the deviation from the frontier, w_{ki} is building-level measurement error, and v_{kij} is apartment-level measurement error.[20] The distributions of $u_i \in [0, \infty)$, $w_{ki} \in (-\infty, \infty)$, and $v_{kij} \in (-\infty, \infty)$ can depend on height, but the lower bound of u_i and the means of w_{ki} and v_{kij}, are all equal to zero independently of height.[21]

The conditional mean of (11) is,

$$E[y|h] = g(h) + E[u|h], \tag{12}$$

as $E[w|h] = E[v|h] = 0$ by assumption. Equation (12) demonstrates the importance of having the parameters of the distribution of u depend on h. Were these parameters, instead, the same across heights, then frontier estimates would equal the height-specific means, up to a common constant, making frontier analysis pointless. Further, in this case, any endogeneity bias present in conditional mean analysis would also be present here. Hence, u (and v and w) are allowed to have separate parameters for each height. However, u's distribution originates in the joint distribution of demand and supply shocks through the equilibrium condition. Thus, unlike frontier costs $g(h)$, the parameters of u's distribution will not be "deep parameters."

[19] Housing prices rose steeply after the Bank of Israel drastically reduced interest rates at the beginning of 2009, as part of the coordinated, worldwide central bank response to the financial crisis. Unanticipated price increases do not threaten identification of the frontier.

[20] The log price is $y = \ln(P) = \ln(G+U) = \ln[G(1+U/G)] \approx \ln G + U/G \equiv g + u$.

[21] Spatial dependence of u_j is considered in the robustness section.

3.2 Variances

Without invoking any distributional assumptions, we identify and estimate the variances of u, v, and w using the hierarchical structure (see Appendix A.2 for formulas). Specifically, conditional on height h, the variance of the apartment-level measurement error v is identified by within building variation in apartment time-adjusted prices, the variance of the building-level measurement error w is identified by within bloc variation in building time-adjusted prices, and the variance of the deviations u is identified by variation in prices (unadjusted for time) across both bloc and time.

3.3 The Frontier

We estimate the frontier by maximum likelihood.[22] At height h, assume that $v_{kij} \sim N(0, \sigma_v^2(h))$ and $w_{ki} \sim N(0, \sigma_w^2(h))$ are normal and that $u_k \sim TN(\mu_u(h), \sigma_u^2(h))$ is the normal distribution truncated from below at zero.[23] Using the hierarchical structure to identify the variances allows us to estimate the error distributions on the basis of second moments only. This is in contrast to a cross-section of data, where skewness in the data is crucial to identification.[24]

The global maximum of the log likelihood, constrained so that average cost decreases to MES and marginal cost increases thereafter, is attained by grid search and Dijkstra's algorithm,

$$\{\widehat{MES}, \widehat{g}, \widehat{\mu}_u\} = \underset{\substack{mes \in \{1,\dots,H-1\} \\ g \in \mathbb{R}^H, \nu_u \in \mathbb{R}^H}}{\operatorname{argmax}} \sum_{h=1}^{H} \mathscr{L}_h(g_h, \nu_{uh}, \cdot), \tag{13}$$

$$\text{s.t. } g_{mes} \leq g_{mes-1} \leq \dots \leq g_1 \text{ and } g_{mes} \leq g_{mes+1} \leq \dots \leq g_H, \tag{14}$$

where $\mathscr{L}_h(g_h, \nu_{uh}, \cdot)$ is the log likelihood at height h (see Appendix A.3 for details and formulas). The constraint allows for $\widehat{MES} = 1$ and so no economies of scale.[25]

[22] We have considered alternative estimators. The commonly used, and convenient, priors of Bayesian-based estimators are not readily compatible with a frontier objective, while minimum-price-adjusted estimators converge slowly (see, Goldenshluger and Tsybakov, 2004).

[23] If $x \sim N(\mu_x, \sigma_x^2)$ then $x \mid a \leq x < b$ is truncated normal. Although the truncated normal is not new to the SFA literature, the half-normal distribution (i.e., $\mu_x = 0$) is more commonly used (e.g., Cai et al., 2021). However, this assumes deviations from the frontier are clustered near it, which we do not find in general. We consider the censored normal and folded normal in the robustness section.

[24] We find that the absolute value of the skewness of u is generally below 0.5, which is considered small, across most heights.

[25] We also present estimates that maximize the log likelihood at each height without constraints and estimates that maximize the log likelihood of a quartic cost function subject to a continuous version of the constraints.

3.4 Regulatory Tax Rates

This section describes how to estimate and bound expected regulatory tax rates of error-free prices. Using the distributions from Section 3.3 that $u \sim TN(\mu_u, \sigma_u^2)$ and $\eta \sim N(0, \sigma_\eta^2)$, where $\sigma_\eta^2 = \sigma_w^2 + \sigma_v^2/J$ for building price and $\sigma_\eta^2 = \sigma_w^2 + \sigma_v^2$ for apartment price, we get,[26]

$$u|u+\eta = y-g \sim TN\Big(\frac{\mu_u\sigma_\eta^2 + (y-g)\sigma_u^2}{\sigma_u^2+\sigma_\eta^2}, \frac{\sigma_u^2\sigma_\eta^2}{\sigma_u^2+\sigma_\eta^2}\Big). \tag{15}$$

Assuming that deviations from the frontier are entirely due to regulatory restrictions (taking into account the discreteness of height) the expected regulatory tax rate based on (1) is,

$$E[\frac{1}{G(h)e^u}\mathrm{RT}(G(h)e^u, h)|y-g(h)]. \tag{16}$$

where u is drawn from (15), conditioned on $y_i - g(h_i)$. However, if quality differs systematically over location then deviations also include quality. In this case, the lower bound based on (9) is,

$$E\Big[\frac{1}{G(h_i)e^{u_i}} \cdot \min_{\kappa_{Si}\in[0,1]} \max_{j\in\Omega_i(d)} \max\{0, G(h_j) - G(h_i+1) - (G(h_j)e^{u_j} - G(h_i)e^{u_i}) + \kappa_T(1-T_{ij})G(h_j)e^{u_j} \\ + \kappa_{Si}(T_{ij}G(h_j)e^{u_j} - G(h_i)e^{u_i})\}|y_k - g(h_k), k=i, j\in\Omega_i(d)\Big], \tag{17}$$

where u_k, for $k = i, j \in \Omega_i(d)$, is drawn independently from (15), conditioned on $y_k - g(h_k)$.

4 Data

Apartment transaction data are obtained from CARMEN, the digitalized repository of buyer reports to the Tax Revenue Authority. The data include the transaction date, price, square meters, apartment floor, number of floors in the building, and year of construction. They also include a unique identifying number from the land registry for the bloc and parcel on which the building sits, where the parcel is a lower level geographical division than the bloc, one or more of which comprise a single bloc. In general, the building and bloc-parcel are coincident. However, for 300 buildings, or 1.6% of the observations, more than one building sits on the same parcel. We exploit these cases to identify the hedonic height effects presented in 2.5 and estimated below in Section 5.1, but drop them for the stochastic frontier analysis. The sample covers the period 1998 to 2017.[27]

[26] Appendix A derives the conditional density when u is truncated normal. Jondrow et al. (1982) derive the conditional density for the half-normal, which is the truncated normal with $\mu_u = 0$.

[27] We drop apartments with nominal prices in the bottom one percent and top one percent of the distribution.

We limit the sample to transactions from CARMEN for which (1) the year of the transaction is the year before, the year of or the year after the construction year, (2) the transaction is for 100% of the asset, (3) the property type is not a single family home, (4) none of the variables listed above is missing, and (5) there is at least one other transaction observed in the building. We adjust prices for apartment floor-space area by expressing them in per-square meters. To account for inflation, we convert prices to real 2017 values. These prices are adjusted for floor and height premia, as described in Section 2.5. To estimate the frontier and regulatory tax, we further adjust for changes in construction input prices (other than land) over time by dividing the real prices by the Israeli Central Bureau of Statistics' residential construction input prices index, expressed in 2017 values.

There are 7,429 blocs, 18,169 buildings, and 270,554 apartments in the sample.[28] The median bloc size is about 0.21km^2. Unconditional on height, the mean number of buildings in a bloc is about 7.5 in our transactions data. Conditional on height and the presence of at least one building, the mean number of buildings in a bloc is 2.4, with about 55% of these bloc-height combinations containing exactly one building.

Table 1 shows apartment-level summary statistics of price (per square meter in real 2017 NIS and adjusted for cost) and the number of floors in the building (i.e., height), and building-level summary statistics of price (average price within a building) and the number of floors in the building. The mean real, input-price, height and floor-adjusted per square meter price is such that a standard 100 square meter apartment would sell for about 1.25 million NIS in 2017 shekels (about 350,000 USD at 2017 exchange rates).

Table 1: Summary statistics

	Mean	St. Dev.	Min	Med	Max
			Apartment		
Log price	9.35	0.38	8.40	9.34	10.53
Price	12,369	5,056	4,457	11,423	37,371
Number of floors	9.36	5.87	1	8	40
			Building		
Log price	9.36	0.39	8.49	9.35	10.50
Price	12,529	5,205	4,852	11,461	36,329
Number of floors	6.65	4.51	1	6	40

Notes: Prices per square meter in real 2017 NIS. There are 18,169 buildings, and 270,554 apartments.

[28]Table 8 in Appendix C.5 shows summary statistics for the number of observations by height.

The points in Figure 4 are building prices by height. There is a large dispersion in prices at nearly all heights, with the average ratio of third to first quartile price equal to 1.6 and the 95% to 5% price ratio equal to 2.7.

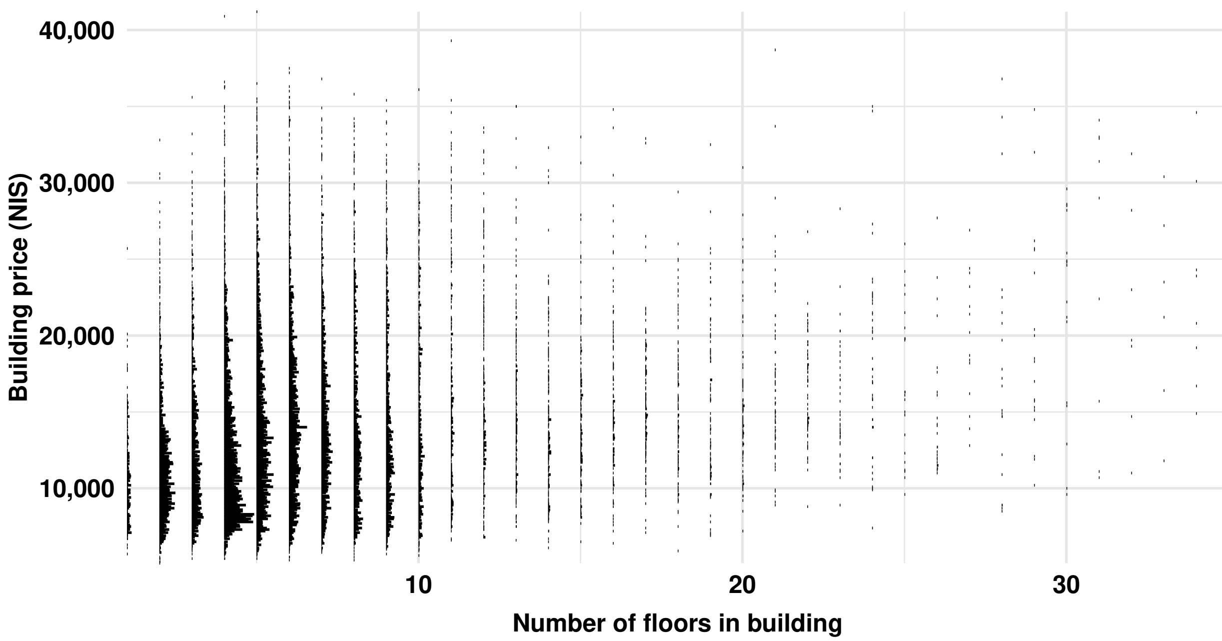

Figure 4: Frequency of building prices in NIS (rounded to nearest 100) by height.

5 Results

5.1 Apartment-Floor, Building-Height Adjusted Price

Adjusting prices for observable attributes is especially important in our context. On the one hand, consumers may be prepared to pay a premium, or demand a discount, for apartments on high floors or in tall buildings. On the other hand, building height varies with location, with taller buildings constructed in more attractive areas, as basic land use theory predicts. The challenge is to obtain an empirical counterpart to p of (10), the price after removing apartment-floor, building-height effects. An insufficiently flexible specification could easily assign apartment floor or building height effects to location effects, thus overstating the increase in the frontier at higher heights; too much flexibility could lead to excessive noise in the estimates. Our solution is to first estimate a fully saturated model of floor and height effects, and then, after inspecting the estimates, choose a reasonable restricted model. The function m in (10) is identified using variation in apartment floor within a building and variation in building height within a parcel, as some parcels have more than one building on them.[29] We then subtract the estimated floor and height effects from the observed price and add back in the effects pertaining to a second-floor apartment in a 4-floor building. This is the price used in the remainder of the analysis.

5.2 Variances

Figure 5 shows the estimated standard deviations, by height, of apartment-level measurement error v (in blue), building-level measurement error w (in red), and deviations from the frontier u (in purple), using (18)-(20) in Appendix A.2. The measurement error variances are estimated using residuals of a nonparametric regression of log price on transaction day. The deviations variance is then estimated using log prices and the estimated measurement error variances. Thus the variance of deviations($\approx$ regulations) is obtained from variation in prices (unadjusted for time) across both bloc and time, while the variances of measurement errors partials out time effects. For some of the higher heights, the degrees of freedom at the building-level are small or zero (see Table 8 in Appendix C.5) so that the estimated building-level measurement error variances do not exist or are negative, and so are missing from the figure. To deal with these cases and to avoid excessively noisy estimates, we smooth the measurement error variances using polynomial series estimates, with the polynomial degrees chosen by cross validation. The resulting curves are relatively flat. We do not smooth the standard deviations of u. Allowing these standard deviations to be unrestricted functions of height avoids imposing any endogeneity bias, as we discussed underneath (12).

The figure shows that the estimated standard deviation of u is on average about 4 times the estimated standard deviation of building error and about 2.5 times the estimated standard deviation of apartment error. Thus the variance of regulation is an order of magnitude larger than the combined measurement error variance. The standard deviations of the measurement errors, however, are clearly nontrivial.

[29] See Appendix A.4 for details. We normalize $m(2,4) = 1$, so that the adjusted price represents a second-floor apartment in a 4-floor building at the given location.

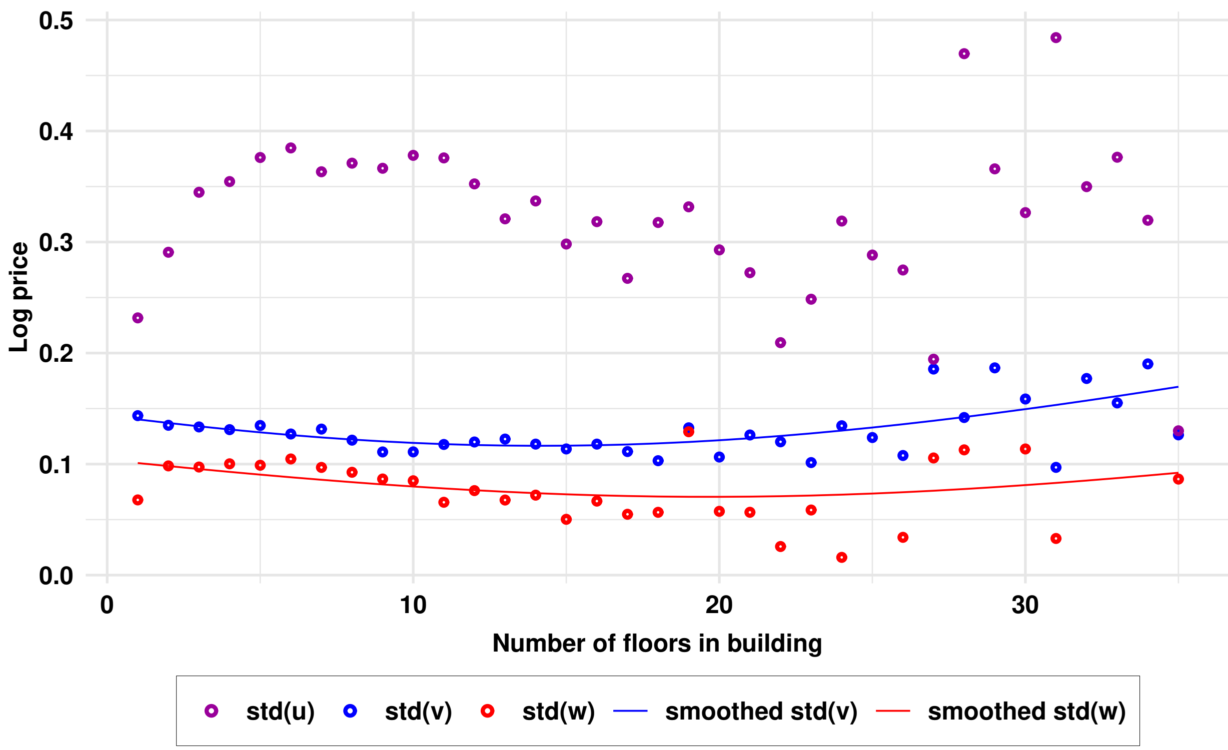

Figure 5: The red, blue, and purple points are estimated standard deviations based on (18)-(20). The red and blue curves smooth the estimates with series estimators.

5.3 The Frontier

Figure 6 shows our constrained ML frontier estimates from (13)-(14) (see also Appendix C.4). The estimates decrease until MES at five floors, increase, and then remain constant before increasing steeply. Although the upper confidence band admits marginal costs that are increasing beyond MES, each parametric bootstrapped sample produced a frontier that had long stretches of constant marginal costs.[30] The figure also shows mean and minimum building prices. The differences between mean prices and the ML estimates, along with the relative sizes of the variances estimated in Section 5.2, show that multi-floor housing markets must be highly regulated, with some building prices more than six times frontier prices. A striking difference between mean prices and the ML estimates, is that the former increase sharply at low heights but the latter decrease. Minimum prices are consistent estimators for the frontier absent measurement error (see Section 2.1) but with measurement error, at low heights, where there are many buildings with data on just two

[30]Let $(\widehat{g}(h), \widehat{\sigma}_v^2(h), \widehat{\sigma}_w^2(h), \widehat{\sigma}_u^2(h), \widehat{\mu}_u(h))$ be the ML estimates. The parametric bootstrap at height h randomly draws v_{kij}^* from $N(0, \widehat{\sigma}_v^2(h))$, w_{ki}^* from $N(0, \widehat{\sigma}_w^2(h))$, and u_k^* from $TN(\widehat{\mu}_u(h), \widehat{\sigma}_u^2(h))$. The bootstrapped observation is $y_{kij}^* = \widehat{g}(h) + u_k^* + w_{ki}^* + v_{kij}^*$.

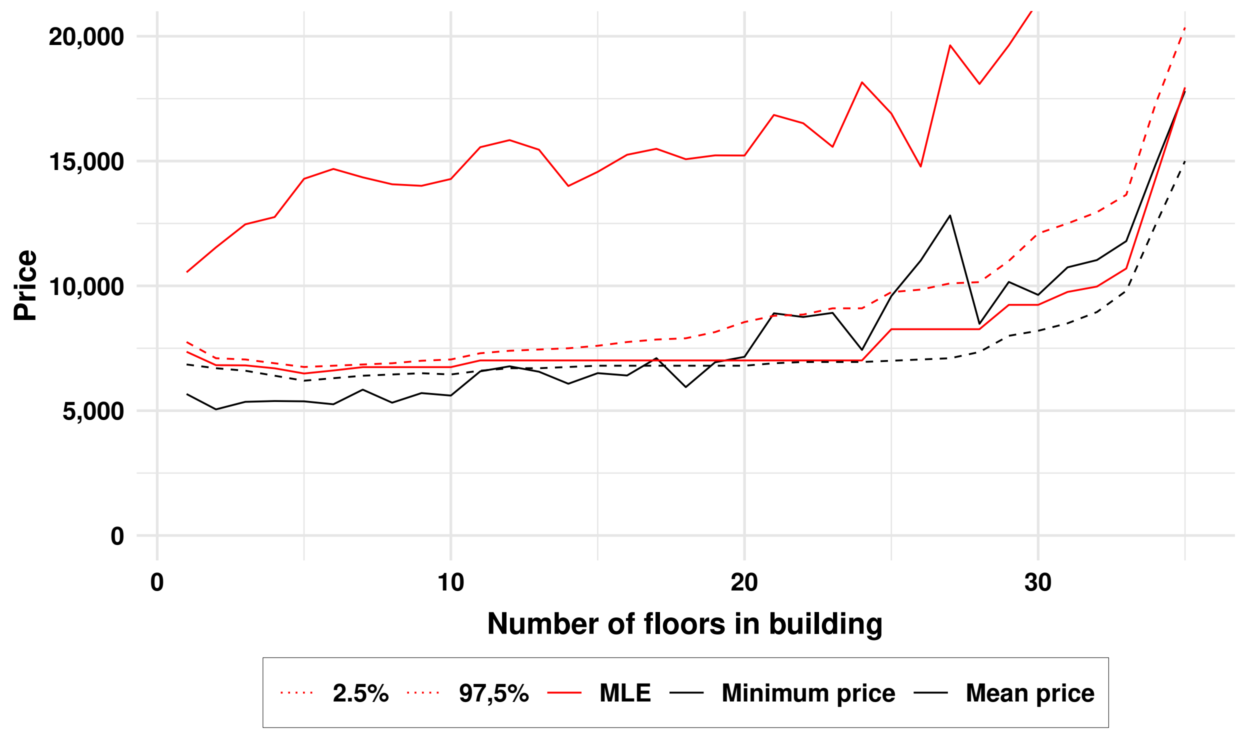

Figure 6: The minimum and mean building prices and constrained ML estimates with 95% confidence bands using 200 parametric bootstrapped samples.

apartments, it is likely that some building has large negative measurement error and is relatively unregulated, making minimum prices biased downwards as frontier estimates. At high heights, there are relatively few buildings and so minimum prices will tend to be biased upwards as frontier estimates.

Figure 7 shows alternative frontier estimates: a scatter plot of ML estimates of the frontier obtained at each height separately by maximizing the log likelihood (13), and smooth AC and MC estimates from the constrained maximum likelihood of a quartic cost function as in (24)-(26) in Appendix A.3. The constrained ML estimates from Figure 6 are also shown. Across all estimates, the average cost at MES is about 10% lower than the average cost of constructing a one-floor building. The marginal cost initially increases, then remains flat, before increasing steeply reflecting that building upwards becomes increasingly difficult at high heights. This is consistent with previous research (e.g., Glaeser et al., 2005) and discussions with industry experts (see footnote 16).

Table 2 compares buildings near the frontier, defined as buildings with average apartment price

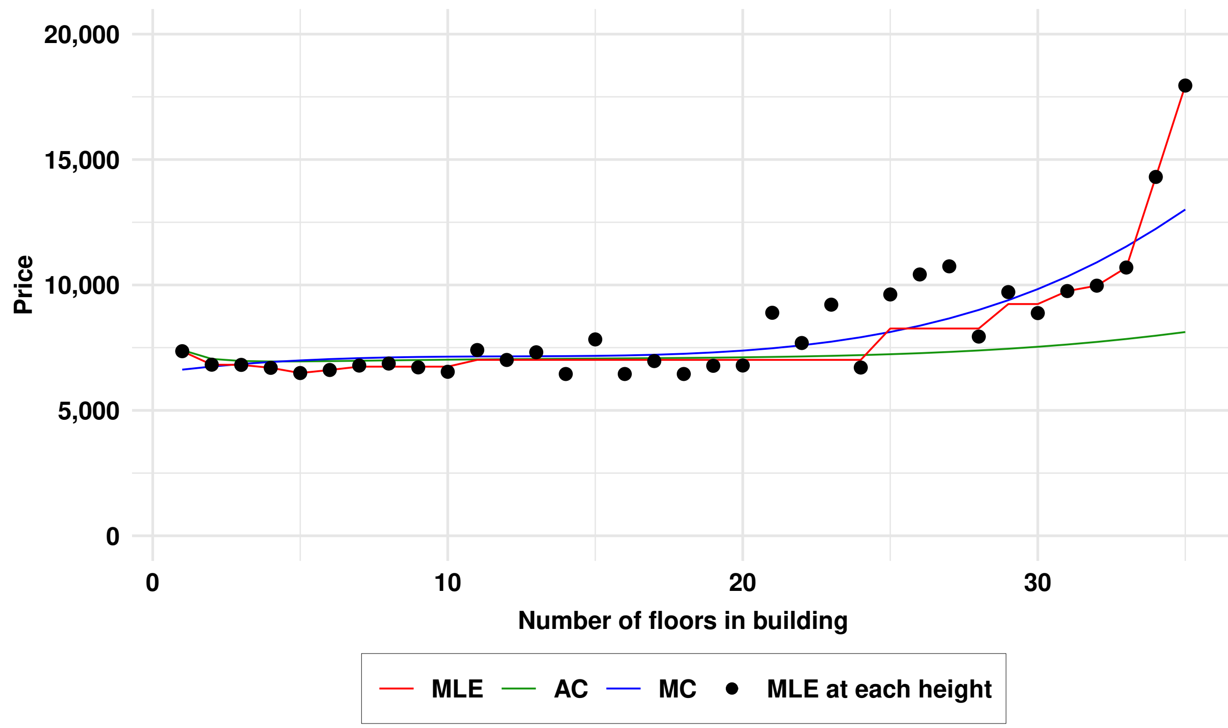

Figure 7: The constrained ML estimates, the smooth ML estimates using a quartic cost function, and the ML estimates for each height separately.

Table 2: Comparison of full sample and near frontier

	Full sample		Near Frontier	
	Mean	St. Dev.	Mean	St. Dev.
		Apartment		
Regulatory tax rate	0.45	0.15	0.12	0.04
Distance to city center	2.43	1.56	1.89	1.22
Density (1km radius)	0.5	0.5	0.31	0.27
Density (4km radius)	0.32	0.27	0.14	0.14
Distance to Tel Aviv city (km)	37.74	35.58	70.59	29.46
		Building		
Regulatory tax rate	0.48	0.16	0.11	0.05
Distance to city center	2.43	1.57	1.82	1.32
Density (1km radius)	0.62	0.57	0.26	0.22
Density (4km radius)	0.41	0.34	0.13	0.12
Distance to Tel Aviv city (km)	37.89	38.50	79.96	28.70

Notes: We remove observations with missing geographical coordinates so that there are 13,102 buildings and 206,835 apartments in the full sample and 350 buildings and 7,215 apartments near the frontier. Distances are in kilometers. Densities are in 10,000's per km^2.

at most 5% greater than the frontier, to the full sample of newly constructed buildings.[31] About 4% of the full sample is near the frontier. Relative to the full sample, housing near the frontier is about twice as far from the city of Tel Aviv, the country's commercial center. Depending on the radius and whether we look at buildings or apartments, 'Near Frontier' housing is in areas with average densities, in 10,000's per km^2, between 0.26 to 0.62 that of the full sample. The smaller standard deviations for 'Near Frontier' indicate a greater homogeneity of this sub-sample relative to the full sample. Although these buildings are further away from Tel Aviv, they are, perhaps surprisingly, closer to their own city centers, but the standard deviation indicates a large degree of disparity.

Consistent with our general view of regulatory variation as extremely local, buildings near the frontier are well represented throughout the country, with 59 of the 160 cities in Table 2 having at least one building near the frontier. Seven districts contain over 99% of buildings near the frontier. The remaining three districts are those closest to Tel Aviv.

5.4 Robustness of the Frontier

In our primary analysis, u_k, representing deviations from the frontier, and approximating regulatory taxes, follows a truncated normal distribution. To test the sensitivity of our results to this assumption, we considered alternative distributions for u_k, including a folded normal, a half-normal, and a zero-censored normal. The latter two assume a prevalence of minimally regulated buildings, which is inconsistent with our estimates, where the mean of u_k is often much larger than its variance. Nevertheless, as can be seen in Figure 8, our estimates are robust to different distributional choices for the deviations.

To assess robustness to spatial dependencies, we specify the spatial autoregressive relationship $u_k = \rho \sum_{l=1}^{K} \omega_{kl} u_l + \zeta_k$. We obtain estimates and bootstrapped confidence intervals (see, e.g., Jin and Lee, 2015) and the results, depicted in Figure 8, affirm that our frontier estimates are robust to spatially correlated regulations.

We also considered building-level regulations, modifying the model to $y_{kij} = g + u_{ki} + w_{ki} + v_{kij}$. Identification now depends on the skewness of the distribution of u_{ki} and the symmetry of the distribution of w_{ki}. The practical application of this model requires σ_u^2 to be sufficiently larger

[31] Table 2 and the analysis in Section 5.7 use the subset of the data with geographical coordinates.

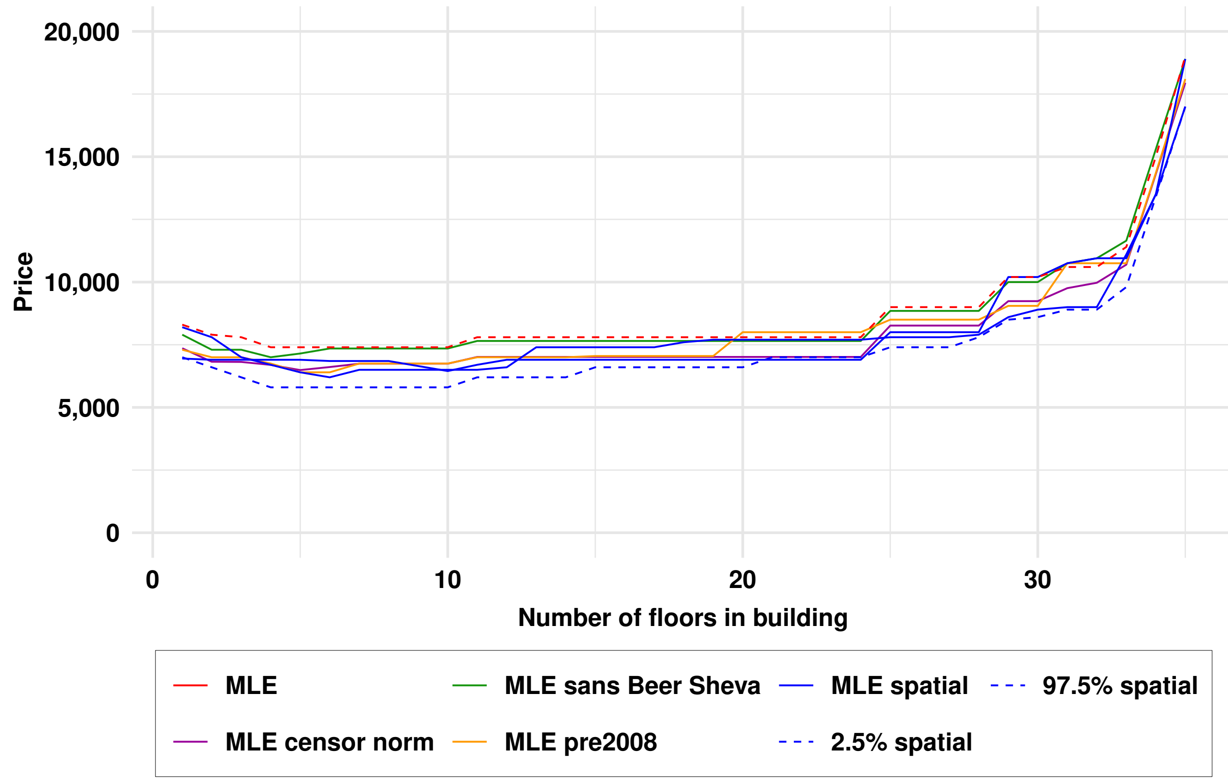

Figure 8: Robustness of ML estimate to ML estimate with censored normal regulations, ML estimate using pre-2008 data, ML estimates without the Beer Sheva district, and ML estimate that allows for spatial correlation in regulations.

than μ_u for the distinction between a truncated normal distribution and normal distribution to be discernible. Figure 8 shows similar estimates using building-level regulations.

To assess the robustness of differences in cost over space, we estimated the frontier excluding the Beer Sheva district, which may have have lower labor costs. The results, also shown in Figure 8, further confirm the robustness of our frontier estimation across different spatial contexts. Removing Tel-Aviv, where building near the acquifier is higher cost at large heights, will have no impact on the frontier since it has no observations near the frontier. Next we examined the robustness of differences in cost over time by estimating the frontier without adjusting for temporal cost differences and separately by restricting our dataset to pre-2008 data, a period marked by significant housing price increases. The results, presented in Figures 8 and 15, demonstrate the stability of our estimates over time.

Additionally, we employed the best linear unbiased estimator (BLUE) and the best linear unbiased predictor (BLUP), with uninformative and normal priors on $g+u_k$, respectively. While these approaches are conventionally used for mean estimation, in our case, they are less suitable since $g+u_k$ represents a minimum. These estimates of the frontier shown in Figure 15 are similar to our ML frontier estimates.

Lastly, we estimated the frontier by a sample size adjustment to the minimum price, as proposed by Goldenshluger and Tsybakov (2004). This involved estimating the frontier as $\widehat{g}_{\mathrm{GTm}} = \min_{k,i}\{\frac{1}{m}\sum_{j=1}^{m} y_{kij}\} + \widehat{\sigma}_{GTm}\sqrt{2\ln(n)}$. The results, are illustrated in Figure 15, with shape similar to our estimates but substantially higher perhaps due to slow convergence rates.

The comprehensive nature of these robustness checks, encompassing distributional assumptions, spatial and temporal variations, and alternative estimation techniques, underscores the reliability and validity of our frontier estimation approach.

5.5 The Frontier Elasticity of Substitution of Land for Capital

The elasticity of substitution of land for capital is typically used to summarize housing production functions. Appendix B shows that it is equal to the elasticity of average to marginal non-land costs $\sigma = d\ln AC/d\ln MC$. The elasticity and isoquant curves implied by the smooth MC and AC estimates are shown in Figures 9a and 9b respectively. The elasticity is equal to zero at MES (AC is at its unique minimum here so $dAC = 0$ and the elasticity is zero), increases sharply because

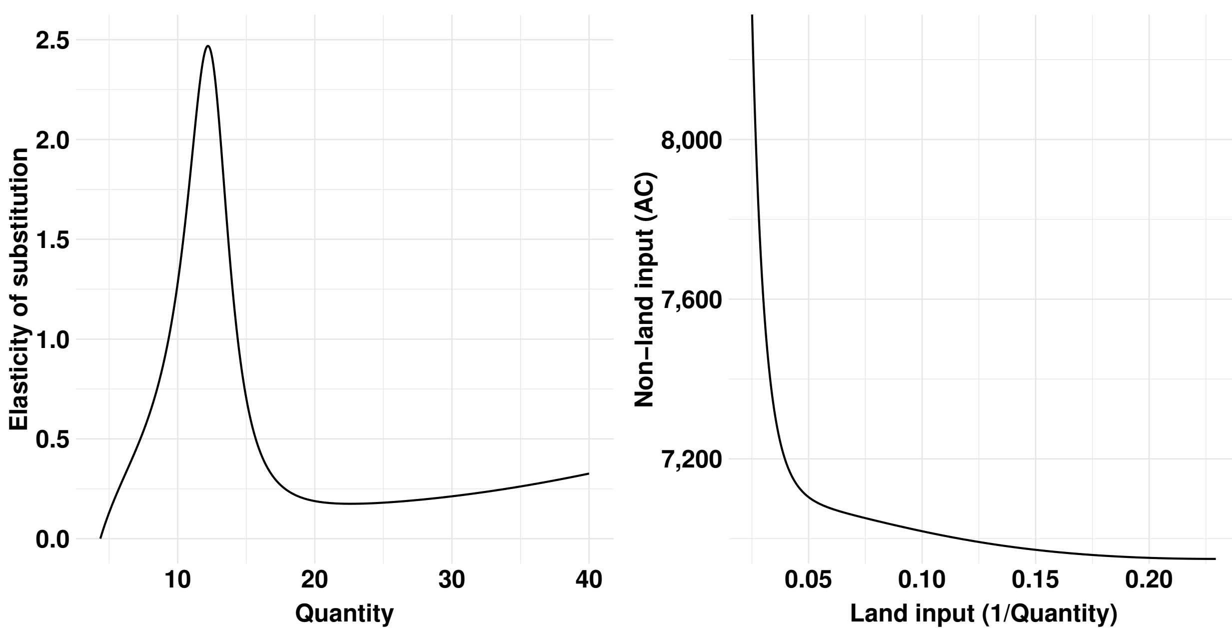

(a) The elasticity of substitution. (b) Isoquant curve.

Figure 9: (a) Elasticity of substitution of land for capital (b) Isoquant curve.

dMC ≈ 0 (this region corresponds to the near linear - i.e., perfect substitutability - segment of the frontier isoquant), then decreases sharply, and remains well below 0.5 thereafter. Most of the literature estimates the elasticity of substitution for small residential structures to be about unity (e.g., Ahlfeldt and McMillen, 2014) and the few elasticity estimates for tall residential buildings are about 0.5 (e.g., Ahlfeldt and McMillen, 2018). Our estimates of the elasticity suggest that substituting capital for land is difficult at low and high heights and easy at medium heights.

5.6 Regulatory Tax Rates

For each building we estimate the upper bound (based on (16)) and lower bounds (based on (17)) for the mean regulatory tax rate. Recall that the upper bound is simply the point estimate. The lower bounds use nearby buildings within distance $d \in \{1\text{km}, 2\text{km}, 3\text{km}, 4\text{km}\}$. The mean number of buildings within 1km, 2km, 3km, and 4km is 80, 195, 315, and 435 respectively. The existing home price regression yields an estimate of 0.0016 for κ_T, as reported in Appendix A.1. The estimated mean value of κ_{Si} is 0.65, with standard deviation 0.35.

We find a substantial extent of regulation, as measured by the regulatory tax. Across all buildings, the mean and standard deviation of the upper bound are 48% and 16%. Across all buildings with heights above MES, the mean and standard deviation of the upper bound are 49% and 17%.[32] Restricting to buildings with geographical coordinates, and using buildings within 1km, 2km, 3km, and 4km respectively, the lower bounds are 19%, 24%, 28%, and 31% with standard deviations 16%, 20%, 23%, and 24%. Restricting further to buildings with heights above MES, the lower bounds are 23%, 29%, 33%, and 37% with standard deviations 17%, 20%, 23%, and 25%.

Figure 10 shows the upper and lower bounds over time for buildings with heights above MES to 30. The lower bounds tighten substantially with time as housing prices increased, post-2008. In 2017, with housing prices having approximately doubled since 2008, the lower bounds for the four radii reach 34%, 40%, 44%, and 45%, with the upper bound at 53%. With a small estimate for κ_T, this demonstrates the greater usefulness of bounds in periods that follow high price growth.

[32]Across all apartments, the upper bound is 45%, with a standard deviation of 15%. Across all apartments in buildings with heights above MES (i.e., five floors), the upper bound is 46%, with a standard deviation of 15%.

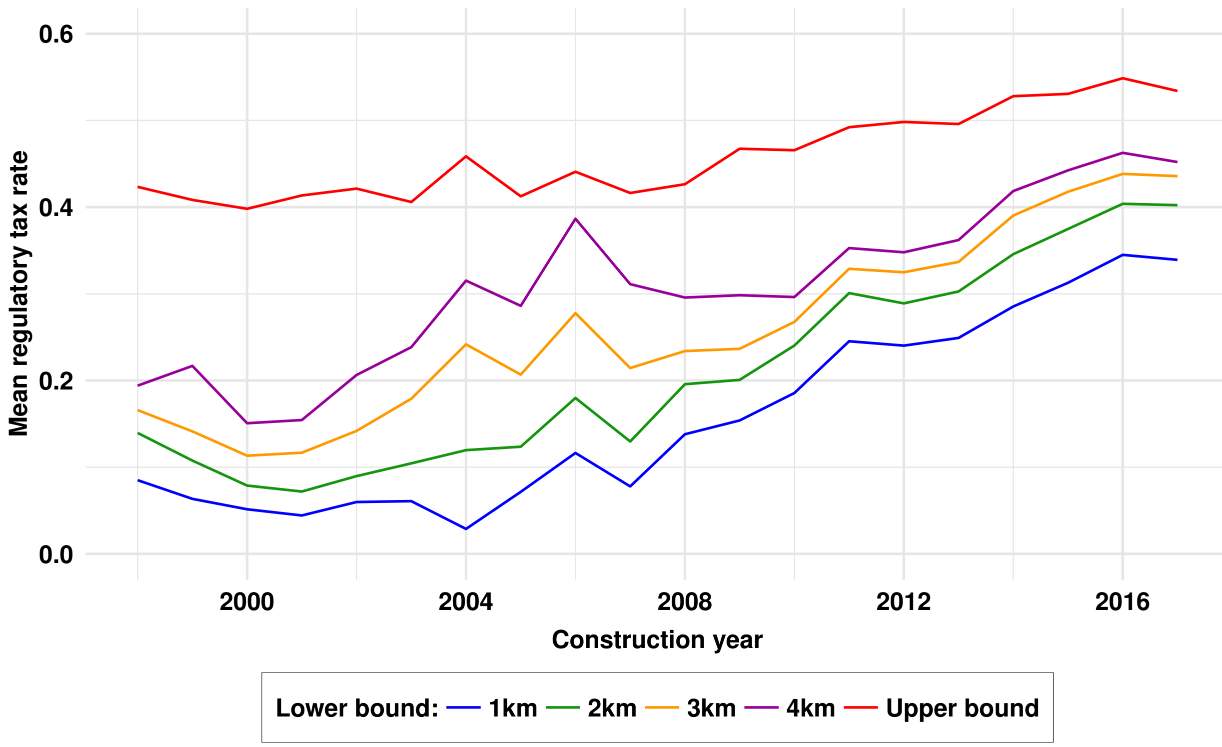

Figure 10: The upper and lower bounds for mean regulatory tax rates for buildings with heights above MES to 30.

5.7 Characterizing Regulatory Tax Rates

We characterize the estimated regulatory tax rate using (16) by the covariates distance to city center, density, and geographical location (summary statistics for these variables are shown in Table 2). The relationships between the regulatory tax and the covariates are shown graphically and through regression estimates in Table 3. These estimates are to be understood as descriptive only, and not causal.

We define the city center as the location within the city with the highest predicted price according to a nonparametric regression of observed building prices on buildings' geographical coordinates (using cross-validation for choice of bandwidth). This definition is consistent with monocentric city models, while obviating the need for non-price data and choosing between employment and consumption as the dominant agglomeration force.

Figure 11a shows the estimated quartic fit of a regression of estimated regulatory tax rates on distance, in kilometers, to the city center for the three largest cities: Jerusalem, Tel Aviv, and Haifa. The figure shows that in general, and where the relationship is precisely measured, the estimated regulatory tax rate decreases with distance to city center. This negative relationship between the

Table 3: Regressions

	Estimated regulatory tax rate					
	(1)	(2)	(3)	(4)	(5)	(6)
			Apartment			
Distance to city center	-	-	-0.004 (0.0002)	-	-	-0.005 (0.0002)
Density (1km radius)	0.090 (0.001)	-	-	0.011 (0.001)	-	-
Density (4km radius)	-	0.278 (0.001)	-	-	0.063 (0.003)	0.070 (0.003)
City fixed effects	No	No	Yes	Yes	Yes	Yes
R^2	0.084	0.229	0.564	0.563	0.564	0.565
			Building			
Distance to city center	-	-	-0.006 (0.001)	-	-	-0.008 (0.001)
Density (1km radius)	0.103 (0.002)	-	-	0.014 (0.002)	-	-
Density (4km radius)	-	0.315 (0.004)	-	-	0.081 (0.009)	0.097 (0.010)
City fixed effects	No	No	Yes	Yes	Yes	Yes
R^2	0.125	0.304	0.669	0.668	0.669	0.672

Notes: Standard errors are in parentheses underneath the coefficients. Distance to city center is in kilometers. Densities are in 10,000's per square kilometer. There ares 13,102 buildings and 206,835 apartments. The top section has outcomes at the apartment level, and standard errors clustered at the building level. The bottom section has outcomes as the average apartment price in a building.

regulatory tax and distance to city center is supported by the regression estimates in Columns (3) and (6) in Table 3. The negative relationship is consistent with Tan et al. (2020), where the city center is defined as the location with the brightest lights at night.

We measure population density at a building's location as the number of people residing in 1995 (three years before the start of our sample period), in 10,000's, within a 1km or 4km radius.[33] Figures 11c and 11d contain scatter plots of estimated regulatory tax rates versus density, with an overlaid quartic fit and 95% pointwise confidence bands. Measuring the density with a 1km radius, Figure 11c shows that the mean tax rate, starting at 0.40 in unpopulated areas, increases until a maximum of 0.58 at about a density of 16,500 people (the 94th quantile of the density).

[33]To be precise, the density is the weighted average of 1995 population densities of census statistical areas within a 1km or 4km radius of the building, where the weight is the statistical area's contribution in area to the intersection of the circle of radius 1km and Israel's land mass.

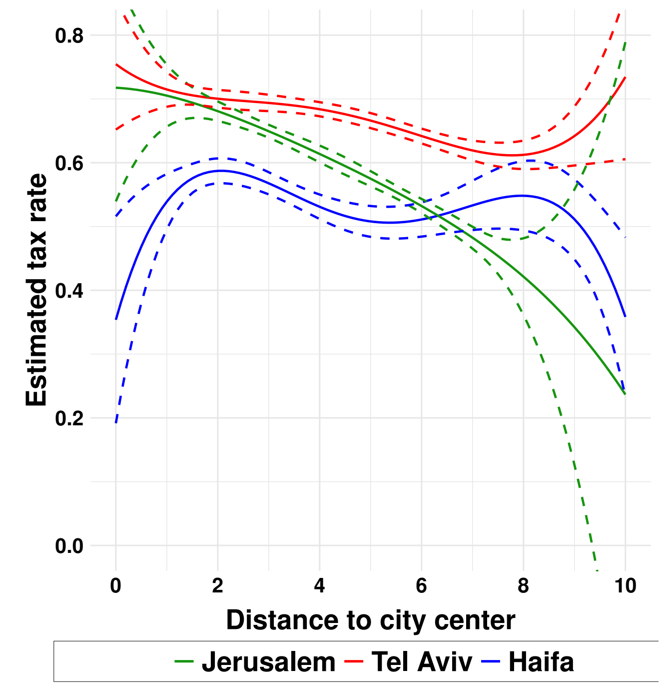

(a) Tax rates by dist. to center.

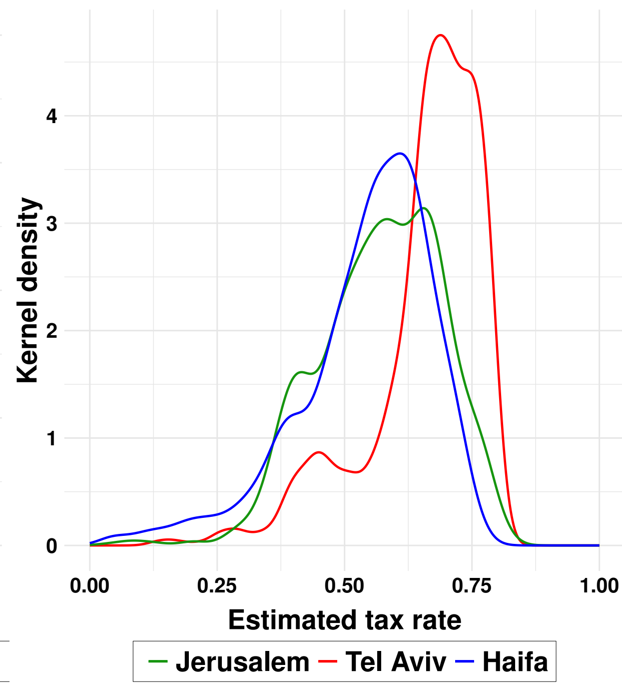

(b) Tax rates by city.

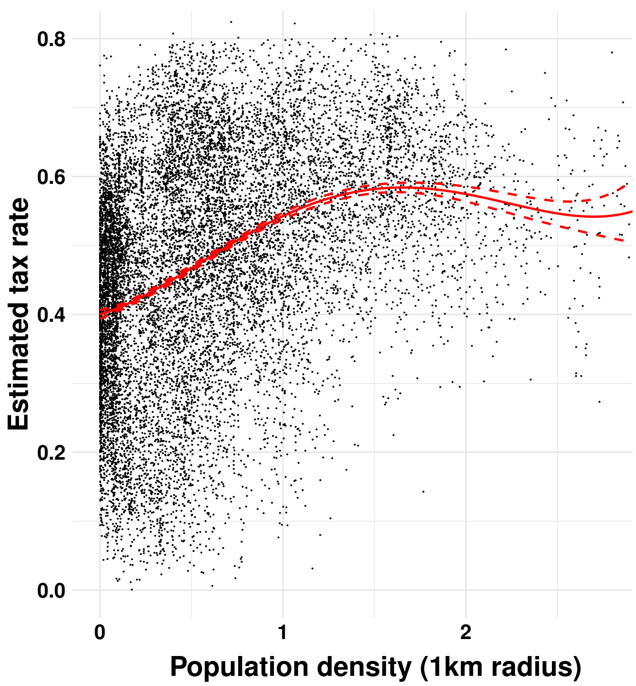

(c) Tax rates by density - 1km radius.

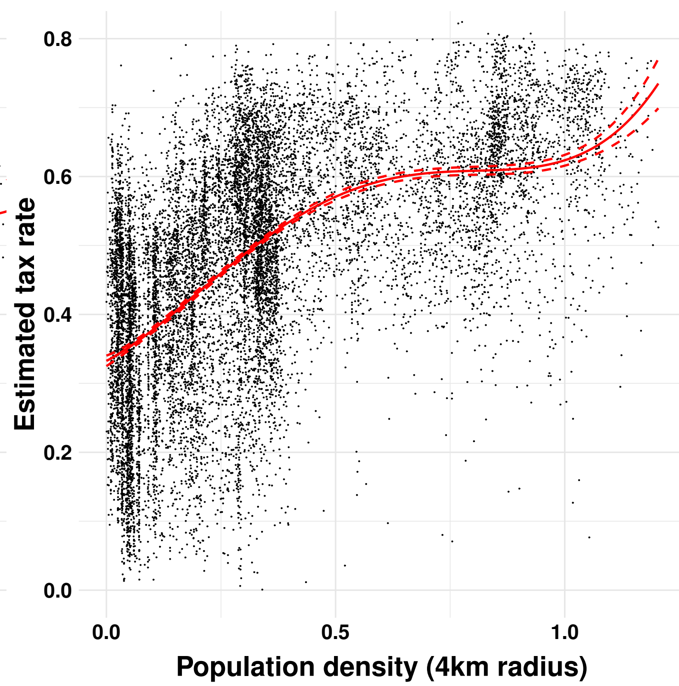

(d) Tax rates by density - 4km radius.

Figure 11: (a) The quartic fit and 95% confidence bands of regressions of the estimated regulatory tax rates on distance to city center for Jerusalem, Tel Aviv, and Haifa, (b) The kernel densities of the estimated regulatory tax rates in these cities, (c) The estimated regulatory tax rate by density, in 10,000's, per km^2 for radius 1km, the quartic fit, and 95% pointwise confidence bands, (d) The estimated regulatory tax rate by density, in 10,000's, per km^2 for radius 4km, the quartic fit, and 95% pointwise confidence bands.

Measuring the density with a 4km radius, Figure 11d shows that the mean tax rate, starting at 0.33 in unpopulated areas, increases to about 0.73. On average, as seen in Columns (1) and (2) of Table 3, for every additional 10,000 people per square kilometer, the tax rate is about twenty percent higher measured with a 1km radius, and three percent higher with a 4km radius. The goodness of fit in the regressions, measured by R^2, improves as the radius increases from 0.075km to 4km. This suggests that the density immediately surrounding a building is less predictive of the tax rate compared to the density of a broader area. A positive relationship between the tax rate and density is reminiscent of Hilber and Robert-Nicoud (2013), who show a positive relationship between the developed share of developable land and the Wharton Index, consistent with their theoretical model of incumbent landowners protecting their asset value. In contrast to the Wharton Index, our measure of regulation is cardinal.

The large increases in R^2 when city fixed effects are added in the latter columns of Table 3 show that the jurisdiction itself, and not just its overall density, is important. Figure 11b shows the kernel density of the estimated regulatory tax rate for the three largest cities: Jerusalem, Tel Aviv, and Haifa. Tel Aviv, which boasts the highest housing prices in the country, has the highest tax rates among the three. This is just an example of a more general relationship in the data, that higher priced cities are characterized by higher regulatory taxes. As the scatter plot in Figure 12 shows,

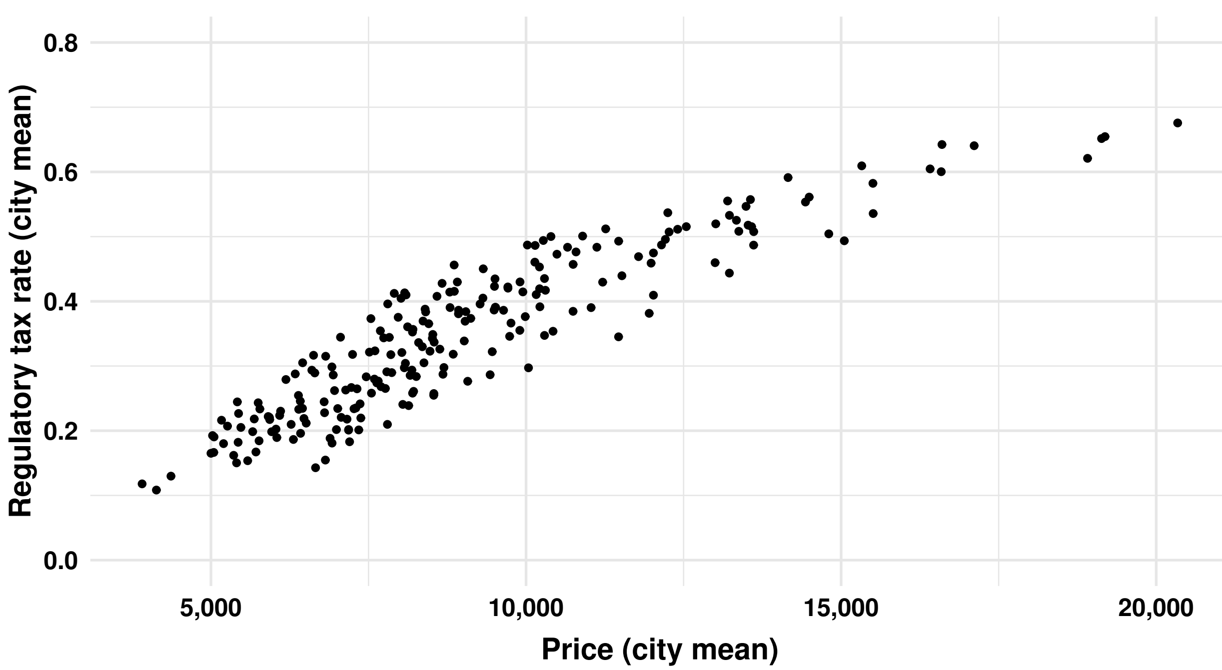

Figure 12: The city mean regulatory tax rate against the city mean apartment price.

the relationship is tight. This is not surprising given the relative flatness of the frontier. However, it is not inevitable - a scenario in which multi-unit housing is restricted in low demand areas only, say the suburbs, would yield a negative relationship. The positive relationship is consistent with predictions in of greater regulation in high amenity cities (Hilber and Robert-Nicoud, 2013).

5.8 Case Studies: Regulation Over Time in Newly Established Cities

The newly established cities of Modiin (situated about halfway between Tel Aviv and Jerusalem) and Elad (about 25 kilometers east of Tel Aviv) offer useful case studies. Modiin and Elad were planned in the 1990s. Modiin's first residents arrived in 1996 and Elad's in 1998. By 2019, Modiin had about 90,000 residents, most of high socioeconomic status, while Elad had about 50,000 residents, most religious and of low socioeconomic status. Since many political economy models of housing regulation locate the source of regulation in home owners' attempts to increase, or at least protect, the asset value of their home, it is interesting to document the degree of regulation in newly established cities, before homeowners become politically influential. Figure 13a shows the mean estimated regulatory tax rates for the full sample (in red), in Elad (in purple) from its year of establishment, and in Modiin (in blue) from two years after its establishment (the first year in our data). Elad's first residents moved in about two years after Modiin's, and Elad's curve shifted three years to the left, and a few points up, basically overlaps Modiin's curve. The figure shows that in their nascent years the regulatory tax rates were, although not zero, much lower than the national

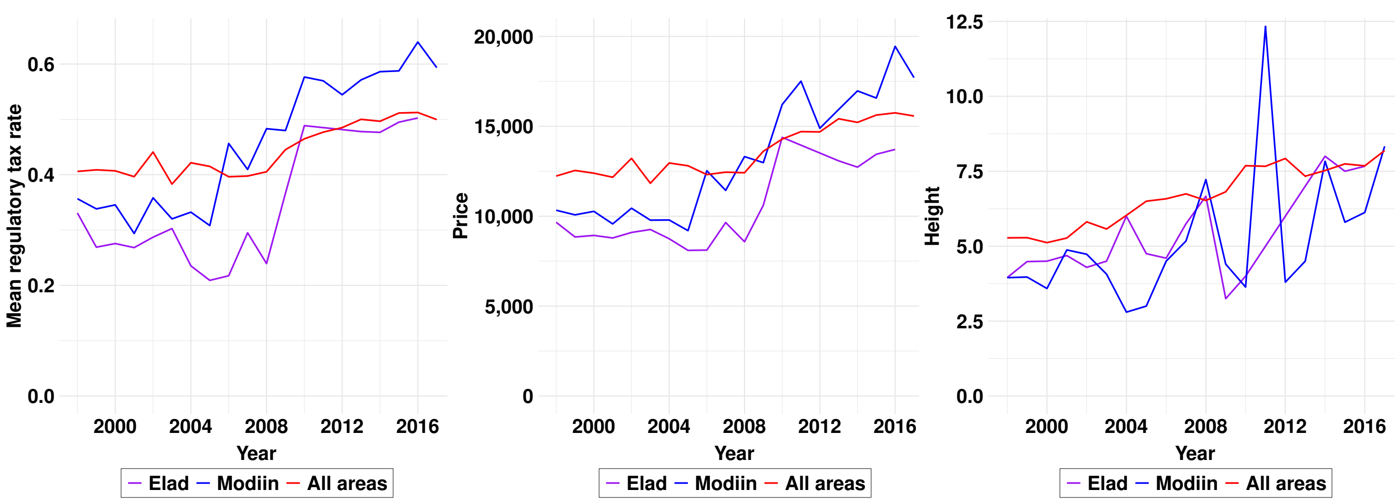

(a) Tax rate over time.

(b) Prices over time.

(c) Height over time.

Figure 13: The mean estimated regulatory tax rates, prices, and building heights over time.

average, and relatively stable. Then about six to eight years after their first residents moved in, the regulatory tax rates essentially doubled. Modiin's rate settled above the national average, while Elad's at the national average. Thereafter, their rates continue to increase at the national rate. Figures 13b and 13c show that the increase in regulation is coincident with a jump up in prices yet relatively stable building heights, suggesting that the sudden increase in the regulatory tax was driven by restrictions that were relatively fixed over time, and became more binding with the price increase.

6 Conclusion

Housing regulation can take many different forms that are often difficult to measure and aggregate, may be arbitrarily enforced, and is endogenous to building location, market conditions and price. Hence, estimating non-land mean costs by a conditional regression embeds unobserved regulatory conditions potentially biasing these estimates. In this paper, we show how to identify and estimate frontier costs in multi-floor housing using just observed prices and heights, identifying frontier marginal costs for heights above MES from variation in demand in unregulated markets and identifying frontier average costs for heights below MES from variation in demand and regulation. We adjust prices based on observed apartment floor and building height, and take into account building-level and apartment-level random housing quality differences and other measurement errors. When quality differs systematically over location and time, we assume local (weak) complementarity between quality and amenities and bound the regulatory tax.

Using data for newly constructed buildings in the Israeli housing market from 1998-2017, we estimate regulatory tax rates, finding a mean rate of 48%, with a standard deviation of 16%. Regulatory tax rates are higher in areas that are higher priced, denser, and closer to city centers. Measurement errors are small compared to regulation. When allowing for systematic differences in quality over location and time, we bound the mean regulatory tax rate in 2017 by 40% (using buildings within a 2km radius) and 53%. Most of that bound is derived from the availability of data on nearby buildings that were built during lower-priced periods and at heights where frontier costs did not significantly decrease. This is contingent on our estimates of a near-zero relationship between temporal demand shocks (period effects) and structural quality (cohort effects).

There is no presumption that regulation is either welfare-enhancing or welfare-detracting—a determination that would require additional sources of information. Nor is there a claim that the elimination of all regulation would lead to price reductions in the amount of the estimated regulatory tax, as the resultant price change would require one to know, at a bare minimum, the elasticity of overall housing demand. Rather the regulatory tax is a measure of the extent of regulation in the market.

Our analysis of regulation is price-based, defining a regulatory tax that relies on vertical deviations from the frontier (i.e., the difference between a building and frontier price at the building height). A quantity-based alternative would rely on horizontal deviations from the frontier (i.e., the difference between a building and frontier height at the building price). For example, in a counterfactual world where there is no regulation, and *holding prices constant*, our point estimates indicate that suppliers would build about 4.6 times higher, constructing about $3,400$ buildings instead of the $18,000$ or so in our sample, and so freeing up about 80% of the building footprint. Assessing the resource savings in this counterfactual world would require values for land and consideration of general equilibrium effects, as well as externalities such as congestion effects. One simple exercise, however, is to consider the land savings from building all apartments in buildings of heights 11 to 24, where marginal costs are constant according to our constrained ML estimates, in 24-story buildings instead. This would require 35% less land, but cost an additional 1% of non-land costs. Likewise, removing regulation so that apartments in shorter than MES-story buildings are built in MES buildings would also require 35% less land, along with saving 1% of non-land costs. We leave further analysis along these lines for future work.

References

Ahlfeldt, G. M. and McMillen, D. P. (2018), 'Tall buildings and land values: Height and construction cost elasticities in chicago, 1870–2010', *Review of Economics and Statistics* **100**(5), 861–875.

Ahlfeldt, G. and McMillen, D. P. (2014), 'New estimates of the elasticity of substitution of land for capital', *Working Paper* .

Albouy, D. and Ehrlich, G. (2018), 'Housing productivity and the social cost of land-use restrictions', *Journal of Urban Economics* **107**, 101–120.

Arrow, K. J., Chenery, H. B., Minhas, B. S. and Solow, R. M. (1961), 'Capital-labor substitution and economic efficiency', *Review of Economics and Statistics* pp. 225–250.

Brueckner, J. K., Fu, S., Gu, Y. and Zhang, J. (2017), 'Measuring the stringency of land use regulation: The case of China's building height limits', *Review of Economics and Statistics* **99**(4), 663–677.

Brueckner, J. K. and Helsley, R. W. (2011), 'Sprawl and blight', *Journal of Urban Economics* **69**(2), 205–213.

Cai, H., Wang, Z. and Zhang, Q. (2017), 'To build above the limit? Implementation of land use regulations in urban China', *Journal of Urban Economics* **98**, 223–233.

Cai, J., Feng, Q., Horrace, W. C. and Wu, G. L. (2021), 'Wrong skewness and finite sample correction in the normal-half normal stochastic frontier model', *Empirical Economics* **60**(6), 2837–2866.

Cheshire, P. C. and Dericks, G. H. (2020), "trophy architects' and design as rent-seeking: Quantifying deadweight losses in a tightly regulated office market', *Economica* **87**(348), 1078–1104.

Cheshire, P. C. and Hilber, C. A. (2008), 'Office space supply restrictions in Britain: The political economy of market revenge', *The Economic Journal* **118**(529), F185–F221.

Cheung, R., Ihlanfeldt, K. and Mayock, T. (2009), 'The incidence of the land use regulatory tax', *Real Estate Economics* **37**(4), 675–704.

Combes, P.-P., Duranton, G. and Gobillon, L. (2021), 'The production function for housing: Evidence from France', *Journal of Political Economy* **129**(10), 000–000.

Conley, T. G. (2010), Spatial econometrics, *in* 'Microeconometrics', Springer, pp. 303–313.

Coulson, N. E. and McMillen, D. P. (2008), 'Estimating time, age and vintage effects in housing prices', *Journal of Housing Economics* **17**(2), 138–151.

Czamanski, D. and Roth, R. (2011), 'Characteristic time, developers' behavior and leapfrogging dynamics of high-rise buildings', *The Annals of Regional Science* **46**(1), 101–118.

Epple, D., Gordon, B. and Sieg, H. (2010), 'A new approach to estimating the production function for housing', *American Economic Review* **100**(3), 905–24.

Florens, J.-P., Simar, L. and Van Keilegom, I. (2020), 'Estimation of the boundary of a variable observed with symmetric error', *Journal of the American Statistical Association* **115**(529), 425–441.

Fu, Y. and Somerville, C. T. (2001), 'Site density restrictions: measurement and empirical analysis', *Journal of Urban Economics* **49**(2), 404–423.

Fujita, M. (1989), *Urban Economic Theory*, Cambridge University Press.

Genesove, D. (2021), 'The Israeli housing market: Structure, boom and policy response', ch. 18 in *The Israeli Economy 1995-2017: Lights and Shadows in the Market Economy*, ed. by A. Ben-Basat, R. Gronau and A. Zussman. Oxford University Press.

Genesove, D., Levy, D. and Snir, A. (2020), 'The elasticity of substitution of capital for land in multi-unit housing', *Working Paper* .

Glaeser, E. L., Gyourko, J. and Saks, R. (2005), 'Why is Manhattan so expensive? Regulation and the rise in housing prices', *The Journal of Law and Economics* **48**(2), 331–369.

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

Goldenshluger, A. and Tsybakov, A. (2004), 'Estimating the endpoint of a distribution in the presence of additive observation errors', *Statistics & Probability Letters* **68**(1), 39–49.

Greene, W. H. (2008), 'The econometric approach to efficiency analysis', in *The Measurement of Productive Efficiency and Productivity Growth*, ed. by H.O. Fried, C.A.K. Lovell and S.S. Shelton. Oxford University Press, pp. 92–250.

Gyourko, J., Hartley, J. S. and Krimmel, J. (2021), 'The local residential land use regulatory environment across U.S. housing markets: Evidence from a new Wharton index', *Journal of Urban Economics* **124**, 103337.

Gyourko, J. and Krimmel, J. (2021), 'The impact of local residential land use restrictions on land values across and within single family housing markets', *Journal of Urban Economics* **126**, 103374.

Gyourko, J. and Molloy, R. (2015), 'Regulation and housing supply', ch. 19 in *Handbook of Regional and Urban Economics*, Vol 5b, ed. by G. Duranton, J. H. Henderson and W. Strange. Amsterdam: Elsevier.

Gyourko, J. and Saiz, A. (2006), 'Construction costs and the supply of housing structure', *Journal of Regional Science* **46**(4), 661–680.

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

Hall, B. H., Mairesse, J. and Turner, L. (2007), 'Identifying age, cohort, and period effects in scientific research productivity: Discussion and illustration using simulated and actual data on French physicists', *Econ. Innov. New Techn.* **16**(2), 159–177.

Hall, R. E. (1971), 'The measurement of quality change from vintage price data', in *Price Indexes and Quality Change: Studies in New Methods of Measurement*, ed. by Z. Griliches. Harvard University Press.

Henderson, J. V., Regan, T. and Venables, A. J. (2017), 'Building the city: Urban transition and institutional frictions', *CEPR Discussion Paper* .

Hilber, C. A. and Robert-Nicoud, F. (2013), 'On the origins of land use regulations: Theory and evidence from US metro areas', *Journal of Urban Economics* **75**, 29–43.

Hilber, C. A. and Vermeulen, W. (2016), 'The impact of supply constraints on house prices in England', *The Economic Journal* **126**(591), 358–405.

Hsieh, C.-T. and Moretti, E. (2019), 'Housing constraints and spatial misallocation', *American Economic Journal: Macroeconomics* **11**(2), 1–39.

Jin, F. and Lee, L.-f. (2015), 'On the bootstrap for moran's i test for spatial dependence', *Journal of Econometrics* **184**(2), 295–314.

Jondrow, J., Lovell, C. K., Materov, I. S. and Schmidt, P. (1982), 'On the estimation of technical inefficiency in the stochastic frontier production function model', *Journal of Econometrics* **19**(2-3), 233–238.

Katz, L. and Rosen, K. T. (1987), 'The interjurisdictional effects of growth controls on housing prices', *The Journal of Law and Economics* **30**(1), 149–160.

Klette, T. J. and Griliches, Z. (1996), 'The inconsistency of common scale estimators when output prices are unobserved and endogenous', *Journal of Applied Econometrics* **11**(4), 343–361.

Kotlarski, I. (1967), 'On characterizing the gamma and the normal distribution', *Pacific Journal of Mathematics* **20**(1), 69–76.

Kumbhakar, S. C., Parmeter, C. F. and Zelenyuk, V. (2022), 'Stochastic frontier analysis: Foundations and advances', *Handbook of Production Economics* pp. 1–40.

Ministry of Finance, Chief Economist Branch (2017), *Survey of the Residential Real Estate Industry for the Second Quarter, 2017* . [In Hebrew].

Molloy, R. (2020), 'The effect of housing supply regulation on housing affordability: A review', *Regional Science and Urban Economics* **80**(C).

Nechyba, T. J. and Walsh, R. P. (2004), 'Urban sprawl', *Journal of Economic Perspectives* **18**(4), 177–200.

Paciorek, A. (2013), 'Supply constraints and housing market dynamics', *Journal of Urban Economics* **77**, 11–26.

Piazzesi, M., Schneider, M. and Stroebel, J. (2020), 'Segmented housing search', *American Economic Review* **110**(3), 720–59.

Pollakowski, H. O. and Wachter, S. M. (1990), 'The effects of land-use constraints on housing prices', *Land economics* **66**(3), 315–324.

Rognlie, M. (2016), 'Deciphering the fall and rise in the net capital share: Accumulation or scarcity?', *Brookings Papers on Economic Activity* **2015**(1), 1–69.

Rubin, Z. and Felsenstein, D. (2017), 'Supply side constraints in the israeli housing market—the impact of state owned land', *Land Use Policy* **65**, 266–276.

Rubin, Z. and Felsenstein, D. (2019), 'Is planning delay really a constraint in the provision of housing? Some evidence from Israel', *Papers in Regional Science* **98**(5), 2179–2200.

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

Schwarz, M. and Van Bellegem, S. (2010), 'Consistent density deconvolution under partially known error distribution', *Statistics & Probability Letters* **80**(3-4), 236–241.

Sutton, J. (1991), *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*, Cambridge, MA: MIT Press.

Syverson, C. (2004), 'Market structure and productivity: A concrete example', *Journal of political Economy* **112**(6), 1181–1222.

Tan, Y., Wang, Z. and Zhang, Q. (2020), 'Land-use regulation and the intensive margin of housing supply', *Journal of Urban Economics* **115**, 103199.

Online Appendix

A Additional Estimation Details

A.1 Estimating κ_T

Our aim is to estimate κ_T through the relationship between period effects (transaction time) and cohort effects (construction time) in a regression of existing home prices on period, cohort, and age (transaction time less construction time, capturing depreciation), where the cohort effects are restricted to be a function of the period effects. In its most general form, this entails estimating

$$y_{its} = \gamma(t) + \delta(\gamma(s)) + \alpha(t-s),$$

where s is construction period, t is transaction period (so that $t-s$ is age), $\gamma(t)$ (which corresponds to its namesake in Subsection 3.4) are period effects, $\delta(\gamma(s))$ are cohort effects, and $\alpha(t)$ are age effects. This restriction on the cohort effects is implied by the model outlined in Subsection 3.4, where cohort effects capture variations in housing quality over time. So long as γ is nonlinear, the restriction provides one solution to the well-known problem of decomposing a variable into age, period, and cohort effect, as period is the sum of cohort and age (e.g., Hall et al., 2007; Hall, 1971). A number of different approaches have been taken in the hedonic pricing literature (e.g., Coulson and McMillen, 2008). Our approach is dictated by our goal of estimating κ_T and the theoretical framework in Subsection 3.4 which motivates that objective.

We set γ and α to be quadratic functions, and, as we are after only a single number for κ_T, set δ as a constant. Nonlinearity is essential, as δ is unidentified if γ is linear. Thus we estimate,

$$y_{its} = \gamma_1 t + \gamma_2 t^2 + \delta(\gamma_1 s + \gamma_2 s^2) + \alpha_1(t-s) + \alpha_2(t-s)^2.$$

A consistent estimate for δ can be obtained by regressing log price on the period of transaction and its square, the square of the period of construction, age (or period of construction) and age-squared. The estimate $\widehat{\delta}$ is the ratio of the coefficient on the square of the period of construction to the coefficient on the square of the period of transaction. Column (1) in Table 4 shows the results of the regression, with parcel fixed effects and the same set of building and apartment attributes as in Table 5 of Appendix A.4, and using the data described elsewhere in the paper but for all transactions with construction years the year after or up to 40 years before the transaction year.

Table 4: Existing Homes Price Regression

Variable	(1)	(2)
Year of Transaction	-0.034	-0.033
	(0.001)	(0.001)
Year of Transaction Squared/100	0.311	0.310
	(0.002)	(0.002)
Year of Construction Squared/100	0.0005	-
	(0.001)	(-)
Age	0.0012	0.0012
	(0.0002)	(0.0002)
Age-Squared/100	-0.0036	-0.0033
	(0.0007)	(0.0007)

Notes: The dependent variable is in prices per square meter in real 2017 NIS. Year is calendar year minus 1997. The number of observations is 776,709.

We estimate $\widehat{\delta} = 0.0005/0.311 = 0.0016$ ($s.e. = 0.0018$), and so $\widehat{\kappa}_T = \widehat{\delta}/(1+\widehat{\delta}) = 0.0016$ (s.e. = 0.0018), indicating that housing quality barely varies with price over time. We obtain similar results for γ and α quartic functions.

Column (2) in Table 4 drops the squared year of construction, substituting instead its interaction with indicator functions for the twenty largest (by number of transactions) cities and an indicator for all other cities. This allows the relationship between period effects and cohort effects to vary across locations. The results are very similar. No city shows an absolute ratio exceeding 0.0460, while the ratio of the weighted mean of the interaction coefficients to the square of the transaction year (with weights equal to the frequency of the cities and the residual category in the regression sample) is -0.0037 (s.e. = 0.0019).

A.2 Variances

Conditioning on height, we estimate the variances of u, v, and w using apartment, building, and bloc hierarchical modeling,

$$\widehat{\mathrm{Var}}(v) = \frac{1}{\sum_{k=1}^{K}\sum_{i=1}^{n_k}(J_{ki}-1)} \sum_{k=1}^{K}\sum_{i=1}^{n_k}\sum_{j=1}^{J_{ki}} (y_{kij}^0 - \bar{y}_{ki}^0)^2, \tag{18}$$

$$\widehat{\mathrm{Var}}(w) = \frac{1}{\sum_{k=1}^{K}(n_k-1)} \Big(\sum_{k=1}^{K}\sum_{i=1}^{n_k} (\bar{y}_{ki}^0 - \bar{y}_k^0)^2 - \widehat{\mathrm{Var}}(v) \sum_{k=1}^{K}\sum_{i=1}^{n_k} \frac{n_k-1}{n_k J_{ki}} \Big), \tag{19}$$

$$\widehat{\mathrm{Var}}(u) = \frac{1}{K-1}\sum_{k=1}^{K} \left(\bar{y}_k - \bar{y}\right)^2 - \frac{\widehat{\mathrm{Var}}(w)}{K}\sum_{k=1}^{K}\frac{1}{n_k} - \frac{\widehat{\mathrm{Var}}(v)}{K}\sum_{k=1}^{K}\sum_{i=1}^{n_k}\frac{1}{n_k^2 J_{ki}}, \tag{20}$$

where y^0_{kij} is the residual of a nonparametric series regression of log price on transaction date (in days), and where the estimated building prices are $\bar{y}^0_{ki} = \frac{1}{J_{ki}}\sum_{j=1}^{J_{ki}} y^0_{kij}$, $\bar{y}_{ki} = \frac{1}{J_{ki}}\sum_{j=1}^{J_{ki}} y_{kij}$, the estimated bloc prices are $\bar{y}^0_k = \frac{1}{n_k}\sum_{i=1}^{n_k}\bar{y}^0_{ki}$ and $\bar{y}_k = \frac{1}{n_k}\sum_{i=1}^{n_k}\bar{y}_{ki}$, and the overall average prices are $\bar{y}^0 = \frac{1}{K}\sum_{k=1}^{K}\bar{y}^0_k$ and $\bar{y} = \frac{1}{K}\sum_{k=1}^{K}\bar{y}_k$.

A.3 The Frontier

Fix height h. To simplify notation, drop the height index h. Since $u \sim TN(\mu_u, \sigma^2_u)$,

$$\mathrm{Var}(u) = \sigma^2_u\Big[1 - \frac{\mu_u}{\sigma_u}\cdot\lambda\Big(\frac{\mu_u}{\sigma_u}\Big) - \Big(\lambda\Big(\frac{\mu_u}{\sigma_u}\Big)\Big)^2\Big], \tag{21}$$

where $\lambda(x) = \phi(x)/\Phi(x)$, and $\phi(.)$ and $\Phi(.)$ are the standard normal probability and cumulative density functions. Combining (20) with (21) we obtain,

$$\widehat{\sigma}^2_u\Big[1 - \frac{\widehat{\mu}_u}{\widehat{\sigma}_u}\cdot\lambda\Big(\frac{\widehat{\mu}_u}{\widehat{\sigma}_u}\Big) - \Big(\lambda\Big(\frac{\widehat{\mu}_u}{\widehat{\sigma}_u}\Big)\Big)^2\Big] = \frac{1}{K-1}\sum_{k=1}^{K}(\bar{y}_k - \bar{y})^2 - \frac{\widehat{\sigma}^2_w}{K}\sum_{k=1}^{K}\frac{1}{n_k} - \frac{\widehat{\sigma}^2_v}{K}\sum_{k=1}^{K}\sum_{i=1}^{n_k}\frac{1}{n_k^2 J_{ki}}. \tag{22}$$

So that given the data and parameters $\widehat{\mu}_u$, $\widehat{\sigma}^2_v$, and $\widehat{\sigma}^2_w$, we obtain $\widehat{\sigma}^2_u$ using (22).

For each of M parameter values for (g, μ_u) and the estimates for σ^2_v and σ^2_w from (18)-(20) we obtain an estimate for σ^2_u and calculate the log likelihood (ignoring constants),

$$\begin{aligned}
\mathscr{L}_h(g, \mu_u, \sigma^2_u, \sigma^2_v, \sigma^2_w; \cdot) &= \frac{1}{2}\sum_{k=1}^{K}\Big(\frac{\mu_k^2}{\sigma_k^2} - \frac{\mu_u^2}{\sigma_u^2} + \frac{1}{\sigma^2_v}\sum_{i=1}^{n_k}\Big(\frac{\sigma^2_w(\sum_{j=1}^{J_{ki}}(y_{kij}-g))^2}{\sigma^2_v + J_{ki}\sigma^2_w} - \sum_{j=1}^{J_{ki}}(y_{kij}-g)^2\Big) + \\
&\ln\sigma^2_k - \ln\sigma^2_u - \sum_{i=1}^{n_k}\Big(\ln(\sigma^2_v + J_{ki}\sigma^2_w) + (J_{ki}-1)\ln\sigma^2_v\Big) + 2\ln\Phi\Big(\frac{\mu_k}{\sigma_k}\Big) - 2\ln\Phi\Big(\frac{\mu_u}{\sigma_u}\Big)\Big), \qquad (23)\\
\mu_k &= \frac{\sigma^2_k}{\sigma^2_u n_k}\sum_{i=1}^{n_k}\frac{\mu_u(\sigma^2_v + J_{ki}\sigma^2_w) + n_k\sigma^2_u\sum_{j=1}^{J_{ki}}(y_{kij}-g)}{\sigma^2_v + J_{ki}\sigma^2_w},\\
\sigma^2_k &= \sigma^2_u n_k\Big(\sum_{i=1}^{n_k}\frac{\sigma^2_v + J_{ki}\sigma^2_w + n_k J_{ki}\sigma^2_u}{\sigma^2_v + J_{ki}\sigma^2_w}\Big)^{-1},
\end{aligned}$$

where μ_k is a weighted average of μ_u and the average distance of log price to the frontier.

Now, the global maximum of the likelihood at height h is obtained by maximizing (23). The global maximum of the likelihood, constrained so that average costs decrease to MES and marginal costs increase thereafter, is attained by a grid search and Dijkstra's algorithm,

$$\{\widehat{MES}, \widehat{g}, \widehat{\mu}_u\} = \underset{\substack{mes\in\{1,\ldots,H-1\}\\ g\in\mathbb{R}^H, \nu_u\in\mathbb{R}^H}}{\operatorname{argmax}} \sum_{h=1}^{H}\mathscr{L}_h(g_h, \nu_{uh}, \cdot),$$

$$\text{s.t. } g_{mes} \le g_{mes-1} \le \ldots \le g_1 \text{ and } g_{mes} \le g_{mes+1} \le \ldots \le g_H.$$

Now we describe how to obtain a smooth ML estimator for a fourth order polynomial cost function, defined on a domain of continuous quantities, which we write as

$$C(h(q)) = \beta_0 + \beta_1 q + \beta_2 q^2 + \beta_3 q^3 + \beta_4 q^4,$$

implying marginal and average cost functions

$$MC(h(q)) = \beta_1 + 2\beta_2 q + 3\beta_3 q^2 + 4\beta_4 q^3 \text{ and } AC(h(q)) = \frac{1}{q}\beta_0 + \beta_1 + \beta_2 q + \beta_3 q^2 + \beta_4 q^3.$$

So $G(h) = \max\{AC(h), MC(h)\}$. The smooth estimator maximizes the likelihood,

$$\{\widehat{MES}, \widehat{\beta}, \widehat{\mu}_u\} = \underset{\substack{mes\in\{1,\ldots,H-1\} \\ b\in\mathbb{R}^5, \nu_u\in\mathbb{R}^H}}{\operatorname{argmax}} \sum_{h=1}^{H} \mathscr{L}_h(\cdot) \tag{24}$$

$$\text{s.t. } MC(mes-1) \leq AC(mes-1) \leq \ldots \leq AC(1), \tag{25}$$

$$AC(mes) \leq MC(mes) \leq \ldots \leq MC(H). \tag{26}$$

We now derive the likelihood in (23). Assume $v_{kij} \sim N(0,\sigma_v^2)$, $w_{ki} \sim N(0,\sigma_w^2)$, and $u_k \sim TN(\mu_u, \sigma_u^2)$. So,

$$f_{v_{kij}}(v) = \frac{e^{-v^2/2\sigma_v^2}}{\sqrt{2\pi\sigma_v^2}}, \qquad f_{w_{ki}}(w) = \frac{e^{-w^2/2\sigma_w^2}}{\sqrt{2\pi\sigma_w^2}}, \qquad f_{u_k}(u) = \frac{e^{-(u-\mu_u)^2/2\sigma_u^2}}{\sqrt{2\pi\sigma_u^2}\cdot\Phi(\mu_u/\sigma_u)}, \qquad u \geq 0.$$

By independence of $u_k, w_{k1}, \ldots, w_{kn_k}, v_{k11}, \ldots, v_{k1J_{k1}}, \ldots, v_{kn_k1}, \ldots, v_{kn_kJ_{kn_k}}$,

$$\begin{aligned}
& f_{u_k+w_{k1}+v_{k11},\ldots,u_k+w_{k1}+v_{k1J_{k1}},\ldots,u_k+w_{kn_k}+v_{kn_k1},\ldots,u_k+w_{kn_k}+v_{kn_kJ_{kn_k}}}(s_{11},\ldots,s_{1J_{k1}},\ldots,s_{n_k1},\ldots,s_{n_kJ_{kn_k}}) \\
&= \int_0^\infty \int_{-\infty}^\infty \cdots \int_{-\infty}^\infty f_{u_k}(u) \prod_{i=1}^{n_k} \Big(f_{w_{ki}}(w_i) \prod_{j=1}^{J_{ki}} f_{v_{kij}}(s_{ij} - w_i - u) dw_i \Big) du \\
&= \int_0^\infty \frac{e^{-(u-\mu_u)^2/2\sigma_u^2}}{\sqrt{2\pi\sigma_u^2}\cdot\Phi(\mu_u/\sigma_u)} \prod_{i=1}^{n_k} \Big(\int_{-\infty}^\infty \frac{e^{-w_i^2/2\sigma_w^2}}{\sqrt{2\pi\sigma_w^2}} \prod_{j=1}^{J_{ki}} \frac{e^{-(s_{ij}-w_i-u)^2/2\sigma_v^2}}{\sqrt{2\pi\sigma_v^2}} dw_i \Big) du \\
&= \frac{\sigma_k \exp\Big(\sum_{i=1}^{n_k} \frac{\sigma_w^2 (\sum_{j=1}^{J_{ki}} s_{ij})^2}{2(\sigma_v^2 + J_{ki}\sigma_w^2)\sigma_v^2} - \frac{\mu_u^2}{2\sigma_u^2} - \sum_{i=1}^{n_k} \frac{\sum_{j=1}^{J_{ki}} s_{ij}^2}{2\sigma_v^2} + \frac{\mu_k^2}{2\sigma_k^2} \Big) \Phi(\mu_k/\sigma_k)}{(2\pi)^{\frac{1}{2}\sum_{i=1}^{n_k} J_{ki}} \sigma_u \Phi(\mu_u/\sigma_u) \sigma_v^{\sum_{i=1}^{n_k}(J_{ki}-1)} \prod_{i=1}^{n_k} \sqrt{\sigma_v^2 + J_{ki}\sigma_w^2}},
\end{aligned}$$

where

$$\begin{aligned}
\mu_k &= \frac{\sigma_k^2}{\sigma_u^2 n_k} \Big(\sum_{i=1}^{n_k} \frac{\mu_u(\sigma_v^2 + J_{ki}\sigma_w^2) + n_k\sigma_u^2 \sum_{j=1}^{J_{ki}} s_{ij}}{\sigma_v^2 + J_{ki}\sigma_w^2} \Big), \\
\sigma_k^2 &= \sigma_u^2 n_k \Big(\sum_{i=1}^{n_k} \frac{\sigma_v^2 + J_{ki}\sigma_w^2 + n_k J_{ki}\sigma_u^2}{\sigma_v^2 + J_{ki}\sigma_w^2} \Big)^{-1}.
\end{aligned}$$

We show $u|u+\eta$ is truncated normal in (15). Assume $u \sim TN(\mu_u, \sigma_u^2)$ and $\eta \sim N(0, \sigma_\eta^2)$.

$$f_{u,u+\eta}(u,s) = \frac{e^{-(u-\mu_u)^2/2\sigma_u^2} e^{-(s-u)^2/2\sigma_\eta^2}}{2\pi\sigma_u\sigma_\eta \cdot \Phi(\mu_u/\sigma_u)},$$

$$f_{u+\eta}(s) = \int_0^\infty f_u(u) f_\eta(s-u)du = \frac{\sigma_* \exp\left(\frac{\mu_*^2}{2\sigma_*^2} - \frac{\mu_u^2}{2\sigma_u^2} - \frac{s^2}{2\sigma_\eta^2}\right)\Phi(\mu_*/\sigma_*)}{\sqrt{2\pi}\sigma_u\sigma_\eta \cdot \Phi(\mu_u/\sigma_u)},$$

$$f_{u|u+\eta}(u|s) = \frac{\exp\left(-\frac{(u-\mu_u)^2}{2\sigma_u^2} - \frac{(s-u)^2}{2\sigma_\eta^2} - \frac{\mu_*^2}{2\sigma_*^2} + \frac{\mu_u^2}{2\sigma_u^2} + \frac{s^2}{2\sigma_\eta^2}\right)}{\sqrt{2\pi}\sigma_*\Phi(\mu_*/\sigma_*)} = \frac{\exp\left(-\frac{1}{2\sigma_*^2}(u-\mu_*)^2\right)}{\sqrt{2\pi}\sigma_*\Phi(\mu_*/\sigma_*)},$$

where $\mu_* = \dfrac{\sigma_\eta^2\mu_u + s\sigma_u^2}{\sigma_u^2 + \sigma_\eta^2}$ and $\sigma_*^2 = \dfrac{\sigma_u^2\sigma_\eta^2}{\sigma_u^2 + \sigma_\eta^2}$.

A.4 Apartment-Floor, Building-Height Adjusted Prices

To obtain the adjusted prices, we begin by regressing the real, cost adjusted, per square meter log price on a full set of floor and building height interactions, dummy variables for transaction year before and transaction year after the year of construction, a nine-degree polynomial in the calendar day of transaction, eight dummies for the legal status of the property, and dummy variables for the building. Identification of the floor effects is possible because of cases in which there are multiple apartments in the same building, but on different floors. Identification of the height effects is possible because of cases in which there are multiple buildings on the same land parcel.[34]

A selected set of the estimates for the floor $\times$ height interactions in buildings with 5 to 10 floors are shown in Figure 14a. For given building height, the relationship between price and floor is J-shaped and right-leaning, with price falling initially, reflecting an initial preference for the ground floor and then more or less linearly increasing, until a penthouse effect at the penultimate and top floor. There is also a building height effect, with shorter buildings preferred to taller ones, especially at higher floors. Figure 14b covers a wider range of heights, grouping each 5 floor range of heights, and shows similar results.

On the basis of these estimates, we choose to model the conditioning on floor and height by a linear term in floor, dummy variables for each of the ground, first, second, and third floors, a linear term in building height, and dummies for the penultimate and top floors, as well as interaction with the sum of those two dummies and the building height. There are also interactions between

[34] These are a small fraction of the data, but of sufficient number that the height effects can be measured.

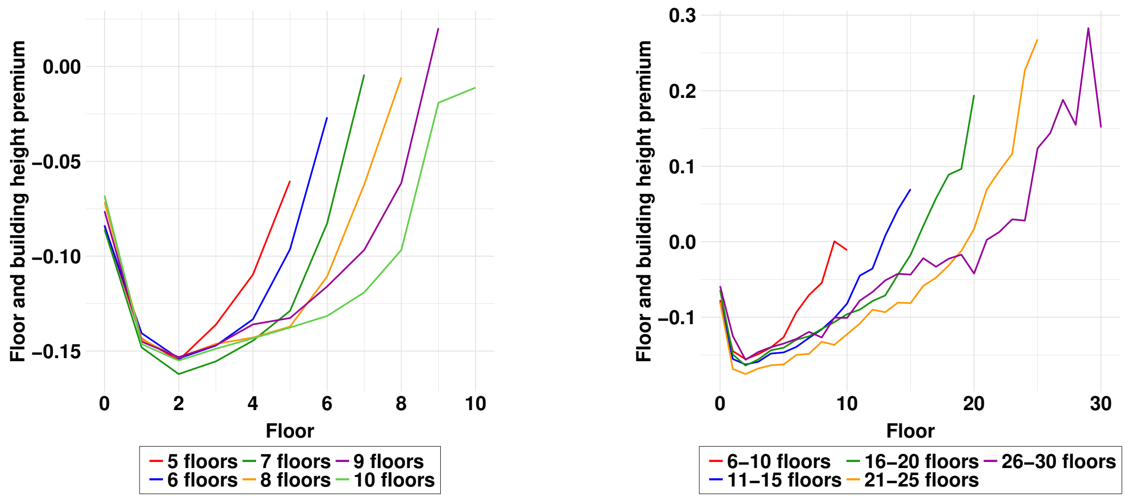

(a) 5-10 floor and building height effects.

(b) 5-30 floor and building height effects.

Figure 14: Floor and building height effects

a dummy for above four floors with the first, second, and third floor dummies, and interactions between heights above 10 floors and the linear term in floor.[35] Table 5 presents the coefficients and standard errors of the main variables.

Table 5: Preliminary stage regression

	Log price
Floor	0.0088
	(0.0003)
Building height	-0.0006
	(0.0001)
Penthouse	0.0361
	(0.0016)
Penthouse - 1	0.0058
	(0.0017)
Penthouse × Building height	0.0027
	(0.0002)
Year before construction year	-0.0037
	(0.0009)
Year after construction year	0.0030
	(0.0007)

Notes: Standard errors are in parentheses. Additional controls: polynomial in calendar time, ground, first, second, and third floor dummies and their interactions with dummies for building heights above 4 and 10 floors, eight legal status dummies, and parcel fixed effects.

[35] These two cutoffs originate in the minimal regulatory requirements for a first and a second elevator.

B The Frontier Elasticity of Substitution of Land for Capital

The elasticity of substitution of the housing production function is the rate at which the cost-minimizing capital to land ratio varies with the marginal rate of technical substitution. This is commonly used to summarize the degree of substitution of one input for the other in housing production. With price-taking firms in input markets, and normalizing the price of capital to 1, the elasticity of substitution is $\sigma = \frac{d\ln k}{d\ln R}$, where k is capital per unit of land, and R is the price of a unit of land (i.e., land rent).

Given price taking firms in the input market, and normalizing the price of capital to 1, the elasticity of substitution is,

$$\sigma = \frac{d\ln k}{d\ln R} = \frac{R}{k} \times \frac{dk}{dR},$$

where $k = K/L$ is the capital to land ratio (or the capital per unit of land), K is capital, L is a given fixed amount of land, and R is the price of one unit of land, i.e., land rent.

With the constant returns to scale production function in land and capital $f_0(K,L)$, per unit of land housing output, equivalently height h, satisfies $h = f_0(K,L)/L = f_0(K/L,1) = f(k)$. Noting that $k = C(h)$, $h = C^{-1}(k) = f(k)$, $C'(h) = 1/f'(k)$, and $C''(h) = -f''(k)/(f'(k))^3$, the elasticity of substitution is,

$$\sigma = \frac{f'(k)(kf'(k) - f(k))}{kf(k)f''(k)} = \frac{C'(h)(hC'(h) - C(h))}{hC(h)C''(h)} = \frac{\overbrace{(MC - AC) \times h}^{R}}{\underbrace{h \times AC}_{k}} \times \frac{\overbrace{MC \times dh}^{dk}}{\underbrace{h \times dMC}_{dR}} = \frac{d\ln AC}{d\ln MC},$$

where the first equality follows from Arrow et al. (1961).

Since in an unregulated market, housing price equals marginal non-land cost, this is also the elasticity of average non-land cost to market price. Furthermore, since price equals total average cost (the long run, zero profit condition) the elasticity of substitution relates the growth of land rent to the growth of non-land costs as height increases.

C Additional Figures and Tables

C.1 Robustness of the ML Estimates

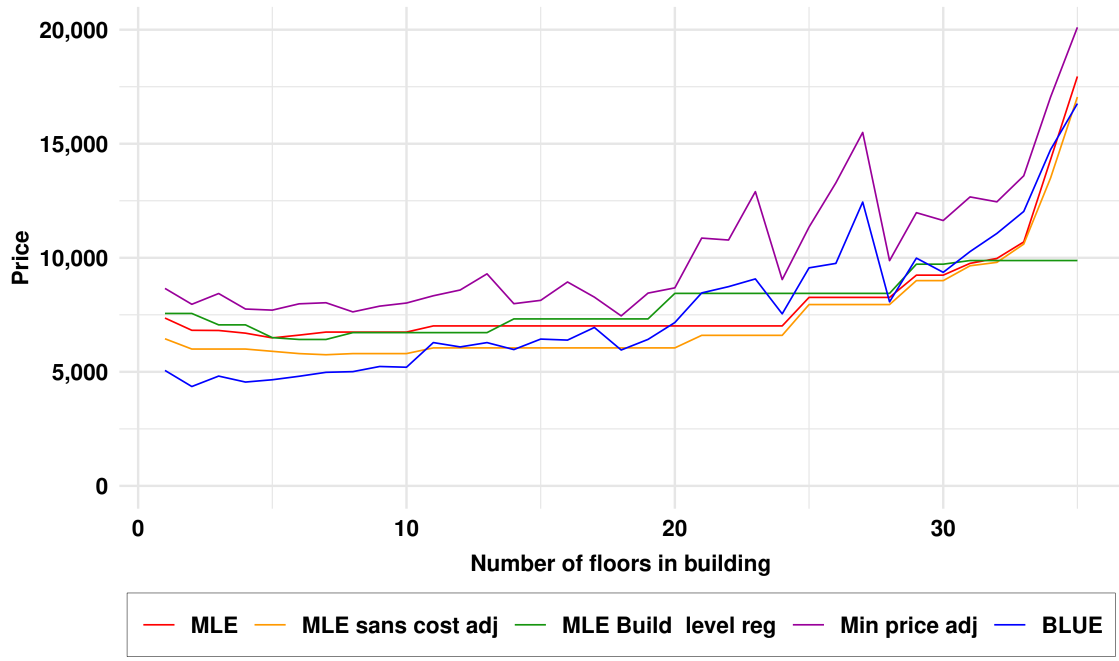

Figure 15: Robustness of ML estimate to ML estimate without adjusting for changes in cost over time, ML estimate using building-level regulation, minimum price at each height adjusted for sample size and error, and BLUE.

C.2 Prices in Cities by Geographical Coordinates

Figures 16a-16c show the heat maps of the estimated prices (using nonparametric local constant regression with bandwidth chosen by cross validation) for the three largest cities - Jerusalem, Tel Aviv, and Haifa.

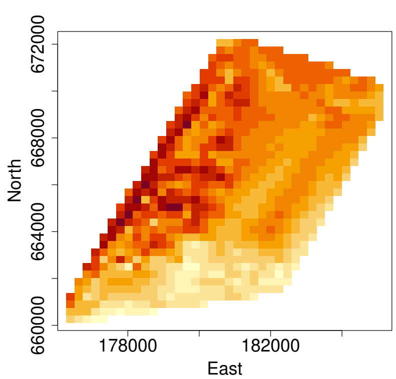

(a) Tel Aviv.

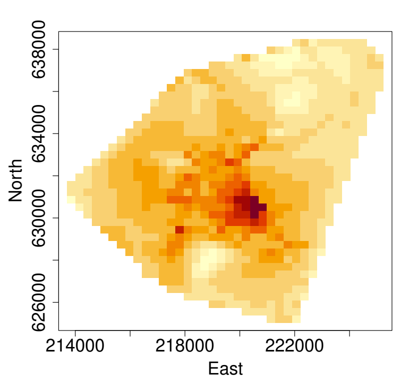

(b) Jerusalem.

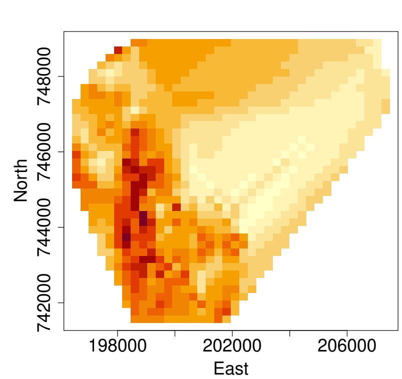

(c) Haifa.

Figure 16: Heat map of prices in the cities Jerusalem, Tel Aviv, and Haifa.

C.3 Local Concentration

As noted in the main part of the paper, the construction industry in Israel is structurally competitive, with a ten-firm national concentration ratio of 0.15 only (Ministry of Finance, Chief Economist Branch, 2017). That the largest firms are known to operate throughout the country, and the country is geographically small, suggests low local concentration as well. Lacking information on the builder's identity in the transaction dataset that is the main source for our analysis, we turn to auctions for construction rights on government owned land for a quantitative statement on local concentration. The auctions, held from 1998 to 2017 for the approximately fifty percent of construction that takes place on government owned land (Rubin and Felsenstein, 2017), can not be reliably matched to our transaction data as their finest, reliable geographical identifier in those data is at the locality - i.e., city, town or regional council - level. Yet they can provide us with a sense of local competition. We calculate a mean locality-level HHI over all zoned-for dwelling units in the auctions of 0.025, equivalent to forty equally-sized firms.

As a rule, larger markets support more competitors and so are characterized by smaller markups Sutton (1991) and less firm heterogeneity, as the more inefficient firms are priced out of the market (Syverson, 2004). This suggests that any bias in the measured regulatory tax arising from varying markups or differential firm efficiency decrease with market size. Thus the extent of any contribution of varying markups and differential firm efficiency to the measured regulatory tax can be gauged by how regulatory tax varies empirically with measures of market size. We consider two measures. The first is population density, used as a proxy for the stock of existing homes in the vicinity of the given apartment or building. Existing homes, whether renovated or not, are substitutes for newly constructed housing, and so both limit markup and, pushing down price, limit inefficient firms. The second measure is the number of buildings constructed in the vicinity of the apartment or building over the period of our sample.

Table 6 shows how the regulatory tax varies with both population density and new building construction. As in our previous results, the regulatory tax rate is always positively associated with population density and so with the stock of existing homes. Furthermore, the regulatory tax rate is also positively associated with new building construction within a 1km radius. Only when we extend the radius to 4km, which in most cases is large enough to encompass a locality, and condition on fixed effect and distance to the city center, do we find a negative relationship between construction and the regulatory tax rate. Even in that case, the coefficient is very small: 1000 more buildings is associated with three percentage points lower regulatory tax. One thousand buildings is essentially the ninety percentile of constructed buildings within a 4km radius. The maximum is 1612, which would be associated with a five percentage points decrease. Such numbers are small with respect to the mean regulatory tax rate of 48 percent that we estimate.

Table 6: Regressions

	Estimated regulatory tax rate			
	(1)	(2)	(3)	(4)
		Apartment		
New building construction (1km radius)	0.0003 (0.00001)	-	0.00003 (0.00001)	-
New building construction (4km radius)	-	0.00009 (0.00001)	-	-0.00006 (0.000003)
Distance to city center	-	-	-0.004 (0.0002)	-0.005 (0.0002)
Density (1km radius)	0.067 (0.0008)	-	0.006 (0.0008)	-
Density (4km radius)	-	0.185 (0.002)	-	0.110 (0.003)
City fixed effects	No	No	Yes	Yes
R^2	0.094	0.246	0.564	0.566
		Building		
New building construction (1km radius)	0.0002 (0.00003)	-	0.00008 (0.00002)	-
New building construction (4km radius)	-	0.00008 (0.00001)	-	-0.00003 (0.00001)
Distance to city center	-	-	-0.006 (0.0007)	-0.008 (0.0007)
Density (1km radius)	0.09 (0.003)	-	0.005 (0.003)	-
Density (4km radius)	-	0.236 (0.006)	-	0.117 (0.011)
City fixed effects	No	No	Yes	Yes
R^2	0.127	0.317	0.670	0.672

Notes: Standard errors are in parentheses underneath the coefficients. Distance to city center is in kilometers. Densities are in 10,000's per square kilometer. There ares 13,102 buildings and 206,835 apartments. The top section has outcomes at the apartment level, and standard errors clustered at the building level. The bottom section has outcomes as the average apartment price in a building.

C.4 Maximum Likelihood Estimates

Table 7 shows heights, estimated quantities, the constrained ML estimates (MLE), ML estimates by height (MLE by height), and the minimum and mean building prices. The equation for the smooth ML estimates appears below the table.

The estimated quartic cost function is $\widehat{C}(q) = 900 + 6472q + 78.43q^2 - 4.1q^3 + 0.0823q^4$.

Table 7: Maximum likelihood estimates

Height	Quantity	MLE	MLE by height	Minimum	Mean
1	1.05	7359	7359	5666	10544
2	2.07	6822	6822	5052	11541
3	3.09	6814	6814	5354	12466
4	4.09	6696	6696	5385	12757
5	5.03	6660	6660	5374	14288
6	6.05	6660	6660	5256	14684
7	7.07	6744	6786	5842	14347
8	8.1	6744	6866	5319	14069
9	9.14	6744	6714	5705	14007
10	10.19	6744	6660	5605	14282
11	11.18	7013	7405	6576	15555
12	12.23	7013	7010	6777	15839
13	13.28	7013	7316	6560	15455
14	14.35	7013	6660	6078	14000
15	15.42	7013	7829	6503	14568
16	16.5	7013	6660	6410	15252
17	17.58	7013	6966	7103	15491
18	18.68	7013	6660	5943	15074
19	19.78	7013	6777	6940	15228
20	20.89	7013	6789	7156	15221
21	22.00	7013	8891	8901	16847
22	23.13	7013	7686	8753	16515
23	24.26	7013	9214	8919	15569
24	25.4	7013	6708	7433	18155
25	26.54	8264	9621	9591	16903
26	27.69	8264	10418	11015	14778
27	28.86	8264	10742	12820	19637
28	30.03	8264	7942	8479	18088
29	31.2	9239	9716	10157	19635
30	32.39	9239	8878	9637	21399
31	33.58	9757	9757	10742	24481
32	34.78	9972	9972	11033	21124
33	35.99	10695	10695	11792	21729
34	37.21	14307	14307	14865	23078
35	38.41	17950	17950	17805	23500

C.5 Number of Observations by Height

Table 8 shows summary statistics for the number of observations by height. The second, fifth, and sixth columns show the number of blocs, buildings, and apartments respectively. The number of observations in each of these column trends downward with height. The third column is the percentage of blocs from column two that contain exactly one building (of a given height) and the fourth column is the mean number of buildings of the same height in these bloc. Given height, in these blocs the median number of buildings is one and the average is about 2.4.

Table 8: Number of observations

Height	Blocs	% of blocs with one building	Mean # of buildings per bloc	Buildings	Apartments
1	182	0.74	1.8	319	1453
2	629	0.53	2.6	1661	8068
3	606	0.57	2.3	1394	10310
4	874	0.45	3.4	2968	28266
5	866	0.47	3.0	2562	27642
6	826	0.49	2.8	2315	27336
7	663	0.51	2.5	1639	24725
8	572	0.53	2.3	1340	24086
9	472	0.52	2.4	1137	24384
10	341	0.55	2.0	674	15682
11	202	0.68	1.6	331	9214
12	155	0.64	1.6	253	7517
13	154	0.76	1.3	207	7303
14	121	0.69	1.7	202	6369
15	112	0.66	1.7	185	7434
16	93	0.68	1.5	142	6024
17	80	0.62	1.8	145	6825
18	76	0.71	1.6	122	4060
19	61	0.66	1.6	97	3407
20	62	0.73	1.5	90	3744
21	49	0.71	1.4	67	3894
22	42	0.69	1.6	69	2373
23	25	0.68	1.6	40	1623
24	36	0.78	1.2	45	1930
25	21	0.95	1.0	22	1252
26	18	0.78	1.4	26	902
27	12	0.83	1.2	14	766
28	15	0.67	1.4	21	925
29	14	0.71	1.4	19	730
30	14	0.86	1.1	16	659
31	7	0.71	1.3	9	309
32	7	1.00	1.0	7	205
33	5	0.80	1.2	6	267
34	6	0.83	1.3	8	267
35	11	0.64	1.5	17	603