

Analyzing Building-Height Restrictions: Predicted Impacts, Welfare Costs, and a Case Study of Bangalore, India

by

Alain Bertaud
Consultant, World Bank
166 Forest Rock Road
Glen Rock, NJ 07452
e-mail: duatreb@msn.com

and

Jan K. Brueckner
Department of Economics
and
Institute of Government and Public Affairs
University of Illinois at Urbana-Champaign
Champaign, IL 61820
e-mail: jbrueckn@uiuc.edu

April 2003

Abstract

This paper analyzes the impacts and the welfare cost of building-height restrictions, providing a concrete welfare-cost estimate for the city of Bangalore. Relying on several theoretical results, the analysis shows that the welfare cost imposed on its residents by Bangalore's FAR restriction ranges between 3 and 6 percent of household consumption. This burden represents a significant share of individual resources, and its presence may push many marginal households into poverty.

Analyzing Building-Height Restrictions: Predicted Impacts, Welfare Costs, and a Case Study of Bangalore, India

by

Alain Bertaud and Jan K. Brueckner*

1. Introduction

Throughout the world, land-use decisions are affected by various forms of government intervention in real estate markets. Such interventions include U.S.-style zoning regulations, which are meant to minimize externalities by separating different types of land uses, as well as greenbelt laws and urban growth boundaries, which limit the spatial expansion of cities (see Fischel (1990) for an overview of such policies). Sometimes, government intervention takes more drastic forms. For example, apartheid policies in South Africa dictated residential location patterns by race, pushing black households to remote locations far from the urban centers (see Brueckner (1996)). In the command economies of the old Soviet bloc, rigid government control sometimes generated perverse land-use patterns, exemplified by Moscow's inverted density contour, which put people far from their places of work (see Bertaud and Renaud (1997)).

Governments in most market economies also exert explicit control over the density of development. While minimum-lot-size rules and other regulations are designed to limit suburban development densities in the U.S., a related regulatory tool is the building-height restriction, which governs land-use in the central parts of many cities worldwide. The most obvious examples are Washington, D.C., where no building can be taller than the U.S. Capitol, and Paris, where height restrictions attempt to preserve the character of the central city.

Although the use of height restrictions in both of these cities is driven by aesthetic considerations, building-height restrictions may also be imposed in an attempt to achieve other goals. A case in point is India, where city governments impose height limitations through restrictions on the "floor area ratio" (FAR) of buildings. The FAR is computed by dividing a building's total floor area by the area of land parcel on which it sits, and an upper limit on this ratio

effectively puts a limit on the building's height.

In controlling FAR, the goal of the Indian planners has been to indirectly limit both job and population densities. It is believed that "excessive" density results in a loss of environmental quality and increased traffic congestion. In addition, higher densities would place greater demands on urban infrastructure, which Indian cities, plagued by weak technical capacities and inadequate tax revenue, feel ill-equipped to provide at appropriate levels.

For example, while low water pressure and high leakage might be acceptable in a low-density neighborhood, the same pressure and leakage performance might create unacceptable living conditions in higher density areas. A similar point applies to waste disposal, where an Indian municipality may succeed at removing only 25 percent of the refuse generated by the city. While this outcome might be unpleasant but survivable in a low-density neighborhood, it may entail grave health consequences at higher densities.

Thus, accepting high densities in central areas would have required a commitment by Indian cities to improve the productivity and performance of municipal services, requiring substantial investments in infrastructure. Faced with this trade-off and limited by their poor technical capacities, Indian municipalities chose to reduce densities via FAR restrictions in order to avoid central-city infrastructure investments. In the last few years, however, the technical performance of many major municipal governments has greatly improved, but the policy of FAR restrictions has not been revised.

Motivated by the widespread use of FAR restrictions in India, the present paper analyzes the effects of such restrictions and attempts to appraise their costs. This task is carried out first on an abstract level, with the standard monocentric-city model used to show how FAR restrictions affect land use. The analysis, presented in section 2, reaches the natural conclusion that, by limiting population densities near the center, an FAR restriction causes the city to expand spatially. This spatial enlargement affects consumers both by raising their average commuting distance and by pushing up the housing prices they must pay. However, the analysis shows that the resulting per capita welfare loss can be measured simply by the increase in commuting cost for an individual living at the edge of the city, who is now more distant from the central workplace.

It should be noted that this welfare-cost measure abstracts from the potential effects of the FAR restriction on infrastructure costs. While the restriction may reduce costs in the center, as explained above, spatial expansion of the city also leads to higher costs by requiring a wider infrastructure network. Because the net change in costs is difficult to estimate, the analysis focuses purely on consumer losses from the FAR restriction, abstracting from infrastructure issues.

Armed with results of the theoretical analysis, the paper turns in section 3 to a case study of a particular Indian city, namely, Bangalore. Using detailed land-use data for this city, the analysis carries out a counterfactual exercise, asking what Bangalore would look like today had it developed in the absence of the FAR restriction. The theory shows how FAR values and dwelling sizes would change throughout the city following a hypothetical removal of the restriction, and these predicted changes can be used to spatially reallocate Bangalore's population, raising central densities and generating a shrinkage in the city's spatial size. The final step is to compute the resulting commuting-cost saving for an individual living at Bangalore's now-closer edge. This saving captures the consumer welfare gain from removal of the FAR restriction, and hence the welfare cost of its presence.

While such a counterfactual exercise is illuminating in its exposure of the costs of having followed the current regulatory policy, it does not provide Indian policy makers with some of the practical information they need. In deciding whether to ease the current FAR restrictions, these policy makers need to know the *future* impacts of such a change, not its retrospective effect. The conclusion of the paper discusses such future impacts, although the discussion is necessarily speculative.

Before proceeding, it is important to acknowledge the earlier work of Arnott and MacKinnon (1977), who carried out the only previous analysis of the effects of building-height restrictions, relying on a numerical version of the monocentric-city model. Their work is similar in spirit to the analysis in section 2, although it is carried out using a much more detailed and realistic specification of the housing production technology. Partly because of this greater realism, Arnott and MacKinnon are unable to present any general theoretical results regarding the impact of an FAR restriction, as is done below. The connections between Arnott and

MacKinnon's findings and those of the present analysis are noted as the discussion unfolds.

2. Theory

This section presents a theoretical analysis of the impact of building-height (FAR) restrictions. After starting with a review of the standard monocentric-city model, the discussion shows how the urban equilibrium is affected by the presence of an FAR restriction. A numerical simulation further illustrates the effects of the restriction, with the results summarized in a series of graphs. The analysis concludes by discussing measurement of the consumer welfare loss from the FAR restriction.

2.1. The monocentric-city model

The monocentric-city model was initially developed by Alonso (1964), Mills (1967), and Muth (1969), and Wheaton (1974) provided a more general analysis of the model's properties. The following discussion relies on Brueckner's (1987) extension of Wheaton's approach. Since the model has been explained many times in the literature, the present treatment is compressed. A more expansive review can be found in Brueckner (1987) and elsewhere.

The city is inhabited by N identical residents, all of whom work in the central business district (CBD), earning income y per period. The inhabitants commute from their residences to the CBD on a dense radial road network, paying t per round-trip mile per period. Letting x denote radial distance, the disposable income of a consumer living x miles from the CBD is then $y - tx$.

Individuals have the common, well-behaved utility function $v(c, q)$, which depends on consumption of a numeraire nonhousing good, denoted c , and housing square footage, denoted q . Everyone is a renter, and p denotes the rental price per square foot of floor space. The consumer budget constraint is then $c + pq = y - tx$. Eliminating c via the budget constraint, utility can be written $v(y - tx - pq, q)$.

Consumers maximize this utility expression by choice of q taking p as given. To insure locational equilibrium, this maximization must lead to a realized utility level that is uniform across space. In other words, it must be true that $\max_q v(y - tx - pq, q) = u$, where u is the uniform utility level. This requirement yields two conditions, a first-order condition for choice

of q and the condition that utility equals u . Together, these conditions determine q and p as functions of the parameters of interest, which in the present case are x and u . These solutions are given by

$$q(x, u), \quad p(x, u). \tag{1}$$

Spatial variation in p , as seen in (1), ensures locational indifference for consumers. To achieve this outcome, p must decline with x to offset the higher commuting costs incurred at greater distances. Thus, $p_x < 0$ holds, where the subscript denotes partial derivative. In response to this price decline, dwelling size increases with x , so that $q_x > 0$.

The utility level is ultimately endogenous, being determined by the equilibrium conditions discussed below. However, a parametric increase in u affects both p and q . The housing price falls as u rises, with $p_u < 0$. The lower price allows the consumer to reach the higher indifference curve with disposable income held constant. As long as housing is a normal good, this movement must involve an increase in q , with $q_u > 0$.

Housing floor space is produced by combining land and capital (building materials) under a constant-returns technology. Because of constant returns, housing output per unit of land can be written $h(S)$, where h is the intensive form of the production function and S is the capital-to-land ratio, referred to as “structural density.” Since $h(S)$ gives housing floor space per unit of land, it effectively represents the floor area ratio (FAR), as discussed in the introduction.

With the price of capital normalized to unity, profit per acre for the housing developer equals $ph(S) - S - r$, where r is rent per unit of land. Given p , the developer chooses S to maximize profit, satisfying the first-order condition $ph'(S) = 1$. Land rent is then determined by the zero-profit condition, which can be written $r = ph(S) - S$. Since S depends on p via the first-order condition, it ultimately depends on the determinants of p , x and u . Moreover, since r depends on p and S , it too depends on these same parameters. Thus, the solutions can be written

$$S(x, u), \quad r(x, u). \tag{2}$$

Since $p_x < 0$ holds and since the lower p depresses the incentive to develop the land, $S_x < 0$ also holds, indicating that structural density declines with x . Similarly, land rent also declines

with x , so that $r_x < 0$. Finally, with a higher u depressing p , both S and r fall as utility rises, yielding $S_u, r_u < 0$. For future reference, the effects of x and u on all the variables are summarized as follows:

$$q_x > 0, \quad p_x, S_x, r_x < 0; \quad q_u > 0, \quad p_u, S_u, r_u < 0. \quad (3)$$

The two urban equilibrium conditions determine the utility level u as well as the distance \bar{x} to the edge of the city. These equilibrium conditions require (i) that the urban land rent at \bar{x} equals the agricultural rent r_a , and (ii) that the city population N fits inside \bar{x} . Letting the equilibrium values of u and \bar{x} be denoted u_0 and \bar{x}_0 , the conditions determining these values can be written

$$r(\bar{x}_0, u_0) = r_a \quad (4)$$

$$\int_0^{\bar{x}_0} \theta x \frac{h(S(x, u_0))}{q(x, u_0)} dx = N, \quad (5)$$

where $\theta \leq 2\pi$ is the constant number of radians of land available for housing. To interpret (5), observe that $h(S(x, u_0))/q(x, u_0)$ equals housing floor space per unit of land divided by floor space per dwelling. This ratio thus gives dwellings per unit of land, which in turn equals population density if each dwelling contains one person.¹ Multiplying this density by the area of the ring at distance x , equal to $\theta x dx$, and integrating across rings out to \bar{x} then gives the population fitting in the city, which must equal N . In conjunction with (4), the resulting condition helps determine the equilibrium values of u_0 and \bar{x}_0 .

It should be noted that the rent earned by the land, both urban and agricultural, disappears from the model. This outcome reflects the standard assumption that absentee landowners, who live outside the city in question, receive the rental income. Models with resident landowners do exist, but they are more cumbersome to analyze (see Pines and Sadka (1986)).

2.2. Introducing an FAR restriction

Under an FAR restriction, the government imposes an upper limit on the square footage of housing that can be produced per unit of land. In the model, such a restriction can be written

$$h(S) \leq \hat{h}, \quad (6)$$

where \hat{h} is the FAR limit. Since S declines with x in an unrestricted city, as seen above, it follows that housing output per unit of land, $h(S)$, declines with x as well. As a result, the FAR restriction will bind in the central part of the city, where $h(S)$ would normally be high, being nonbinding farther from the center.

Let \hat{x} denote the value of x where the FAR restriction becomes binding. Then, letting u_1 and \bar{x}_1 denote the u and \bar{x} values for the FAR-restricted city, the equilibrium conditions are given by

$$r(\bar{x}_1, u_1) = r_a \quad (7)$$

$$h(S(\hat{x}, u_1)) = \hat{h} \quad (8)$$

$$\int_0^{\hat{x}} \theta x \frac{\hat{h}}{q(x, u_1)} dx + \int_{\hat{x}}^{\bar{x}_0} \theta x \frac{h(S(x, u_1))}{q(x, u_1)} dx = N. \quad (9)$$

To understand (8), note that the left-hand side gives the FAR value that would be freely chosen by developers at $x = \hat{x}$, which is based on the chosen structural density $S(\hat{x}, u_1)$. Thus, (8) says that the freely-chosen FAR at \hat{x} equals the restricted value \hat{h} , indicating that this location is where the FAR restriction becomes binding. To interpret (9), note that the second integral has the same form as (6) but a lower limit of \hat{x} instead of zero, and that the first integral embodies the fact that h must equal \hat{h} inside of \hat{x} . By forcing h to remain constant instead of rising as x falls over this range, the FAR restriction tends to reduce population density in the central part of the city.

Note finally that, since land rent follows the $r(x, u_1)$ function outside \hat{x} , (7) is the appropriate boundary condition involving r_a . However, this land rent function is not relevant inside

\hat{x} since S is not freely chosen over this range. Instead, using the zero-profit condition from above, r is given by $p(x, u_1)\hat{h} - \hat{S}$ inside \hat{x} , where \hat{S} is the structural density corresponding to \hat{h} (satisfying $h(\hat{S}) = \hat{h}$). This alternate land-rent formula does not, however, enter the equilibrium conditions (7)–(9).

These three equations determine the equilibrium values u_1 , \bar{x}_1 and \hat{x} for the FAR-restricted city. The key question concerns the comparison between these values and those for the unrestricted city, focusing on the differences in u and \bar{x} . Using the properties of the model established above, it can be shown that the FAR restriction hurts the urban residents and causes the city to spread out. In particular,

Proposition 1. *The restricted city takes up more space and has lower consumer welfare than the unrestricted city, with $\bar{x}_1 > \bar{x}_0$ and $u_1 < u_0$ holding.*

To establish this conclusion, suppose first that u and \bar{x} were the same in both cities, with $u_1 = u_0 \equiv u^*$ and $\bar{x}_0 = \bar{x}_1 \equiv \bar{x}^*$. Then, note that equality of u_1 and u_0 means that population density (h/q) is the same between \hat{x} and \bar{x}^* in the restricted and unrestricted cities. However, since $\hat{h} < h(S(x, u^*))$ holds for $x < \hat{x}$, it follows that the restricted city has lower density inside \hat{x} than does the unrestricted city over this same range of distance. With the lower densities in the center and the same density outside \hat{x} , it follows that the restricted city fits fewer people than the unrestricted city, which implies that it cannot accommodate the population N . A similar argument rules out the case where $u_1 > u_0$ and $\bar{x}_1 < \bar{x}_0$, establishing the Proposition.²

2.3. A numerical simulation

To further illustrate differences between the restricted and unrestricted cities, it is helpful to rely on a numerical simulation. The numerical example is based on U.S. parameter values, which can be chosen with some degree of confidence. However, the results apply qualitatively to any city, including the Bangalore case considered below.

The simulation is based on Cobb-Douglas forms for the utility and production functions. The number of households N for the city is set at 800,000, implying an effective population of over 2 million in the case where a household realistically contains 2.6 persons. Income per household is set at the 2000 U.S. census value of \$42,151. About two-thirds of the circular

area of the city is available for housing, with θ in (6) and (9) set equal to 4 (as opposed to 2π) radians. Other parameter values are discussed in the appendix.

The unrestricted city has an \bar{x}_0 value of 21.4 miles. In the city, $h(S)$ ranges from a value of around 17.5 at the CBD to 0.4 at \bar{x}_0 . To interpret these values, note that FAR, as represented by $h(S)$, is roughly equal to the number of floors in a building, in the case when the parcel contains no open space around the structure. On the other hand, if a building covers only half of a parcel's land area, then FAR is equal to roughly half of the number of floors. In the latter case, the above numbers suggest that buildings in the unrestricted city have more than 30 floors at the center and are single-storied at the edge of the city. This FAR pattern is shown in Figure 1.

For the restricted city, the upper bound on FAR is set at 3.75. Under this restriction, the city radius expands by 2.1 miles, reaching an \bar{x}_1 value of 23.5 miles. The value of \hat{x} , where the FAR restriction becomes binding, is 11.7 miles, about halfway to the edge of the city.

The resulting FAR pattern can be seen in Figure 1. While FAR is flat at 3.75 inside \hat{x} , note that FAR values are higher outside \hat{x} than in the unrestricted city. Formally, the reason for this outcome is that lower utility in the restricted city means that $h(S(x, u_1)) > h(S(x, u_0))$ holds over the x range where S is freely chosen (recall $S_u < 0$).³

To understand this increase in FAR intuitively, note that by effectively tightening housing supply, the FAR restriction raises the price per square foot of housing at each location, as seen in Figure 2. Formally, this rise is a consequence of the decline in u combined with $p_u < 0$, which yield $p(x, u_1) > p(x, u_0)$. With p higher everywhere in the restricted city, the incentive for development grows, leading to an increase in FAR wherever the restriction is not binding.

The increase in p also leads to smaller dwellings in the restricted city, as seen in Figure 3. Once again, in response to the tightening of housing supply caused by the FAR restriction, consumers cut back on their consumption of floor space. It should be noted that, because the units of measurement of housing consumption are arbitrary in the simulation, the values of q and p can be rescaled appropriately. The scaling is chosen to give plausible values at the CBD in the unrestricted city, with dwellings of about 500 square feet, renting at about \$8 per square foot per year.

Figure 4 shows the effect of the FAR restriction on land rent, with rent given as a value per square mile. Near the center, rent is higher in the unrestricted city, but this relationship is reversed beyond a distance of $x = 4$. To understand this pattern, note that the decline in utility raises rent in the restricted city outside \hat{x} , with $r(x, u_1) > r(x, u_0)$ holding given $r_u < 0$. But once x falls below \hat{x} , rent rises more slowly than in the unrestricted city, a consequence of the fixity of housing output and structural density (see the above rent formula). This phenomenon can be seen in Figure 4, where the restricted land rent function becomes more nearly linear around $x = 12$, mirroring the shape of the housing price curve in Figure 2. Because it increases less rapidly as x falls, restricted rent is eventually overtaken by rent in the unrestricted city, which rises above it.

Finally, Figure 5 shows the effect of the FAR restriction on population density (measured per square mile). Strictly speaking, this Figure gives dwelling density (dwellings per square mile), which equals population density only when household size is unity. Outside \hat{x} , the lower u in the restricted city implies higher density (recall $S_u < 0$ and $q_u > 0$). But, as in the case of land rent, the FAR restriction slows the increase in density once x falls below \hat{x} . Over this range, the only force raising density is the decline in dwelling size q , with the effect of a higher FAR ceasing to operate (this fact also accounts for the kink at \hat{x}). As a result, density in the unrestricted city eventually overtakes the restricted density as x falls, rising above it inside $x = 8$. Density at the CBD in the unrestricted city is about 15,000 dwellings per square mile, while central density is only 4300 in the restricted city.

With the exception of the crossing pattern of the land rent curves in Figure 4, all of the patterns in the Figures are completely general, being guaranteed to hold under any parameterization of the model. In the case of land rent, the pattern shown in Figure 4, where central rent is higher in the unrestricted city, is not assured.

2.4. The welfare cost of the FAR restriction

As seen above, imposition of the FAR restriction reduces the utility level of urban residents. It is helpful to compute a dollar measure of this welfare loss, and the subsequent analysis is directed toward this end.

The analysis focuses on the individual living at the edge of the city, who resides farther

from the CBD after the FAR restriction is imposed. In comparing the circumstances of this individual with and without the restriction, a crucial fact is useful: the edge resident pays the same price per unit of housing p regardless of the location of \bar{x} .⁴ Intuitively, the housing price at the edge of the city depends on the prices of land and capital at this location, which determine production cost. But these prices are independent of the actual value of \bar{x} , being equal to r_a and 1, respectively.

Since the housing price paid by the edge resident stays the same as the FAR restriction pushes \bar{x} outward, it follows that the only effect felt by this individual is an increase in commuting cost, which rises by $t(\bar{x}_1 - \bar{x}_0)$. Thus, a compensatory payment of this amount would be sufficient to offset the harmful effect of the FAR restriction. In other words, if the edge resident received a compensation of $t(\bar{x}_1 - \bar{x}_0)$, then his utility would remain at the original level of u_0 instead of falling to u_1 . Thus, $t(\bar{x}_1 - \bar{x}_0)$ represents the welfare cost of the FAR restriction for the edge resident.

The impact of the FAR restriction for the city's interior residents is more complex than for the edge resident, suggesting that the above welfare-cost measure may not be relevant. As seen in Figures 2 and 3, housing prices rise following imposition of the FAR restriction, and dwelling sizes fall in response. Consumption of the nonhousing good c also changes, and households relocate, moving on average farther from the center. However, since everyone in the city must be equally well-off in equilibrium, these complexities can be ignored. In particular, it can be shown that the welfare cost computed for the edge resident applies exactly to the city's interior residents, even though their circumstances change in a more complex fashion. The following result is relevant:

Proposition 2. *For each urban resident, the welfare cost of the FAR restriction equals $t(\bar{x}_1 - \bar{x}_0)$, the increase in commuting cost for the edge resident.*

A formal argument establishing this result is given in the appendix.

It should be noted that, while harming consumers, the FAR restriction also affects the welfare of absentee landowners by altering total land rent. Arnott and MacKinnon's (1977) simulations show that the FAR restriction raises total differential land rent (urban rent minus

r_a), and the same outcome occurs in the simulation of section 2.3. However, since the direction of the impact on absentee landlords cannot be established in general, the reverse outcome is possible in principle.

3. Applying the Analysis to Bangalore

As explained in the introduction, Bangalore has developed under a long-standing FAR restriction, which has limited building heights in the city center. This section attempts to measure the consumer welfare cost of Bangalore's FAR restriction using the theoretical results from section 2.

This exercise has two components. The first step is to predict what Bangalore would look like today had it developed in the absence of the FAR restriction. In this counterfactual exercise, particular interest lies in the reduction of Bangalore's spatial size, and the consequent reduction in \bar{x} . The second step is to use the implied \bar{x} reduction along with a measure of commuting cost per mile to find the commuting-cost savings for the edge resident. These savings measure the welfare gain from the absence of the FAR restriction, and hence the welfare cost of its presence. These two components are discussed sequentially in the next two sections.

3.1. Computing the change in Bangalore's spatial size

Reversing the numerical exercise discussed in section 2, removal of the FAR restriction generates three important impacts on land-use: FAR values rise at central locations, they fall at locations where the restriction is not binding, and dwelling sizes rise throughout the city. In order to predict the change in Bangalore's spatial size, the magnitudes of these three effects must be quantified. Unfortunately, this task involves an element of guesswork, although the numerical results from section 2.3 provide some guidance.

Before delving into the details of this task, it is useful to discuss the available land-use data for Bangalore, which pertain to 1990 (the appendix explains the origins of the data). The city is divided into rings, with each comprising a kilometer's worth of radius. Data for each ring include its total residential population, the average FAR value for its structures, the permitted FAR value, and the ring's built-up land area. In 1990, Bangalore's population extended out

to the twelfth ring, whose outer radius is 12 km. The city’s total 1990 population was nearly 3.4 million.

Figure 6 shows FAR values for the 12 rings, with the horizontal scale starting at ring 1. The dotted curve shows the permitted FAR value, which is roughly constant across rings, and the solid dark curve shows the actual value. The rough correspondence of the two curves over their initial ranges suggests that the FAR restriction is binding in rings 1 through 4, and perhaps in ring 5 as well. The simulation results in Figure 1 predict that FAR will rise over most of the binding range when the restriction is removed, falling at the edge of this range (as happens near $x = 12$ in Figure 1).

While FAR values are thus predicted to increase in rings 1–4, the counterfactual equilibrium values are, of course, unknown. These values are chosen so that the unrestricted FAR curve has a plausible appearance, roughly matching the pattern in Figure 1. FAR values in rings 1–4 are set at 4.2, 3.1, 2.3, and 1.6, yielding the smooth unrestricted curve in Figure 6 (the lighter solid curve). Consistent with the pattern in Figure 1, the FAR curve is shifted downward over the range where the restriction is not binding, as seen in Figure 6. The downward shift is set at 7 percent.

Dwelling sizes rise when the FAR restriction is removed, as seen in Figure 3, and the magnitude of change is theoretically linked to the FAR decline shown in Figure 6.⁵ In the simulation of section 2.3, parameter choices led to percentage changes that are almost exactly equal in magnitude, with FAR falling by about 35 percent over the nonbinding range and q rising by about 35 percent throughout the city. For simplicity this equality in percentage changes is preserved in the calculations for Bangalore. Thus, dwelling sizes are assumed to increase by 7 percent in each ring, matching the 7 percent decline in FAR values over the nonbinding range. These are changes considerably smaller than those in the simulation, and the origin of this difference is explained below.

Recalling that population density is equal to $h(S)/q$, with h represented by FAR, the proportional change in density is given by the proportional change in the FAR/ q ratio. This change can be used to compute the change in ring populations following removal of the FAR restriction. Assuming that the residential land area in a ring stays constant, the new ring

population is found by multiplying the original population by the proportional change in the FAR/ q ratio, which represents the change in density. This proportional change is computed using the FAR values in Figure 6 along with the assumed 7 percent city-wide increase in q . Note that the original level of q is not required for this calculation.

This exercise is carried out using data on the existing ring populations for Bangalore, which are shown by the dark columns in Figure 7. Applying the proportional adjustment from above yields the new ring populations, which are shown by the lighter columns in the Figure. Populations in the last few rings in the city involve a further adjustment, which is explained momentarily.

Ring populations rise near the center of the city, falling farther from the center, and this pattern is easily understood. The substantial rise in the FAR values in rings 1–4 dominates the slight rise in q , leading to an increase in FAR/ q and thus in density. These density changes, which mirror the simulation results shown in Figure 5, lead to higher populations in rings 1–4. In the city’s outer rings, the decline in FAR and the rise in q reduce densities, again mirroring Figure 5, and ring populations fall. However, the population gain near the center dominates the declines farther out, so that the city’s land area shrinks. While the edge resident lives in ring 12 in the existing city, the edge resident in the restricted city lives in ring 9. Thus, the edge resident is 3 km closer to the center following removal of the FAR restriction. This predicted reduction in \bar{x} provides the key to computing the welfare cost of the FAR restriction, a task that is carried out in the next section.

A population adjustment is needed in the outer rings, as mentioned above, because the growth pattern of real-world cities differs from the idealized model. Instead of filling in a uniform circle as population grows, development at the outer edges of cities is “ragged,” with growth proceeding outward along major road arteries but lagging in other areas. The resulting city looks more like a star than a circle, with the points of the star lying on key transportation corridors.

This pattern can be seen in Figure 8, which shows the square kilometers of built-up area contained within each ring. Built-up area increases proportionally with distance through ring 7, reflecting development of all but a pie slice of the available area. Beyond this distance,

however, built-up area falls even though the total available ring area continues to increase. In effect, over this range of distance, development moves outward along the points of the urban star, which become increasingly narrow as the city's outer limit is approached.

Given this pattern, it would be illegitimate to compress Bangalore's population into the inner portion of the existing built-up area as the city centralizes. In doing so, the points of the urban star would effectively be truncated, with the population compressed into a circle that lacks the ragged, star-shaped periphery of real cities.

To preserve the ragged development pattern, the following approach is taken. Once the new ring populations are computed using the method described above, the city extends out to ring 8. Since the resulting population is unrealistically compressed, inconsistent with ragged development at the city's edge, the pattern is adjusted by arbitrarily moving a portion of the population in the two outer rings (7 and 8) into the vacant ring 9 lying just outside. This reallocation is carried out so that population smoothly declines across rings 7–9.

As part of this adjustment, built-up areas are appropriately reduced in rings 7 and 8, which lose population, and built-up area is allocated in the new outer ring 9. As can be seen in the lighter curve from Figure 8, this adjustment preserves the city's hump-shaped pattern of built-up areas. Without the adjustment, built-up areas would follow the dark curve, which would be truncated near its peak, thus lacking the realistic downward-sloping range indicative of ragged development at the city's edge.

While affected by this adjustment, the population pattern in the unrestricted city is obviously also sensitive to the assumed changes in FAR and dwelling sizes, as described above. The magnitudes of these changes were chosen to produce an unrestricted city that broadly matches the predictions of the simulation exercise of section 2.3.

To understand the choice process, observe that in the simulated city, the FAR restriction is binding over 24 percent of its land area.⁶ Based on the data shown in Figure 8, Bangalore's restriction is binding over nearly the same fraction of its built-up area, 23 percent (the share of rings 1–4). As a result, the impact of removing the FAR restriction, as measured by the reduction in the city's built-up area, should be similar between Bangalore and the simulated city. For the simulated city, this percentage reduction equals 17 percent,⁷ and under the

maintained assumptions, Bangalore's built-up area falls by exactly the same percentage (the area falls from 191 to 159 square kilometers).

Thus, the maintained assumptions, which include the specified FAR increases in rings 1–4, the 7 percent decline over the nonbinding range, and the 7 percent increase in dwelling size, generate an outcome that matches key aspects of the simulation. These percentage changes are much smaller than the 35 percent values from the simulation (see above), but using a similarly large percentage produces an implausible result, with the city growing spatially rather than shrinking when the FAR restriction is removed.

3.2. Computing the welfare cost of the FAR restriction

Figure 7 shows that, as a result of the city's spatial shrinkage following removal of the FAR restriction, the edge resident moves 3 km closer to the CBD. However, with Figure 7 also showing only a handful of residents in ring 12, it may instead be proper to view the distance decline for the edge resident as only 2 km. Welfare cost, which equals the resulting saving in the edge resident's commuting cost, is computed for both distance values.

While the 3 km distance reduction applies to the 1990 city, it is easier to estimate the resulting reduction in commuting cost using data from nearer the year 2000, which are more readily available. The magnitude of the commuting-cost saving is ultimately judged, however, by expressing it as a percentage of income. Under the assumption that incomes and commuting costs in Bangalore have grown in step with one another since 1990, this percentage value can be then used to measure the welfare cost for the 1990 city.

The key requirement for the welfare-cost calculations is an estimate of the commuting-cost parameter t based on Indian data. Commuting cost has two components, money cost and time cost, and as explained by Calfee and Winston (1998), traditional studies have assumed that commuters value the time spent in commuting at approximately 60 percent of the wage rate. Using a 1999 Indian wage estimate of 12.7 Rs. per hour, the implied time cost of commuting is then 7.6 Rs. per hour.⁸ Converting this figure into a cost per kilometer requires an estimate of traffic speed. Urban roadways in India are highly congested, and a common speed estimate is 20 to 30 km per hour. Using the lower value, the time cost of commuting is then $7.6/20 = .38$ Rs. per km.

The money cost of commuting depends on the transport mode chosen. Because estimates of operating costs for private vehicles are not readily available for India, a money cost estimate for bus commuters is used instead. Since money cost for this mode is presumably cheaper than for any type of private vehicle, the estimate serves as lower bound on the average money cost of commuting, thus biasing the welfare-cost computation in a conservative direction.

The estimate makes use of information provided by Zhi Liu of the World Bank, based on personal contacts with officials at the Bangalore Metropolitan Transit Corporation. The information shows that Bangalore bus fares are distance-based, and that the average cost per km traveled is 0.40 Rs. Adding this figure to the above time cost estimate yields an overall commuting cost of 0.78 Rs. per km.

This number must be multiplied by 2 to put it on a round-trip basis and multiplied again by 300 work days per year to convert it to an annual basis.⁹ These calculations yield a commuting cost of 468 Rs. per km per year, using late-1990 prices.

If commuting distance for the edge resident falls by 3 km, as suggested by Figure 7, then the reduction in commuting cost equals 1404 Rs. per year. If the smaller decline of 2 km is relevant, then the commuting-cost reduction is 936 Rs. per year. The magnitude of these savings can be judged by comparing them to per capita income for Bangalore in 1999, which equaled approximately 28,300 Rs.¹⁰ The commuting-cost savings range from 3.3 to 5.0 percent of per capita income, representing a significant share of individual resources.

While per capita income provides a convenient benchmark, its use may overstate the importance of the commuting-cost savings on a household basis given that not all members are wage earners. Recognizing this possibility, an alternative way of gauging the magnitude of these savings is to compare them to total household consumption expenditure, which is measured via survey methods. The NSS Consumer Expenditure Survey for 1999-2000 documents an average household consumption for urban India of 3867 Rs. per month, or 46,400 Rs. per year. If a household contains two wage earners, with each experiencing the above savings in commuting cost, then the combined savings range between 1872 and 2802 Rs., representing between 4.0 and 6.0 percent of household consumption. If the average number of wage earners is instead 1.5, then savings range between 1404 and 2106 Rs., representing between 3.0 and

4.5 percent of household consumption.

These percentage values have roughly the same order of magnitude as those based on the per capita approach, contrary to expectations. The reason is that household consumption figure of 46,400 Rs. is only 64 percent larger than the per capita income value of 28,300 Rs. even though the average Indian household contains approximately 5 people.¹¹ According to Salman Zaidi of the World Bank, this discrepancy reflects a growing divergence between measures of Indian household purchasing power based, alternatively, on survey data and the national income accounts. However, since the survey data may provide a more realistic picture of household circumstances, the above savings percentages, which lie in the 3–6 percent range of household consumption, may be credible.

As noted above, this percentage range is derived using an estimate of t based on late-1990's data, not on data for the 1990 city. However, if commuting costs and incomes grew by the roughly the same proportion over the period since 1990, then use of 1990 data would have yielded similar percentage values. As a result, the 3–6 percent range for welfare cost as a share of housing consumption expenditure is relevant for 1990. Finally, note that since these percentages represent the gains from removal of the FAR restriction, they capture the welfare cost of its presence.

A welfare cost of 3–6 of household consumption is substantial in size, being larger in magnitude than the measured welfare cost of some key distortions in Western economies.¹² Moreover, in a country like India, where vast numbers of people live on the edge of impoverishment, this welfare loss may represent the difference between poverty and non-poverty status for many households.

4. Conclusion

This paper has analyzed the impacts and the welfare cost of building-height restrictions, providing a concrete welfare-cost estimate for the city of Bangalore. Relying on several theoretical results, the analysis shows that the welfare cost imposed on its residents by Bangalore's FAR restriction ranges between 3 and 6 percent of household consumption. This burden represents a significant share of individual resources, and its presence may push many marginal

households into poverty.

The paper has carried out a counterfactual exercise, asking what Bangalore would look like today had it developed in the absence of an FAR restriction. Policymakers, however, may be concerned with a more practical question: conditional on Bangalore's history, how would the city's future development be affected if the FAR restriction were removed today? Recognizing that Bangalore's population will continue to grow as land-use adjusts to a removal of the restriction, an answer might be provided as follows. First, instead of growing at the boundary, the city would accommodate new population growth over an initial period through redevelopment of its central area at higher densities. Once this central redevelopment process was completed, further population growth would then occur at the city's edge. Because of its higher central densities, Bangalore at each future point in time would be more spatially compact, and housing prices would be lower, than if the FAR restriction had remained in place. As a result, consumer welfare would be higher at each point in time than in the presence of the restriction.

As for infrastructure investment, removal of the FAR restriction would require increased spending in the city's central areas, as explained in the introduction. However, with the city more compact at each point in time, the need for spatial expansion of its infrastructure network would be reduced, producing an offsetting reduction in costs.

Overall, the future gains from removal of the FAR restriction thus mirror in a qualitative sense the benefits documented in the counterfactual exercise carried out in the paper. Quantifying these gains, however, would be difficult given uncertainty about the city's future population path and other aspects of the future economic environment.

Finally, several extensions of the theoretical analysis in Section 2 might prove illuminating. The first would eliminate absentee land ownership, assuming instead that urban residents own equal shares of the region's land area. In this case, changes in per capita rental income would enter the welfare-cost calculation along with the factors identified above. A second extension would introduce traffic congestion, allowing commuting costs to rise with traffic flows. By leading to inefficient spatial expansion of a city, unpriced traffic congestion amplifies the tendency already generated by an FAR restriction (see, for example, Wheaton (1998)). As

a result, one might expect the restriction to generate a higher welfare cost in a congested city. A final extension would analyze the effect of an FAR restriction in a city where the location of employment is endogenous, allowing the emergence of subcenters. If an FAR restriction were added to Lucas and Rossi-Hansberg's (2002) variant of the Fujita-Ogawa (1982) model, which allows variable densities, then the equilibria with dispersed rather than concentrated employment would presumably be more likely. Because of their complexity, the latter two extensions would probably require the use of numerical methods.

Appendix

A.1. Simulation assumptions

The Cobb-Douglas utility function is written $v(c, q) = c^{1-\alpha}q^\alpha$, and α is set equal to 0.1, yielding a relatively small weight on housing consumption. Maximizing v subject to the budget constraint and substituting the resulting c and q demand functions into the condition $v(c, q) = u$ yields a solution for p , which is written $p(x, u) = A(y - tx)^{1/\alpha}u^{-1/\alpha}$, where A is a constant involving α . Substituting the $p(x, u)$ into the demand function for q yields $q(x, u) = B(y - tx)^{(\alpha-1)/\alpha}u^{1/\alpha}$, where B is a constant involving α .

The housing production function is also Cobb-Douglas, yielding an intensive form given by $h(S) = gS^\beta$, with $\beta = 0.6$ and $g = 0.0005$. The latter coefficient is adjusted to generate realistic FAR values and population densities, although it has no effect on the values of \bar{x} and u . Solving the first order condition $ph'(S) = 1$ for S and substituting $p(x, u)$ yields $S(x, u) = C(y - tx)^{1/\alpha(1-\beta)}u^{-1/\alpha(1-\beta)}$, where C is a constant involving α , β and g . Substituting this solution along with $p(x, u)$ into the zero-profit condition then yields $r(x, u) = D(y - tx)^{1/\alpha(1-\beta)}u^{-1/\alpha(1-\beta)}$, where D is a constant that depends on α , β and g . These functions are then substituted into the equilibrium conditions (4)–(5) and (7)–(9), which are solved numerically using Mathematica.

As mentioned in the text, the parameter assumptions underlying the solutions include $N = 800,000$, $y = \$42,151$, and $\theta = 4$. The estimate of r_a is based on the average U.S. agricultural land value for the year 2000, which equals \$1,210 per acre. Assuming a 5 percent discount rate, this value yields a rent per acre of \$60.50. This figure in turn implies a rent per square mile of \$38,720, which is the r_a value used in the simulations. To derive the time cost of commuting, the \$42,151 household income is used to compute an hourly wage of \$16.86, assuming 2000 work hours per year per worker and 1.25 workers per household. Then, assuming that commute time is valued at 60 percent of the wage, and that rush hour traffic moves at 30 miles per hour, the implied value of commuting time is \$0.34 per mile per worker. Adding an operating cost of \$0.36 per mile (the current Federal allowance), total commuting cost per

mile is then \$0.70. Multiplying by 1.25 workers per household, by 250 work days per year, and again by 2 to convert to a round-trip basis, commuting cost per mile for a household on a yearly basis is \$437.50 per year, which is rounded up to yield a t value of \$450.

A.2. Proof of Proposition 2

To establish Proposition 2, let the indirect utility function corresponding to $v(c, q)$ be given by $\tilde{v}(y - tx, p)$. Then, for the edge residents with and without the FAR restriction, $\tilde{v}(y - t\bar{x}_1, \bar{p}) = u_1$ and $\tilde{v}(y - t\bar{x}_0, \bar{p}) = u_0$ hold. Moreover, the compensation Z required to offset the effect of the FAR restriction for the edge resident, which must satisfy $\tilde{v}(y - t\bar{x}_1 + Z, \bar{p}) = u_0$, is given by $Z = t(\bar{x}_1 - \bar{x}_0)$. To show that the same Z is appropriate for interior residents, replace \bar{x}_1 with x in the last utility condition, and totally differentiate it with respect to x . This computation yields $\tilde{v}_y(-t + Z_x) + \tilde{v}_p p_x = 0$, or $Z_x = -p_x(\tilde{v}_p/\tilde{v}_y) + t$. However, since $\tilde{v}_p/\tilde{v}_y = -q$ holds given the properties of the indirect utility function, this last expression reduces to $q(p_x + t/q)$. But using the well-known formula for slope of the housing price function, which is written $p_x = -t/q$, this last expression equals zero. Therefore, $Z_x = 0$ holds, implying that the compensation required to offset the effect of the FAR restriction is spatially invariant and thus equal to $t(\bar{x}_1 - \bar{x}_0)$ for all residents.

A.3. Data description

The existing land use for Bangalore was disaggregated into successive rings centered on the CBD at 1 km intervals. The data were extracted from a GIS database containing census populations by subdistrict, a land use map showing the built-up area, and a regulatory map showing the permitted FAR values in different parts of the city. An estimate of housing floor space per person was used to compute actual FAR values, as follows. First, the estimate was generated by assuming that floor space follows an increasing exponential function, with values of 7 square meters at the CBD and 13.4 square meters at the edge of the city. Multiplying floor space per person by the ring population then yields total ring floor space. Finally, dividing this quantity by net residential land in the ring yields the ring's FAR value. Net residential land is computed by taking an appropriate share of the ring's total built up area.




Figure 1: Simulated FAR patterns




Figure 2: Simulated housing price patterns




Figure 3: Simulated dwelling size patterns




Figure 4: Simulated land rent patterns




Figure 5: Simulated population (dwelling) density patterns




Figure 6: Bangalore FAR values




Figure 7: Bangalore population patterns




Figure 8: Built-up area

References

- ALONSO, W., 1964. *Location and land use* (Harvard University Press, Cambridge).
- ARNOTT, R.J., MACKINNON, J.G. 1977. Measuring the costs of height restrictions with a general equilibrium model. *Regional Science and Urban Economics* 7, 359-375.
- BERTAUD, A., RENAUD, B., 1997. Socialist cities without land markets. *Journal of Urban Economics* 41, 137-151.
- BOADWAY, R.W., WILDASIN, D.E., 1984. *Public Sector Economics* (Little, Brown and Company, Boston).
- BRUECKNER, J.K., 1987. The structure of urban equilibria: A unified treatment of the Muth-Mills model. In: Mills, E.S. (Ed.), *Handbook of Regional and Urban Economics*, Vol. 2, North Holland, Amsterdam, pp. 821-845.
- BRUECKNER, J.K., 1996. Welfare gains from removing land-use distortions: An analysis of urban change in post-apartheid South Africa. *Journal of Regional Science* 36, 91-109.
- CALFEE, J., WINSTON, C.. 1998. The value of automobile travel time: Implications for congestion policy. *Journal of Public Economics* 69, 83-102.
- FISCHEL, W.A., 1990. Do growth controls matter? A review of empirical evidence on the effectiveness and efficiency of government land use regulation (Lincoln Institute of Land Policy, Cambridge MA).
- FUJITA, M., OGAWA, H., 1982. Multiple equilibria and structural transition of non-monocentric urban configurations. *Regional Science and Urban Economics* 12, 161-196.
- LUCAS, R.E., ROSSI-HANSBERG, E., 2002. On the internal structure of cities. *Econometrica* 70, 1445-1476.
- MILLS, E.S., 1967. An aggregative model of resource allocation in a metropolitan area, *American Economic Review* 57, 197-210.
- MUTH, R.F., 1969. *Cities and Housing* (University of Chicago Press, Chicago).
- PINES, D., SADKA, E. 1986. Comparative statics analysis of a fully closed city, *Journal of Urban Economics* 20, 1-20.
- WHEATON, W.C., 1974. A comparative static analysis of urban spatial structure. *Journal of*

Economic Theory 9, 223-237.

WHEATON, W.C., 1998. Land use and density in cities with congestion. *Journal of Urban Economics* 43, 258-272.

Footnotes

*This paper was commissioned by the World Bank, but its contents reflect the opinions of the authors. We thank Richard Arnott and Robert Buckley for helpful comments. Any shortcomings in the paper, however, are our responsibility.

¹Otherwise, h/q must be multiplied by persons per dwelling to generate population density. If this value is constant throughout the city, then the only effect of this change is to introduce a multiplicative factor into (5).

²With utility higher in this case, density now must be lower outside \hat{x} in the restricted city than in the unrestricted city (recall $S_u < 0$ and $q_u > 0$). In addition, density is again lower in the restricted city inside of \hat{x} . This conclusion follows because $u_1 > u_0$ implies that $h(S(\hat{x}, u_0))/q(\hat{x}, u_0) > h(S(\hat{x}, u_1))/q(\hat{x}, u_1) = \hat{h}$ holds in the unrestricted city. Since $h(S(x, u_0))/q(x, u_0)$ exceeds $h(S(\hat{x}, u_0))/q(\hat{x}, u_0)$ for $x < \hat{x}$, $h(S(x, u_0))/q(x, u_0) > \hat{h}$ must then hold over this range. Thus, the restricted city is everywhere less dense, and since it is assumed to be smaller spatially than the unrestricted city, it again cannot accommodate its population.

³Arnott and MacKinmon (1977) present a diagram similar to Figure 1.

⁴To establish this fact, let the values of p and structural density S at \bar{x} be denoted \bar{p} and \bar{S} . Then note that these values are determined by two conditions: the first-order condition $\bar{p}h'(\bar{S}) = 1$ and the zero-profit condition $\bar{p}h'(\bar{S}) - \bar{S} - r_a = 0$. Since \bar{x} itself does not explicitly enter these equations, \bar{p} does not depend on its magnitude.

⁵Using the solutions from the appendix, FAR falls by the proportion $(u_0/u_1)^{-\beta/(\alpha(1-\beta))}$ and q rises by the proportion $(u_0/u_1)^{1/\alpha}$.

⁶This percentage is given by $\hat{x}^2/\bar{x}_1^2 = 11.7^2/23.5^2 = .24$

⁷This percent is given by $\bar{x}_0^2/\bar{x}_1^2 = 21.4^2/23.5^2 = .83$.

⁸This estimate makes use of a wage-estimation methodology developed by the U.S. Government Import Administration and documented at <http://ia.ita.doc.gov/wages>. The methodology uses wage and per capita GNP data for a cross section of 56 countries to compute the following regression: Wage = 0.462 + .000432*(per capita GNP). Using the per capita income estimate of 28,300 Rs. for Bangalore (see footnote 10), this regression predicts a wage of approximately 12.7 Rs.

⁹The assumption of 300 work days per year can be justified by assuming 50 work weeks per year and 6 work days per week.

¹⁰The Bangalore Directorate of Economics and Statistics provides per capita income figures for the city, with the 1998-1999 value equal to approximately 28,300 Rs. (as quoted on the website <<http://www/bangaloreit.com/html/aboutbug/bangprofile.htm>> operated by the Karnataka state government).

¹¹Asian Demographics Ltd. <<http://www.asiandemographics.com/DI231102.htm>> provides a household-size estimate of 5.07 persons.

¹²For example, Boadway and Wildasin (1984, p. 391) cite empirical estimates showing that the deadweight loss from capital income taxation in the U.S. ranges between 2 and 12 percent of tax revenue. Expressed as a share of the economy's overall income, the relevant percentages would be much smaller.