

Bidding and Sorting

PAI 735/ECN 635

State and Local Government Finance

Professor Yinger

Introduction

One of the central questions in local public finance is: How do households select a community in which to live? These notes provide an intuitive overview of the academic literature on this question.⁽¹⁾ Scholars have also considered two related questions: How do communities select the level of local public services and tax rates? Under what conditions is the community-choice process compatible with the local voting process? These two questions are not addressed here.⁽²⁾

The literature on community selection is often traced back to a famous paper by an economist named Charles Tiebout, who argued in 1956 that people shop for a community, just as they shop for other things.⁽³⁾ This process came to be known as "shopping with one's feet." The Tiebout model became very influential both because it identified an important type of behavior to be studied and because it implied that the process of allocating people to communities has certain desirable properties. To be specific, Tiebout claimed that this process, like any unfettered market process (without externalities) is efficient. As it turns out, Tiebout's analysis was highly simplified, and indeed left out both the housing market and the property tax, which is the main source of local revenue in the United States. Thus, many people regard this analysis as unrealistic and reject its efficiency conclusion, but others think that this conclusion is still correct, at least to a first approximation.

These notes review the literature since Tiebout. They start by addressing the way a household decides how much to pay for housing in a given community, called bidding, and what happens when different types of households compete for housing, called sorting. The notes also consider

an important empirical phenomenon, called capitalization, which is a test of bidding and sorting models, and explore the normative issues first raised by Tiebout.⁽⁴⁾

Bidding

The central positive question addressed in these notes is: How does the housing market allocate households to communities when local public services and property taxes vary from one community to the next? A broad consensus has emerged concerning the appropriate way to analyze this question. This consensus has general applicability to any country with active housing markets, reasonably mobile households, and multiple local governments that exhibit some variation in public service levels or tax rates. It evolved from the insight made famous by Tiebout, namely that households care about local public services and local taxes and compete for entry into the most desirable communities. This consensus has two components: bidding and sorting.

Bidding analysis builds on a variety of assumptions that roughly characterize urban areas in the United States. First, this analysis assumes that households fall into distinct income/taste classes and that all households care about their consumption of housing, public services, and other goods. Households within a class are considered to be identical in their demands for these things, but many classes may exist. Households are also assumed to be mobile, that is, to be able to move costlessly from one jurisdiction to another. This assumption implies that an equilibrium cannot exist unless all people in a given income/taste class achieve the same level of satisfaction or, to use the economist's term, of utility. In other words, any household that does not reach as high a utility level as do similar households will have an incentive to move, and this type of moving behavior will lead to a situation in which all similar households have the same utility (and no household has an incentive to move).

All households who live in a jurisdiction are also assumed to receive the same level of public services, and residence in a jurisdiction is assumed to be a precondition for the receipt of public services there. Each urban area is assumed to have many local jurisdictions, which have fixed boundaries and vary in their local public service quality and property tax rates. Finally, in most

models, households are homeowners, not renters, and local public services are financed through a local property tax.

These assumptions are used to understand where people live and how much they pay for housing. Thus, they are introduced in a basic analysis of the housing market in which households compete with each other for access to the most desirable locations. In this context, the most desirable locations are those with the best combination of high-quality public services and a low property tax rate. Households compete for entry into desirable locations by bidding against each other for housing. To proceed with our analysis, therefore, we must consider bidding behavior in a housing market.

An analysis of bidding behavior depends on four key concepts. Housing is measured in units of housing services, which provide an index of the size and quality of a housing unit. For our purposes, we can think of housing services as simply the size a house in square feet, although the concept is easily generalized to consider other features of housing. The price per unit of housing services is the associated price concept per unit time, say per year. The rent for a housing unit equals the price per unit of housing services multiplied by the number of units of housing services the unit contains. If the unit is an apartment, this rent is equivalent to the annual contract rent. If the unit is owner-occupied, this rent is not observed in the market place but is implicit. The value of a housing unit is the amount someone would pay to own that unit; it equals the present value of the flow of net rental services associated with ownership.⁽⁵⁾

The price per unit of housing services is assumed to be constant at a given jurisdiction; that is, each unit of housing services in a given jurisdiction costs the same amount. However, this price (and hence both rent and value) varies across jurisdictions as the quality of public services and the local property tax rate varies. Consider first the case of a single income/taste class, that is, of an urban area in which all households are alike. Because households are mobile, as well as alike, each household must achieve the same utility level. As a result, households who live in jurisdictions with relatively desirable service-tax packages must pay for the privilege in the form of higher housing prices; otherwise, these households would be better off than households in other jurisdictions and those other households would have an incentive to move. Another way to

put this is that households who end up in jurisdictions with relatively undesirable service-tax packages must be compensated in the form of relatively low housing prices. This analysis can be summarized in the form of a bid function, which indicates the maximum amount a household would pay to live in a jurisdiction as a function of the desirability of the service-tax package there.

This logic is summarized in Figures 1 and 2. In Figure 1, a household's bid, expressed as a price per unit of housing services, is on the vertical axis and the quality of local public services is on the horizontal axis. In this figure, the property tax rate is held constant. As the quality of public services goes up (holding taxes constant), households' bids for housing go up. In other words, people are willing to pay more for housing in a jurisdiction with better public services. Figure 2 plots household bids as a function of the property tax rate (holding public service quality constant). The higher the property tax rate, the lower the bid, all else equal.

Thus, the consensus view described in these two figures predicts that rents will increase with public service quality and that property values will be higher in communities with higher public service quality and or lower property tax rates, all else equal. One way to test this view, therefore, is to determine whether rents and property values do, in fact, depend on these variables. The empirical evidence on this phenomenon, which is called "capitalization," is discussed in a later section.

Sorting

Figures 1 and 2 describe the housing bids for one type of household but do not reveal how different types of households are sorted into jurisdictions. The key to understanding sorting is to recognize that bid functions like the one in Figure 1 are steeper for some types of households than for others. The steepness of a bid function indicates the extent to which a household type's bids for housing increase when the quality of public services increases. The steepness of a bid function matters because landlords (or housing sellers) prefer to sell to the household type that is willing to pay the most per unit of housing services. Thus, as shown in Figure 3, households with relatively steep bid functions win the competition for housing in locations where the quality of public services is relatively high, and households with relatively flat bid functions win the

competition for housing in locations where the quality of public services is relatively low. For example, the group with the steepest bid function in this figure wins the competition for all levels of public service quality above S_3 .

Under normal circumstances, high-income households have steeper bid functions than low-income households. In other words, high-income households are willing to pay more for an increment in public service quality. This does not mean that high-income households intrinsically value service quality more than do low-income households (although this might be true); instead, it simply means that high-income households, who after all have greater resources, increase their bids for housing by a greater amount when service quality goes up.

This relationship between income and bid-function steepness implies that high-income households live in locations with relatively high quality public services, a situation often called "matched sorting." This situation is illustrated by Figure 3 when the steeper bid functions (the ones to the right) belong to a higher-income classes and the flatter bid functions (the ones to the left) belong to a lower-income class. In this case, it is clear that the highest-income households win the competition for housing in the highest-service jurisdictions.

Figure 3 also reveals that low-income households win the competition for housing in jurisdictions with relatively low public service quality, such as the jurisdictions providing quality below S_1 . How, one might ask, can low-income people ever outbid high-income people for housing? The key to answering this question is to remember that bids are expressed per unit of housing services. Hence low-income people win the competition because they bid a high amount per unit of housing services but consume a relatively low quantity of housing services. If housing services are measured in square feet, then low-income people select small apartments or double up so that they consume very few square feet per capita. In Figure 3, low-income households bid more per square foot than high-income households in jurisdictions with low levels of S , and a landlord (or seller) in those jurisdictions can make more money by renting (selling) to low-income people who double up than to high-income people who do not.

The logic of sorting implies that jurisdictions tend to be homogeneous in terms of income and preferences, but it does not rule out the possibility of heterogeneous jurisdictions. In fact, if a jurisdiction provides a level of public service quality at which the bid functions of two groups cross, both groups are content to live there. This situation is illustrated by S_1 , S_2 , and S_3 in Figure 3. At each of these public service levels, two groups bid the same amount for housing, so a jurisdiction delivering any of these service levels will contain two household types. Moreover, it is possible to have a sorting outcome in which many different household types bid the same amount per unit of housing services in a relatively large jurisdiction, such as a central city, and hence all live there.

The sorting analysis in Figure 3 does not consider property taxes. As it turns out, this is entirely appropriate because property taxes do not affect sorting even though they affect bids. The outcome reflects the fact that all household types, regardless of income or preferences, are willing to pay \$1 to avoid \$1 of property taxes. The same cannot be said for public service quality. Some groups are willing to pay more than others for an improvement in service quality, which is precisely why the bid functions in Figure 3 do not all have the same slope.

Although property taxes do not affect sorting, it is instructive to draw another version of Figure 3 that has property taxes built in. To take this step, we must first consider Figure 4, which presents a household's demand curve for local public services. Recall that a demand curve indicates the marginal benefit from some good or service, that is, the amount a household (or set of households) is willing to pay for an additional unit of that good or service. In Figure 4, a household receiving a relatively low service quality is willing to pay a relatively high amount for an increment in service quality, whereas a household receiving a relatively high service quality is only willing to pay a relatively small amount for a comparable increment. In other words, the marginal willingness to pay for local public service quality declines as service quality increases.

Although the marginal willingness to pay declines with service quality, the marginal cost of an increment in service quality is constant. In other words, it takes the same property tax increase to fund a one-unit increase in service quality whether one starts from a low or high public service quality in Figure 4. As service quality increases, therefore, the net willingness to pay for an

increment in public service quality declines and may eventually become negative. At very high levels of service quality, in other words, the small benefit from an increment in quality may be outweighed by the required property tax increase.

This situation is illustrated in Figure 5, which now has community income on the horizontal axis and housing bids on the vertical axis.⁽⁶⁾ In this figure, one must distinguish between the income of a class of households searching for housing, which determines the slope of their bid function, and the income of an existing community, Y , which determines the service quality and property tax rate there. As community income increases, the demand for public service quality increases, but any increase in service quality must be funded with higher property taxes. As a result, bid functions now increase at relatively small levels of service quality, but then become flatter and eventually decline as service quality increases. Figure 5, like Figure 3, illustrates that household types with relatively steep bid function (presumably high-income households) win the competition for housing in jurisdictions with relatively high service quality.

An alternative approach to sorting was proposed by an economist named Bruce Hamilton in 1975.⁽⁷⁾ This approach adds several additional assumptions.⁽⁸⁾ First, it assumes that housing suppliers respond to any variation in the price of housing across communities by expanding the supply of housing in communities with higher prices. This assumption implies, unlike the conventional view, that the price of housing will not be higher in a community with higher public service quality or lower property tax rates. Second, this approach assumes that a jurisdiction can use zoning restrictions to set the consumption of housing at exactly the optimal level.⁽⁹⁾ This assumption is important because it rules out the possibility that low-income households want to move into high-income communities in order to take advantage of the large tax bases there. After all, the property tax a household pays depends not only on its own property value but also on the average value of the property in its jurisdiction.

A more formal statement of this issue requires a new concept, namely, a household's tax price for local public services. In the case of a private commodity, a household can purchase as much as it wants at the posted market price. In the case of a local public service, however, the amount a household pays for an increment in service quality depends on the property tax system. The tax price is simply the household's share of any increase in property taxes needed to fund an increase

in public service quality. Even for a given type of household, this tax price can vary widely across jurisdictions, in part because some jurisdictions have more expensive houses and in part because some jurisdictions have far more commercial and industrial property than others. The tax price is lower in a district with a great deal of commercial and industrial property because much of the burden of any school tax increase falls on commercial and industrial taxpayers, not on homeowners.

The easiest way to derive an expression for a tax price is to combine a single voter's property tax payment with the local government budget constraint. A household's property tax payment, T , equals an effective property tax rate, t , multiplied by the value of the household's property, V , or $T = tV$. A local government must set spending per household, E , equal to the property tax rate multiplied by average property value per household, V^* , or $E = tV^*$.⁽¹⁰⁾ This government budget constraint implies that the tax rate must equal spending divided by the tax base or $t = E/V^*$. Substituting this tax rate into the household's tax payment, we find that $T = E(V/V^*)$. This result implies that the tax-price equals (V/V^*) . The amount someone spends on a private good is the quantity of the good multiplied by the price. In this case, the quantity of the local public services is E , and its "price" is (V/V^*) .

Now we can return to the issue of sorting. Consider a household with a house value (V) equal to \$100,000. If it lives in a community where the average property value (V^*) equals \$50,000, then its tax-price is 2, whereas if it lives in a community where the average property value equals \$200,000, its tax price is only 0.5. In other words, moving from the first community to the second is equivalent to cutting the price it pays for local public services by 75 percent! Obviously, the household will make this move unless the price of housing is higher in the second community. The Hamilton assumption rules out this type of move, and hence the need for price variation across communities, by forcing all households in a given community to purchase the same value of housing. The household in this example cannot move into the second community under the Hamilton assumption because everyone in that community must purchase a house worth \$200,000. Thus, an equilibrium can be attained with the same housing price in both communities.

With the Hamilton assumptions, the sorting diagram takes yet another form, as shown in Figure 6. Two features of this figure should be emphasized. First, every household lives in the jurisdiction that provides the service-tax package it is willing to pay the most for, which is, by definition, its optimal service-tax package. This outcome is efficient because no household can gain by moving to another jurisdiction. Thus, the Hamilton assumptions reproduce Tiebout's main normative result, namely that a system of local governments is efficient.

Second, in Figure 6 the price of housing is the same in every community. In other words, the Hamilton assumptions ensure that there is no capitalization. The price of housing does not reflect the quality of local public services or the local property tax rate. A simple test of the validity of the Hamilton model, therefore, is to determine whether capitalization exists.

Capitalization

A famous article by an economist named Wallace Oates in 1969 first tested the hypothesis that house values depend on local public service quality and on local property tax rates, a phenomenon known as capitalization.⁽¹¹⁾ His strong evidence for capitalization in a sample of communities in New Jersey stimulated dozens of additional studies on the topic. This section briefly reviews recent findings on the capitalization of both local property taxes and local public service levels.

Property Tax Capitalization

Property tax capitalization arises from the basic equality between the value of an asset, in this case a house, and the present value of the net benefits from owning it. The amount someone is willing to pay for a house is the present value of the rental benefits minus the present value of the property tax payments.⁽¹²⁾ As noted earlier (see footnote 5), the value of a long-lived house with an annual rental value equal to R (and no property taxes) can be written as $V = R/r$, where r is a discount rate. Annual property taxes equal tV , and the present value of this tax stream is tV/r . Thus, with property taxes, the house value formula is $V = (R/r) - (tV/r)$. Solving this equation for

V yields $V = R/(r + t)$. As illustrated in Figure 2, this relationship between house values and property taxes also can be derived from a bidding model.

This formulation assumes that a \$1 increase in the present value of property taxes leads to \$1 decrease in the value of a house, which is equivalent to assuming full (or 100 percent) capitalization. The empirical literature on this topic is designed to determine whether capitalization exists and to estimate the degree of capitalization. Let β stand for the degree of capitalization. If β equals 0.5, for example, then a \$1 increase in the present value of property taxes leads to a \$0.50 decrease in the value of a house. With this amendment, the above formula can be re-written as follows:

$$V = \frac{R}{r} - \beta \left(\frac{tV}{r} \right) = \frac{R}{r + \beta t} \quad (1)$$

The objective of the literature on property tax capitalization is to estimate "beta."

Although this equation is fairly simple, it has proven to be difficult to estimate for several reasons.⁽¹³⁾ First, it involves a non-linear relationship between t and V , even after taking logarithms, so it cannot be estimated with linear regression methods. As a result, existing studies use various approximations or non-linear estimating techniques. When data on tax and house value changes over time are available, simple but exact functional forms are possible. Let subscripts indicate time periods, then (1) leads to

$$\frac{\Delta V}{V_1} = \frac{-\beta \Delta t}{r + \beta t_2} \quad \text{or} \quad \Delta V = \frac{-\beta \Delta T}{r} \quad (2)$$

where "delta" indicates a change over time. If t_2 is roughly constant (as it might be within a single jurisdiction after a comprehensive revaluation of all property), then both of these equations reduce to a constant multiplied by the change in a tax variable. The second equation

can be estimated, for example, when tax payments undergo a large change, as in the case of Proposition 13, a property tax limitation passed in California in 1978.

Second, the value of the discount rate, r , is not observed, and the form of equations (1) and (2) precludes separate estimation of r and β . Most studies follow Oates by estimating a value of β/r , assuming a value for r , and then calculating the implied value of β .⁽¹⁴⁾ The trouble with this approach is that the value of β depends on an untested assumption and different studies use different values of r . In fact, the most extreme estimates in the literature, in either direction, are driven largely by extreme assumptions about r .

Third, the asset-pricing logic behind equations (1) and (2) requires assumptions about house buyers' expectations. To be specific, this logic predicts that a \$1 increase in the present value of future property taxes will lead to a \$1 decline in house value (i.e. "beta" = 1), but it does not say that current tax differences will be fully capitalized if they are not expected to persist. Virtually all the literature estimates the capitalization of current property tax differences. Under the assumption that current differences will persist indefinitely, the assumption that "beta" = 1 makes perfect sense. In fact, however, current differences may not be expected to persist. In this case, the capitalization of current tax differences, the "beta" in (1) and (2), is related to the capitalization of the present value of the expected tax stream, say β' , as follows:

$$\beta = \beta' [1 - (1 + r)^{-N}] \quad (3)$$

where N is the length of time current tax differences are expected to persist. The theory indicates that "beta prime" = 1, but the estimated clearly need not equal 1, and indeed need not equal the same value under all circumstances. For example, if current property tax differences across communities are expected to disappear in 10 years and $r = .03$, then (3) implies that the estimated "beta" will be only 26 percent even if "beta prime" = 1.

Virtually every study of property tax capitalization finds a statistically significant negative impact of property taxes (or property tax changes) on house values. The vast majority these studies use data from the United States, but a few use data from Canada. Estimates of vary widely, but if r is set at 3 percent, the estimates of "beta" for the best studies fall between 15 and 60 percent. No study yet provides a definitive estimate of "beta prime," but several recent studies provide some evidence that it is close to 100 percent. These results are consistent with the consensus bidding/sorting model, but strongly contradict the Hamilton approach.

Public Service Capitalization

Figures 1 and 3 show that the price of housing services, P , is a function of local public service quality, S . It follows that apartment rents and house values are also functions of S . Many studies, again starting with the study by Oates, have tested this prediction. As it turns out, finding public service capitalization is intrinsically more difficult than finding property tax capitalization, largely because the demand for public services varies across households. As noted earlier, any household will pay \$1 to avoid \$1 in property taxes, but households differ in the amount they are willing to pay for another unit of public services. Moreover, the quality of local public services is notoriously difficult to measure, and existing data often do not provide information on many dimensions of service quality.

One approach is to use spending per capita as a measure of public service quality. This was the approach used by Oates, who focused on education spending. As it turns out, spending is a fairly poor measure of service quality because it assumes, contrary to extensive evidence, that \$1 of spending per capita buys the same public service quality in all communities. Several more recent studies use student test scores or crime rates or similar measures to determine the quality of local public services. This is a more reasonable approach if the required data are available.

Virtually all studies of public service capitalization find a positive, statistically significant impact of public service quality on house value (or rent).⁽¹⁵⁾ One study finds, for example, that moving from the last to the first decile on a fourth grade reading score would increase median property values in a community by \$4,300, all else equal.

A few studies have identified special circumstances that yield particularly compelling evidence about the capitalization of local public services. One study compares the values of houses that are located in the same municipality but in different school districts, a situation made possible by unusual jurisdiction boundaries in the Cleveland area. After controlling for a long list of housing and neighborhood characteristics and accounting for differences in school property taxes, this study finds that people are willing to pay quite a lot, up to \$2,171 per year, to live in a good school district instead of a poor one. Another study compares the market values of houses on opposite sides of attendance-zone boundaries in the same school district. These houses obviously share many neighborhood characteristics other than school zone, thereby minimizing problems that can arise from omitted variables. This study also finds that housing prices increase with school quality.

In short, virtually every study in this literature finds evidence that the quality of local public services is capitalized into apartment rents and house values. These results therefore support one of the key predictions of the consensus view of bidding and sorting and contradict the Hamilton approach.

Normative Analysis

The lasting fame of Tiebout derives primarily from his claim that "voting with one's feet" is analogous to shopping for a commodity and leads to efficient provision of local public services. The literature has struggled with questions of efficiency and equity ever since.

Efficiency

Although the Tiebout model did not incorporate either a housing market or a property tax, Hamilton showed that the same result can be obtained from a model with these features, so long as his assumptions about housing suppliers and zoning are satisfied. In particular, the Hamilton assumptions imply that every community is homogeneous, in the sense that it contains a single income/taste class, and that voters select the level of local public services as if it had a constant price, just like a private good. Thus, the Tiebout analogy to shopping in a private market place is preserved and the efficiency properties of that type of shopping are attained.

The Hamilton approach has been widely cited as proof that our system of local governments is at least approximately efficient. However, as shown earlier, the key prediction of the Hamilton model, namely, no capitalization, is resoundingly rejected by dozens of studies on the topic. As a result, it makes no sense to accept the Tiebout/Hamilton conclusion that a system with many independent local governments is fully efficient.

Instead, two key conclusions emerge from the literature. The first, which is in the spirit of Tiebout and Hamilton, is that setting up a federal system with many independent local governments has profound efficiency advantages over a centralized system in which the quality of local public services is the same everywhere. This efficiency advantage arises because the demand curve for local public services, which is illustrated in Figure 3, varies widely across households.

Consider the two households in Figure 7. Household 1 has a low income and a low demand curve, D_1 , and household 2 has a high income and a high demand curve, D_2 . Each household's demand curve indicates its marginal benefit from public services, so its total benefit from any particular public service level is the area under its demand curve up to that service level. Now suppose MC is the cost per unit of public services for both households.⁽¹⁶⁾ Then each household would prefer to have the level of service at which its marginal benefit (its demand curve) equals its marginal cost ($= MC$). Thus, household 1 would prefer to live in a jurisdiction with service level S_1 , and household 2 would prefer to live in a jurisdiction with service level S_2 . If each household succeeds in finding such a jurisdiction, the outcome is efficient, in the sense that no other arrangement (short of redistribution) would make either household better off.

With a single local government, however, both households would be forced to consume the same level of local public services, say S_3 . In this situation, the level of local public services would be too high for household 1 and too low for household 2. To be specific, the units of public services between S_1 and S_3 would be delivered to household 1 even though the marginal benefit from each of those units falls below the marginal cost household 1 must pay. The loss to household 1 is the indicated shaded area. The units of public services between S_3 and S_2 would not be delivered to

household 2 even though the marginal benefit from each of those units falls above the marginal cost household 2 must pay. The loss to household 2 is the indicated shaded area. In short, moving from separate, homogeneous jurisdictions to a single jurisdiction leads to significant losses for both household types. Or, to put it another way, allowing one jurisdiction to break up into smaller, more homogeneous jurisdictions results in net benefits to both types of households and is therefore an efficient thing to do.

The second key conclusion is that despite the efficiency of many local governments compared to a single local government, our actual system of local governments is not nearly as efficient as it could be. In fact, dozens of capitalization studies reject the Tiebout/Hamilton assumptions and hence reject their conclusion of full efficiency. Instead, our system has several sources of inefficiency that could, at least in principle, be corrected or offset by the actions of higher levels of government.

Two potential sources of inefficiency are particularly important.⁽¹⁷⁾ First, local governments are not nearly as homogeneous as the idealized Tiebout/Hamilton view implies. Indeed, as shown earlier, heterogeneous communities are possible, even likely, in the consensus bidding/sorting analysis. Moreover, it can be shown a large central city is likely to provide a level of public services below the efficient level. If so, state or federal subsidies to city governments could actually enhance efficiency.⁽¹⁸⁾ Second, unless all the Hamilton assumptions are satisfied, the property tax, like other taxes, distorts household behavior and leads to a deviation from the efficient level of local public services. Scholars do not agree on the best way to eliminate this distortion.⁽¹⁹⁾

Equity

When local governments have considerable autonomy, as in the United States, and sorting occurs, some jurisdictions have much higher incomes and tax bases than others and end up with much higher quality public services. This effect is magnified by the fact that high-income jurisdictions tend to have favorable environments for providing public services and hence relatively low public service costs.⁽²⁰⁾ Moreover some jurisdictions have extensive commercial

and industrial property, which lowers their tax-price and thereby raises the quality of public services voters select.

Because the system of local governments is established by higher levels of government, by the states in the United States, for example, higher levels of government bear ultimate responsibility for the nature of this system. As a result, higher levels of government may be concerned with variation in local public services, and may want to compensate local governments for unfavorable fiscal factors that are outside their control, such as a low tax base, high input prices, or a harsh environment. This compensation by a state can take the form of an intergovernmental aid program that accounts for tax-base and cost differences across communities or of institutional changes, such as regional tax-base sharing or allowing cities to tax suburban commuters.

Fairness issues assorted with bidding and sorting have long been recognized in the case of education, both by academics and policy makers. Moreover, many American states, often in response to a court ruling, now provide higher grants per pupil to school districts with lower property values per pupil. Some grant systems are designed to bring all school districts up to a minimum spending level, and others are designed to ensure that all districts that levy the same property tax rate will receive the same spending per pupil. No existing grant system eliminates the correlation between property wealth and school spending, but some of them undoubtedly lower this correlation significantly.

The consensus bidding model also raises a different sort of equity issue that is related to the timing of ownership, not to the income distribution. To be specific, an unanticipated change in service quality or property taxes in one jurisdiction relative to others leads to capital gains or losses for the owners of property there at the time of the change. In fact, if capitalization is complete, the capital gain or loss equals the present value of the stream of annual changes in service benefits or taxes, and these owners bear the entire burden of the change. Moreover, households who buy a house from one of these owners in the future bear no burden at all because these future buyers will be compensated for higher taxes or lower service quality in the form of lower housing prices. As a result, capitalization creates classes of households based on the timing

of their home purchase, instead of their income or wealth. Some policies are unfair because they arbitrarily benefit or harm people in some of these time-dependent classes relative to others.

Consider, for example, a community in which poor assessing practices lead to effective property tax rates, that vary widely from one house to another.⁽²¹⁾ In this situation, an unanticipated revaluation that brings all houses to the same effective tax rate will result in capital losses for houses that were previously underassessed and to capital gains for houses that were previously overassessed. These gains and losses are largely arbitrary and hence unfair; a long-time owner whose house was previously underassessed could be said to be paying back tax breaks in the form of a capital loss, but the same loss falls on a recent buyer who gained nothing at all from the past underassessment. One cannot avoid this problem, however, by retaining the poor assessing system, because such a system hands out regular, small, unannounced effective tax rate cuts or increases as it allows assessments to diverge from market values. The resulting incremental gains and losses also are arbitrary, and hence unfair. The only way out of this dilemma is to pay the one-time fairness cost of revaluation and then keep assessments up to date in the future.

Conclusions

In sum, state and federal governments making decisions about a federal system face many challenges in designing the appropriate policies. Many people have the general presumption that decentralization promotes efficiency at the expense of fairness. However, this presumption should not be pushed too far. Some policies to promote decentralization may enhance efficiency while violating some standards for fairness, but other decentralization policies may enhance efficiency and fairness simultaneously. Moreover, the conclusion that decentralization tends to promote efficiency does not imply that any particular federal system is as efficient as it could be. Indeed, state and federal governments may find many ways to promote efficiency even in a fairly decentralized federal system.

1. For a detailed review of this literature, with extensive references, see Stephen Ross and John Yinger, "Sorting and Voting: A Review of the Literature on Urban Public Finance," in *Handbook of Urban and Regional Economics, Volume 3, Applied Urban Economics*, edited by P. Cheshire and E. S. Mills (North-Holland, 1999), pp. 2001-2060.
2. The literature on these two questions is reviewed in Ross and Yinger, op. cit.
3. Charles M. Tiebout, "A Pure Theory of Local Expenditures," *Journal of Political Economy*, vol. 64 (1956), pp. 416-424.
4. Recall that "positive" analysis refers to theories and empirical tests of theories concerning the way people behave. Any "positive" statement can, in principle, be tested against the evidence. In contrast, "normative" analysis, which may draw on positive analysis, concerns what is desirable and depends on the analyst's values.
5. For an asset with a long life, such as a house, the present value formula reduces to a simple form. Specifically, the value of a house, V , equals the annual rental flow, R , divided by an annual discount rate, say r ; that is, $V = R/r$. In this context, r is sometimes called the capitalization rate, because it translates an annual flow, R , into an asset or capital value, V .
6. A more general version of this analysis employs an index of the demand for local public services, not just income, on the horizontal axis. See Ross and Yinger, op.cit.
7. Bruce W. Hamilton, "Zoning and Property Taxation in a System of Local Governments," *Urban Studies*, vol. 12 (1975), pp. 205-211.
8. Hamilton makes another important assumption in addition to the two highlighted here, namely, that a unit of public services costs the same in every jurisdiction. This assumption rules out the possibility that the cost of public services depends on the characteristics of the people in a jurisdiction. Many studies have found that the cost of public services does depend on residents' characteristics (see Ross and Yinger, op. cit.), so this assumption is contradicted by the evidence.
9. Zoning tools do not, of course, literally determine housing consumption. However, Hamilton and others argue that the wide range of zoning tools (set-back requirements, lot-size restrictions, and so on) enable local governments to control housing consumption with a high degree of precision. See Ross and Yinger, op. cit.
10. This analysis can easily be extended to consider state aid and other local revenue sources. See Ross and Yinger, op. cit.
11. Wallace E. Oates, "The Effects of Property Taxes and Local Public Services on Property Values: An Empirical Study of Tax Capitalization and the Tiebout Hypothesis," *Journal of Political Economy*, vol. 77 (1969), pp. 957-971.
12. This formulation also explains the origin of the term "capitalization." As pointed out in

footnote 5, a present-value formula translates an annual flow into a capital value. This translation is called capitalization.

13. Two technical problems not considered here are that the tax rate variable is endogenous and the equation may be subject to omitted variable bias. For more on issues that arise in estimating capitalization, see John Yinger, Howard S. Bloom, Axel Boersch-Supan, and Helen F. Ladd, *Property Taxes and House Values: The Theory and Estimation of Intra-jurisdictional Property Tax Capitalization*, New York: Academic Press, 1988.

14. The required value of r is a real discount rate, defined as a nominal rate minus anticipated inflation. A reasonable real discount rate would be in the 2 to 4 percent range.

15. More details on these studies can be found in Ross and Yinger, *op. cit.*

16. This assumption implies that the two households have the same tax price, and hence implicitly assumes that they have the same house value. This assumption simplifies the analysis but is not necessary for the conclusions in the text.

17. Several other sources are discussed in the literature. First, the sorting process might not allocate households to communities in an optimal way. See D. E. Wildasin, *Urban Public Finance* (London: Harwood Academic Publishers, 1986). Second, local governments might themselves behave in an inefficient way (again, see Wildasin). Third, property tax and public service capitalization may prevent local voters from acting in an efficient manner (see Ross and Yinger, *op. cit.*).

18. For more on this argument, see Ross and Yinger, *op. cit.*

19. Great Britain experimented with a community-specific head tax as a way to avoid this distortion, but found that it was not only unpopular, but also ran into some of the efficiency problems discussed in footnote 17. See Ross and Yinger, *op. cit.*

20. Environmental factors are discussed at length in Ross and Yinger, *op. cit.*

21. This example is discussed in detail in Yinger, et al., *op. cit.*