

NBER WORKING PAPER SERIES

THE PRODUCTIVITY OF SCHOOLS AND
OTHER LOCAL PUBLIC GOODS PRODUCERS

Caroline M. Hoxby

Working Paper 6911
<http://www.nber.org/papers/w6911>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1999

The views expressed here are those of the author and do not reflect those of the National Bureau of Economic Research.

© 1999 by Caroline M. Hoxby. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Productivity of Schools and Other
Local Public Goods Producers
Caroline M. Hoxby
NBER Working Paper No. 6911
January 1999
JEL No. H70, I22, L32, L51

ABSTRACT

I construct an agency model of local public goods producers with special reference to public schools. The model assumes that households make Tiebout choices among jurisdictions, but it has more realistic assumptions about information and the cost of residential mobility. I examine producers' effort and rent under local property tax finance and centralized finance. I show that, if there are a sufficient number of jurisdictions to choose among, conventional local property tax finance substantially reduces the agency problem and associated loss of productivity. Specifically, I demonstrate that local property tax finance can attain about as much productivity as a social planner with centralized finance can, even if the social planner is armed with more information than a real social planner could plausibly have. The key insight is that decentralized Tiebout choices make some information the social planner would need *verifiable* and other information *unnecessary*.

Caroline M. Hoxby
Department of Economics
Harvard University
Cambridge, MA 02138
and NBER
choxby@harvard.edu

I. Introduction

This paper is motivated by our need to understand the productivity consequences of different methods of financing and providing local public goods. Both in the United States and abroad, systems such as vouchers, privatization, and decentralization are being considered as means of improving schooling, crime prevention, and other services that are commonly provided as local public goods. Much of the debate focuses how these systems would affect productivity. For a given amount of spending, do these systems cause providers to produce more of the desired service? Interest in productivity is partly driven by empirical work that shows that the productivity consequences of vouchers, privatization, and decentralization are important.¹ In industrialized nations, interest in productivity sometimes stems from real cost increases that are not matched by improved outcomes. For instance, over the period from 1970 to 1997, the real per-pupil cost of elementary and secondary education in the United States rose by 81 percent but student achievement was nearly flat.² In developing countries, interest in productivity often stems from a recognition that budgets for local public goods are necessarily small, so that efficient provision is especially valuable.

Yet, economic theory on local public goods has focused almost exclusively on allocative efficiency—how much of the local public good should each person pay for and receive? Even the most sophisticated literature about vouchers and decentralization commonly assumes that all spending on local public goods is equally productive, regardless of the system by which the money is raised and distributed. The almost exclusive focus of the local public goods literature on allocative efficiency is out of step with the need for theory that balances productivity and allocative efficiency issues. This paper tries to redress that imbalance.

¹ See Hoxby (1994, forthcoming 1999), Silanes-de-Lopez, Shleifer, and Vishny (1997), and Peltzman (1993,1996).

² The source is the Digest of Education Statistics, National Center for Education Statistics, United States Department of Education. Student achievement is assessed by the National Assessment of Education Progress testing program.

I build a model of productivity that emphasizes how imperfect information allows local public goods providers to take rent. I show that information is structured in such a way that an agency problem exists between the residents of a jurisdiction and their local public goods provider. Key features of the model are that (1) cost conditions vary among jurisdictions but are not fully observed, (2) some cost conditions are a function of local residents' own characteristics, (3) moving costs are high enough to keep households from moving frequently among jurisdictions, and (4) some qualities of local public goods are observable to consumers even though they are unverifiable (a contract cannot specify them sufficiently to allow a court to enforce them).

I show that a social planner who attempted to manage productivity in this setting would have to engage in costly information gathering. The productivity he would achieve would depend crucially on the information he was able to collect. I also show that, so long as households have a sufficient number of jurisdictions to choose among, conventional local property tax finance attains similar productivity as a social planner who has information that is so generous as to seem far-fetched: perfect information about quality, the specification of the cost function, and local residents' characteristics. The key insight is that a Tiebout residential market changes the information structure by effectively making demand information verifiable and eliminating the need for information about cost conditions that are a function of residents' own characteristics. Moreover, local property tax finance manages productivity in a completely decentralized way, requires only a small percentage of households to move exogenously at any given time, and does not depend on details of the local political process of determining property tax rates (such as whether there is a referendum or vote by representatives). The most surprising result, overall, is that conventional local property tax finance has a number of qualities that we would want if we were to *design* a mechanism to manage productivity.

A few points about the paper are worth emphasizing. It is not about allocative efficiency. The bulk of the local public goods literature studies allocative efficiency and assumes away productivity issues: the procedure is reversed here. It may be useful to think of allocative efficiency problems as having been pre-solved, perhaps with lump-sum redistribution. Ultimately, social welfare calculations should take account of

both the allocative efficiency and productivity consequences of each system of financing local public goods. Second, when I consider what the Tiebout process can do, I consider a housing market with many jurisdictions and few interjurisdictional externalities. This description of reality is clearly better for some metropolitan areas than for others. Boston has 70 jurisdictions within a 30-minute commute of downtown and 130 jurisdictions in its metropolitan area, while Miami has only one jurisdiction in its metropolitan area. I consider a market where the Tiebout process hat works well in order to show what information the process *can* generate with local property tax finance. I thus create a standard for considering how productivity is affected by departures from a Tiebout market, such as consolidation of jurisdictions, centralization of public goods finance, or vouchers. Finally, for simplicity, there is only one local public good in the model. I consistently use elementary and secondary schooling as the example of that good. I chose schooling because it is the most common target of proposed reforms to local public goods provision and because it provides accessible, realistic examples.

II. Productivity in the Local Public Goods Literature

Why has productivity been studied so little in the local public goods literature? Is productivity a non-issue--that is, are all public good providers equally productive regardless of how they are financed and how their area's Tiebout process functions? Is managing productivity purely a political problem? The answers to these questions clarify the importance of information and mobility.

By proposing a mechanism that made choice over public goods mimic choice over private goods, Tiebout's 1956 paper effectively made the theory of the consumer a tool for managing local public goods. Without the Tiebout mechanism, local public goods can be managed only by politics, which are the means by which non-local public goods are managed. Indeed, there is a substantial literature about how politics manage allocation and productivity issues in public goods.³ But allocation issues in *local* public goods are widely

³ This literature includes seminal work by Niskannen (1971), Brennan and Buchanan (1977, 1978), and Romer and Rosenthal (1978). Inman (1987) presents a detailed survey of the literature.

understood to be managed by the Tiebout process or by interaction between the Tiebout process and local politics.⁴ There is no similar understanding that productivity is managed by the Tiebout process, or interaction between the Tiebout process and local politics. The local public goods literature tends to assume either that a dollar spent always translates into a dollar of value or that politics are the only means by which productivity can be managed (as though local public goods were no different from other public goods).⁵

Tiebout characterized local public goods providers as cost-minimizing entrepreneurs, and much of the literature has treated this characterization as an *assumption*.⁶ It is not a mere assumption, however, but the logical outcome of the Tiebout process combined with the assumption that local public goods are either verifiable or repeatedly purchased at short intervals. These assumptions are reasonable if we consider local clubs, one of the popular early applications of Tiebout's paper. If we envision local public goods producers as swimming pool providers who sell memberships, then people might reasonably have an annual opportunity to re-purchase and/or club quality might be specified well enough to prevent serious problems of non-verifiability. In such a case, the Tiebout process could generate both efficient allocation of swimming pools and cost minimization by pool providers. A provider earning positive rent would be replaced by another who offered marginally less expensive memberships.

The important local public goods are not, however, swimming clubs but schooling, policing, and similar services. For these goods, the entrepreneurial model is unrealistic. Technology dictates that they be

⁴ Recent work on allocation in local public goods consistently emphasizes the interaction between the Tiebout process and politics. Examples include Epple and Romano (1998), Nechyba (1997), and Fernandez and Rogerson (1996). All of these papers assume, however, that a dollar spent on local public goods is equally productive regardless of the mechanism for allocating and financing local public goods.

⁵ There are some existing models, described by Rubinfeld (1987), that relate local jurisdictions to productivity, but they do so via economies of scale. Eberts, Schwartz, and Stone (1990) present evidence that economies of scale are exhausted once jurisdictions reach a size that is small for metropolitan jurisdictions. Stiglitz's well-known (1977) model is an economies of scale model, since the crucial (but unrealistic) assumption is that people must work in the jurisdiction where they consume local public goods.

⁶ See Rubinfeld (1987), p 575.

closely attached to residential location, and moving costs prohibit annual re-purchase. Indeed, the technology of schooling is such that, even in the absence of financial moving costs, a family could not frequently re-purchase without jeopardizing their children's schooling. Moving costs by themselves do not, however, preclude entrepreneurial producers. Real estate developers are entrepreneurs who deal with the same purchase patterns. Why are there no local public goods entrepreneurs who buy up jurisdictions with overpriced local public goods, improve quality for cost, and then re-sell the houses, extracting the consumers' full valuation of the improved conditions? The answer is that only a minority of the characteristics of local public goods are verifiable, so that even the best-specified contract could not prevent substantial variation in quality over the relevant purchase period. Indeed, educators are aware that quality goes far beyond the verifiable characteristics; they sometimes call the verifiable characteristics "custodial care" (children being in classrooms with teachers for a certain number of hours each day, grades and attendance records being kept, and so on). Nevertheless, many aspects of school quality that are difficult or impossible to verify are not difficult to observe. For example, there is evidence that parents and teachers can identify the teachers who, when students are tested, turn out to have the highest value-added.⁷

In short, Tiebout's entrepreneurial model does not work because quality is largely non-verifiable over the relevant purchase period. Yet, the entrepreneurial model demonstrates that the Tiebout process has implications for producers as well as consumers, and consumers' ability to *observe* local public goods has never really been doubted. This paper attempts to salvage the Tiebout insight while admitting realities that make the entrepreneurial model fail.

This paper owes a debt to the literature on regulation, especially Laffont and Tirole's (1993) work on optimal incentive mechanisms for regulated providers who operate in conditions of imperfect information. The paper also owes much to Epple and Zelenitz (1983), who ask whether politics are redundant or a mere second

⁷ See Black (1998).

line of defense since a well-functioning Tiebout process makes optimal allocation decisions. They show that politics play a meaningful role even when the Tiebout mechanism is functioning perfectly because misallocation and rent extraction can occur so long as the supply of land to particular jurisdictions is inelastic.⁸ Their argument is similar to the argument (made above) that incumbent owners are vulnerable to rent extraction. It is useful, however, to think about incumbents as former buyers rather than as inelastic suppliers because the previous buying helps us recognize that information asymmetries are what allow rent extraction to occur. If buyers could assess providers' future opportunities for rent extraction at the time they made their purchases, they would simply deduct the present value of future rents from the price they paid and providers would earn zero rents.⁹

III. The Foundations of the Model: Households, Jurisdictions, Agents, and Costs

Apart from crucial assumptions about information and mobility, the model employs simplifying assumptions that are common in the Tiebout literature on allocative efficiency.

Households

Each household buys a home and consumes one unit of the local public good available in its

⁸ In much recent literature, Epple and Zelenitz (1981) is implicitly used to justify examination of the interaction between the Tiebout process and politics. Schizophrenically, the same literature tends to assume that conditions exist that would allow the Tiebout process alone to make optimal allocation decisions. Epple, Filimon, and Romer (1983,1984), for instance, show the existence of equilibria when consumers can relocate among jurisdictions costlessly and continuously and when the preferences of the median voter in each jurisdiction determine the level of the local public good.

⁹ The distinction between this paper and papers on "taming the leviathan," is instructive. Such papers examine whether a tax regime can cause a government to over- or underprovide local public goods. Rauscher (1997), for instance, demonstrates that if local jurisdictions tax factors that are mobile among jurisdictions (capital) instead of immobile (labor in a world with no migration or land), then a revenue-maximizing government spends less on social capital that will increase the future tax base since it cannot prevent its future tax base from moving away. Rauscher's example is subsidized higher education. In a more general though less transparent model, Glaeser (1995) shows the same result in reverse. He demonstrates that if local governments tax factors that are immobile among jurisdictions (land) instead of mobile (income), then amenities that last many periods are more likely to be provided by local government. In short, Rauscher and Glaeser are concerned with allocative efficiency.

jurisdiction. In the case of schools this means that each household has a school-aged child. At any given time, a small minority of households randomly moves for exogenous reasons so that their moving costs are sunk relative to their choosing a jurisdiction. For all other households, the cost of choosing a new jurisdiction is the cost of moving, which is assumed to be high. Households can be classified by type; within each type, households have roughly the same income, preferences, and children--in the sense that their children cost a similar amount to educate and have similar achievement if given identical schooling. Types are indexed by i . (Both discrete and continuous types of households are common in the Tiebout literature. Continuous types are interesting in allocative efficiency models, but they are unrealistic and introduce unnecessary complications into the study of productivity, where the question is how to get multiple providers who serve similar households to be productive.¹⁰)

Households maximize expected utility, and their utility depends on local public goods quality, services generated by property (land and housing), and other consumption:

$$U_i(q, w, x) \tag{1}$$

where q is local public goods quality, w is property services, x is other consumption, and $U()$ has typical properties.¹¹ The utility function is modified by i to indicate that preferences vary by type. The budget

¹⁰ Epple and Platt (1998) describe households who vary continuously, both in their preference for schooling and in income. If types are continuous, there are always many more types than jurisdictions, and allocative efficiency results naturally focus on how stratification (the continuous version of segregation) occurs and how fiscal externalities created by different households being the same jurisdiction are internalized. Hamilton (1976) presents a model in which systematic inter-household fiscal externalities are internalized by capitalization.

If we are not interested *per se* in such problems, continuous types are neither interesting nor realistic. Households do not observe extremely fine distinctions in other households' types, as they would have to if we were to take stratification results literally. Households actually fit other households into broad categories, based on objective information (income, professional/blue-collar occupation, college education, and so on) and more subjective observations. Thus, in a metropolitan area with many jurisdictions, a newcomer is often given a list of "suitable" jurisdictions by friends or colleagues.

¹¹ The utility function satisfies:

$$\frac{\partial U_i}{\partial x} > 0, \frac{\partial^2 U_i}{\partial x^2} < 0, \frac{\partial U_i}{\partial q} > 0, \frac{\partial^2 U_i}{\partial q^2} < 0, \frac{\partial U_i}{\partial w} > 0, \frac{\partial^2 U_i}{\partial w^2} < 0 \quad \forall i.$$

constraint is specified below for alternative methods of financing local public goods.

Jurisdictions

In each housing market (metropolitan area), there are many jurisdictions (public school districts), each of which provides the local public good (schooling). Jurisdictions have fixed boundaries, and there are multiple jurisdictions for each type of household. The supply of potential jurisdictions is elastic (subject to a caveat described below), but I use this assumption only to ensure that the supply prices of land and houses (that is, the part of property prices that are independent of local public goods) are set by alternative uses of land, building materials, and builders. One can think of potential jurisdictions being agricultural or commercial but convertible to residential use if the land owners are paid for their land and builders are paid to build houses.

Agents

There are an arbitrarily large number of risk-neutral agents who can offer to provide local public goods for a jurisdiction. I follow numerous other authors in assuming that agents are rent-maximizing and can take rent in various forms: reduced effort, improved job benefits, salary in excess of that at which they are willing to supply their labor, gifts from people who benefit from their patronage, and so on. An agent who is not working for a jurisdiction has outside job opportunities that give him zero rent (in other words, that pay him the wage at which he is willing to supply his labor, but no more), so an agent will quit if a jurisdiction offers him a contract on which he would earn negative rent. I treat agents as though they provided public goods single-handedly, although they are, in reality, managers. An agent's rent is given by

$$R = b - C - \varnothing(e) \quad (2)$$

where R is rent, b is the budget or payment he receives for the local public goods, C are the costs he pays to inputs that go into producing the local public good, and e is his effort. Examples of effort on the agent's part are monitoring of employees and hard bargaining with (or shopping among) input suppliers. $\varnothing(e)$ is the money value of the disutility of effort, which is non-negative for all non-negative levels of effort and increases in effort at an increasing rate: $\varnothing'(e) > 0$ and $\varnothing''(e) > 0$.

Costs

Per-person costs, C , are a function of local public goods quality (q), the agent's effort (e), the type of households (i) being served, and cost conditions that are specific to the jurisdiction ($\acute{a}$):

$$C(q, e, i, \acute{a}) . \quad (3)$$

All else equal, per-pupil costs increase in quality at an increasing rate:

$$\frac{\partial C}{\partial q} > 0 \text{ and } \frac{\partial^2 C}{\partial q^2} > 0 . \quad (4)$$

Per-pupil costs are decreasing in effort, and extra effort makes supplying the marginal unit of quality less costly:

$$\frac{\partial C}{\partial e} < 0 \text{ and } \frac{\partial^2 C}{\partial q \partial e} < 0 . \quad (5)$$

The type of households being served affects per-pupil costs because one of the most common claims about public goods provision is that different groups of people are differentially expensive to serve. For instance, it may be more expensive to supply a given quality of schooling to children from poor households because their parents are less likely to provide them with inputs that are complementary to schooling (supplies, nutrition, quiet study space, a computer, and so on). It is costly for a school to supply substitutes for these inputs.

Finally, I allow per-pupil costs to be a function of idiosyncratic cost conditions in each jurisdiction--that is, cost conditions that are specific to a jurisdiction rather than a function of its residents' characteristics. For instance, a jurisdiction might have teachers who are idiosyncratically good given the wages they are paid, geography that makes school transportation cheap, or buildings that are inexpensive to maintain.¹² A higher value of $\acute{a}$ is assumed to increase costs and make supplying the marginal unit of quality more costly:

$$\frac{\partial C}{\partial \acute{a}} > 0 \text{ and } \frac{\partial^2 C}{\partial q \partial \acute{a}} > 0 . \quad (7)$$

¹² More precisely, costs are idiosyncratic if they are not a *predictable* function of residents' characteristics. If one family had a child whose special education program made him extraordinarily expensive to educate, the associated costs would be idiosyncratic.

Cost conditions, $\acute{a}$, are distributed among jurisdictions with density $f(\acute{a})$ over the interval $[\acute{a}, \bar{\acute{a}}]$. Regardless of how households move among jurisdictions or which potential jurisdictions engage in local public goods provision, the distribution of cost conditions across the total population of households *remains the same*. This assumption ensures that we focus on how the agents' choices affect productivity, rather than how people discard current jurisdictions in favor of potential jurisdictions that are idiosyncratically low-cost.

Per-pupil costs, C , are not a function of jurisdiction size--that is, there are no economies of scale. This simplifying assumption can be relaxed easily since the property market would ensure that predictable economies of scale would be capitalized into a jurisdiction's property prices.¹³

IV. The Information Environment

Information is usually defined as verifiable if a contract citing the information could be legally enforced. Underlying this definition, however, is the notion that a piece of information is verifiable if it is *socially knowable*--that is, if the information is observed with certainty by social planners or if the public all perceive that information identically. I use the terms "verifiable" and "socially knowable" interchangeably.

I define basic provision to be the aspects of local public goods that are readily verifiable. In the case of schools, for instance, basic provision would include custodial care. The cost of basic provision is a constant in the cost function that is common to all jurisdictions that serve the same type of household. Quality is defined to be the aspects of local public goods that exceed basic provision.

A household is assumed to know its own type and to be able to tell whether other households of its type reside in a jurisdiction. Households' types are not, however, assumed to be socially knowable: there is no definitive census of households' preferences and the costs of educating their children. Although social planners

¹³ If everyone knew that there was a function that related jurisdiction size to per-pupil costs, then economies-of-scale would be capitalized or re-districting might occur. If, on the other hand, economies of scale were so complex that they appeared to be idiosyncratic, they would be included in $\acute{a}$.

may form estimates of a households' types at a cost, even costly estimates will not be fully informative. For instance, real policy makers use decennial census data to estimate median household income by jurisdiction, but such measures are crude indicators of preferences and educational costs.

Households are assumed to observe a variety of outcomes that result from local public goods provision, some of which can be objectively and others only subjectively measured. For instance, parents comparing schools consider a wide range of outcomes: test scores, grades, work habits, extracurricular activities, attitudes towards further education, and other behavior. Outcomes are not merely a function of quality, however, they are also a function of a household's own characteristics. Students' test scores, for example, are partly determined by school quality and partly determined by the resources they have at home. Since households are assumed to be able to recognize other households of their own type, they can compare the relative quality of jurisdictions that serve households of their type by comparing outcomes. Thus, local public goods quality is observable to households within jurisdictions that serve their type.

Local public goods quality is not verifiable, however, for several reasons. First, many outcomes are inherently subjective and cannot be consistently measured, as a social planner would require. Second, household types may directly affect outcomes, and a social planner cannot partial out the effect of quality without knowing all households' types. Social planners are assumed to be able, however, to estimate some objective outcomes at a cost. He can also pay for estimates of how types affect the estimated, objective outcomes. Even using these costly estimates, a social planner can partially assess quality--both because his estimates are inaccurate and because he has no measures of subjective outcomes.¹⁴

The budgets that agents receive to provide goods in a jurisdiction are observable and verifiable. Households and social planners, however, are assumed to be unable to observe agents' efforts or jurisdictions'

¹⁴ The paragraph refers to social planners, but households are also poor judges of quality when comparing jurisdictions that serve households of very different types. The error in a household's assessment of a jurisdiction increases as the households served that are further away (in type space) from the household's own type.

idiosyncratic cost conditions.¹⁵ For instance, local residents have to take their agent's word on how much effort he puts into recruiting and screening teachers. If he is outside a jurisdiction, an agent is no more able to observe its idiosyncratic cost conditions than households or social planners. Once an agent is commissioned to provide local public goods to a jurisdiction, however, he learns its idiosyncratic cost conditions and takes them into account when choosing quality.¹⁶

In order to preserve the integrity of the information assumptions, I assume that idiosyncratic costs change "often enough" that they do not become socially knowable. This assumption excludes strategies whereby, for instance, society pursues one method of local public goods finance purely to find out what idiosyncratic costs are so that it can then pursue another method of finance (that presumably requires better information to work well) with better information. In the same spirit, I assume that all people of the same type chose to buy the same amount of property services, w , regardless of the jurisdiction they live in. That is, idiosyncratic costs are assumed to change "often enough" that people do not, say, enlarge their houses in response to their local agent's actions. Below, the choice of w is specified more exactly.

A brief discussion of centralized finance rules illustrates how productivity is affected by the information environment described. Productivity management requires simultaneous downward pressure on costs and upward pressure on quality. Centralized finance rules get caught between these competing aims, and the tension can be relaxed only by more information.

Centralized rules for financing local public goods usually start from one of two bases, costs or outcomes.¹⁷ Cost-based rules use the production technology that agents report. For instance, cost-based school

¹⁵ This is one condition, not two. If effort were known, then idiosyncratic cost conditions would be known, and *vice versa*.

¹⁶ An agent who is in a jurisdiction that allows him to take positive rent will not attempt to shop around for better rent opportunities. The "market for lemons" logic applies, and the market for positive rent jobs disappears.

¹⁷ The rules I consider are not merely illustrative; they are used in some American states and advocated for many others. California and New Mexico have centralized finance rules built purely upon a cost basis. See *Public School*

finance rules allocate money to school districts based on how many teachers they have at each level of experience, and so on. Cost-based rules are poor at managing productivity because they give no incentive to agents to minimize costs for any given level of quality. If it is cost minimizing to have fewer, but better-paid teachers, an agent has no incentive to use this technology. Cost-based rules also give agents no incentives to provide quality, except when it occurs as a by-product of the inputs for which they are compensated. An agent, for example, has no incentive to find the best teachers among the pool of applicants attracted by the salary he offers. Social planners can only improve cost-based rules by paying for better estimates of the production function and better information on household types and outcomes.

Outcome-based rules are built on the outcomes agents report. For instance, school districts might be given a basic budget plus incentive payments determined by the achievement of their students. Outcome-based finance rules are considered often, but they are enacted only occasionally and are usually short-lived. Why? Consider outcome-based rules for schools, which--in order to be forceful and inexpensive to administer--typically employ outcomes like students' scores on state-wide, multiple-choice tests. Such rules encourage agents to sacrifice outcomes that are not tested to outcomes that are. Schools may reduce time devoted to skills like essay-writing, neglect subjective outcomes, and overinvest in test-taking tactics (including tactics like getting test exemptions for students who are likely to perform poorly). Moreover, a rule that provides good incentives for some agents is likely to provide poor incentives for others. An agent who serves students who are expensive to educate may have no incentive to provide quality because the slope of the reward function is too small to compensate him for the effort and expense involved in raising his students' outcomes. The same rule may unduly reward agents who serve students who are cheap to educate. Social planners can improve outcome-based rules only by investing in more information on objective outcomes, household types, and the relationship between household types and outcomes. Once outcomes are announced, however, even well-

Finance Programs of the United States and Canada. Clotfelter and Ladd (1996) discuss some centralized finance rules built upon performance.

constructed rules can become politically unsustainable. It is politically awkward to direct *less* money to students whose achievement is worse or to assign "handicaps" to household types (because they amount to announcing that poor children, for instance, are *expected* to have lower test scores).

V. Local Property Tax Finance

How does local property tax-based finance with a Tiebout residential market manage productivity under the same assumptions? Consider the problems solved by moving households, non-moving (incumbent) households, and agents.

Moving Households

Moving households equilibrate the property market, taking local public goods quality into account. They choose property in a jurisdiction, taking property prices, property services, tax rates, and the quality of local public goods qualities as given. Thus, their problem is to

$$\max U_i(q_j, w_j, x) \quad \text{subject to} \quad (1 + t_j)p_j^w w_j = y_i - x \quad (8)$$

over choices of x and jurisdictions indexed by j . That is, they choose among combinations: $[(q_1, w_1), \dots, (q_j, w_j)]$.

p_j^w is the "annual price" of property--in other words, the annual payment a household makes on a property.

The property tax rate, t_j , is defined on the annual price of property.¹⁸ It is convenient to define a gross-of-tax annual price of property, $p_j = (1 + t_j)p_j^w$. Define the indirect utility function:

$$V_i(q_j, p_j, y_i) = U_i(q_j, w_j(p_j, y_i), y_i - p_j w_j(p_j, y_i)) \quad (9)$$

Having stuck closely to the literature on allocative efficiency, we can invoke a number of standard allocation

¹⁸ Everything in the equation must be defined over the same period, and a year is the most convenient period. t_j is the property tax rate on the annual property cost, where a typical statutory property tax rates is defined on the asset price of a property. If, for example, the annual property cost were the asset price times the interest rate (as it would be if the property had an infinite lifespan), then t_j would be equal to the statutory property tax rate divided by the interest rate. A statutory property tax rate that is common where property taxes provide the majority of local public financing is 20 mils (0.020); this would be equivalent to a t_j in the range of 0.2 to 0.33.

results. Jurisdictions will be segregated by household type.¹⁹ Any two jurisdictions that serve the same type of households offer packages between which the households are indifferent. This gives us an arbitrage condition for jurisdictions, j and $j+1$, serving same-type households:

$$V_i(q_j, p_j, y_i) = V_i(q_{j+1}, p_{j+1}, y_i) . \quad (10)$$

Jurisdictions that serve households of different types offer packages such that households are not attracted to the packages offered by jurisdictions serving another type. If jurisdiction j serves type i , and jurisdiction $j+2$ serves type $i+2$, then:

$$V_i(q_j, p_j, y_i) > V_i(q_{j+2}, p_{j+2}, y_i) \quad \text{and} \quad V_{i+2}(q_j, p_j, y_{i+2}) < V_{i+2}(q_{j+2}, p_{j+2}, y_{i+2}) . \quad (11)$$

The property market is in equilibrium when property prices (p_j^w) have adjusted so that equation (10) and the inequalities in (11) hold simultaneously with the first-order conditions of the moving households' problem. The key condition that must hold in equilibrium is:²⁰

$$\left. \frac{dp}{dq} \right|_{V_i=\bar{V}_i} = - \frac{\frac{\partial V_i(q, p, y)}{\partial q}}{\frac{\partial V_i(q, p, y)}{\partial p}} . \quad (12)$$

Equation (12) indicates that, among jurisdictions that serve type i , the gross-of-tax price reward for an increase in the quality of local public goods must be equal to the increase in indirect utility (for households of type i) from that increase in quality. That is, the actions of moving households on the property market guarantee that

¹⁹ Stricter versions of these results are generally referred to as the "stratification" and "increasing bundles" results. See Epple and Platt (1998) for illustration of these results in a model that, like this paper's model, allows both preferences and incomes to vary among households. They assume, however, that preferences and incomes are continuously distributed and, as a result, need stronger assumptions to generate stratification and increasing bundles than are necessary in this model.

²⁰ The other two conditions are:

$$\left. \frac{dp^w}{dy} \right|_{V_i=\bar{V}_i} = - \frac{\frac{\partial V_i}{\partial y}}{\frac{\partial V_i}{\partial p}} \quad \text{and} \quad \left. \frac{dq}{dy} \right|_{V_i=\bar{V}_i} = - \frac{\frac{\partial V_i}{\partial y}}{\frac{\partial V_i}{\partial q}} .$$

house prices reflect changes in households' valuation of quality.

Although the arbitrage condition is in terms of the gross-of-tax price, the actions of households on the property market affect only the net-of-tax price. A more intuitive version of equation (12) is, therefore:

$$dp^w = - \frac{\frac{\partial V_i(q,p,y)}{\partial q}}{\frac{\partial V_i(q,p,y)}{\partial p}} \cdot \frac{1}{1+t} \cdot dq . \quad (13)$$

Figure 1 shows indifference maps between p^w and q for two different types of households. All jurisdictions that serve the same type of household must offer packages that lie on the same indifference curve.

A convenient implication of the segregation result is that, henceforth, we need follow what happens to only one type of household in order to understand productivity. For each type of household, there will be several jurisdictions, the offerings of which will be constrained, in equilibrium, by equation (13), the inequalities in (11), and the first-order conditions in footnote 22. Because we are interested in productivity and not allocation, the model will hinge upon competition among jurisdictions serving the same type of household, and--once we have determined that jurisdictions will be segregated and ordered--we can focus on one household type at a time.

Non-Moving Households, Part I

The role of non-moving households is to choose the property tax rate, by means of a local political process. Here, the difference between incumbents and moving households becomes important. Moving households, who do not yet own property in the district they are moving to, care only about local public goods quality and gross-of-tax property prices. But, incumbents, whose wealth is invested in their property, consider the impact of tax rates on net-of-tax property prices. Since incumbents are always in the majority, their preferences will determine the outcome of conventional political processes (referendums, representative governments, *et cetera*) that determine tax rates. Thus, we only need to consider incumbents' preferences over

property tax rates.

A household knows that its jurisdiction will contain households of its own type, knows its own preferences and income (and, thus, the preferences and income of its co-residents), but does not know whether its jurisdiction has idiosyncratically high or low costs. It has to choose an appropriate tax rate. I first show how incumbents would choose the tax rate if they thought they were in a jurisdiction with the worst possible cost conditions: $\acute{a}=\tilde{a}$. Later, I show why incumbents are very likely to stick to this choice.

In a jurisdiction with the worst possible cost conditions, property prices will serve only the role of clearing the markets for the supply of land and housing. That is, property prices will not capitalize local fiscal externalities such as better cost conditions. Let $\bar{p}^w$ represent the supply price of property. Suppose, for a moment, that households knew the form of the cost function for agents with $\acute{a}=\tilde{a}$. They would then know the price of property ($\bar{p}^w$), the prices of other goods, and the shape of the zero-rent iso-rent curve for $\acute{a}=\tilde{a}$. Armed with this information, households could make a conventional consumption decision, determining how much property and quality to demand ($\bar{w}_i^*$ and $\bar{q}_i^*$) and what total payment to make for local public goods. They would not so much chose a tax rate as *back out* the tax rate needed to generate their desired total payment from the property they purchased.

Figure 2 illustrates the solution. ID_i is an indirect indifference curve that shows the trade-off between quality and total payments for quality. So long as households back out the tax rate that generates their desired total payment from $\bar{p}^w \bar{w}_i^*$, the residential market will actually set $p^w=\bar{p}^w$ and the agent's budget will be equal to the total payment households are willing to make for the quality he offers.

But, can households approximate the zero-rent iso-rent function (at least in the neighborhood of point A in figure 2)? Households have difficulty approximating rent functions because, in general, they cannot distinguish between true and false information in agents' behavior. Agents with good cost conditions might mimic the behavior of agents with $\acute{a}=\tilde{a}$ in order to extract rents. Fortunately, so long as agents earn zero rents

in alternative jobs, their quitting reveals true information, especially about the cost function for $\acute{a}=\tilde{a}$ jurisdictions. If households offer contracts that state, "produce at point X or get fired," and point X is below the zero-rent iso-rent curve, an agent will quit. For any given contract, quits will occur in $\acute{a}=\tilde{a}$ jurisdictions if they are going to occur anywhere.

Furthermore, utility is very low if an agent quits since no local public goods provision (no even basic provision) will occur. Thus, households will be wary of offering contracts that previously caused a quit, and experience of quits will put a floor on the tax rates incumbent households choose.

Agents

Each agent chooses what quality to supply and what effort to make, given the cost conditions he gets and the incentives he is given by the budget-setting process. Each agent faces a per-pupil budget rule:

$$b_j = t_j p_j^w w_j \quad (15)$$

which states that each agent receives the revenue from the local property tax as his budget. He takes as given the idiosyncratic cost conditions of the jurisdiction, the household type his jurisdiction serves, the property tax rate, and local property (w_j). The agent, however, recognizes that property prices respond predictably (that is, according to equation (13)) to changes in quality at a given tax rate.²¹ In short, he solves:

$$\max_{q,e} R = t_j p_j^w w_j - C(q, e, i, \acute{a}) - \phi(e) , \quad (16)$$

subject to:

$$dp^w = - \frac{\frac{\partial V_i}{\partial q}}{\frac{\partial V_i}{\partial p}} \cdot \frac{1}{1+t} \cdot dq .$$

The first-order conditions are:

²¹ Property prices are also affected by many factors (such as the state of the local economy) that the agent does not control. So long as these factors are exogenous to his decisions, they introduce noise but do not change his maximization problem. If agents were risk averse, they would have to be compensated for assuming the risk associated with taking up a local public goods provision job.

$$tw \frac{dp^w}{dq} - \frac{\partial C(\dots, i, \acute{a})}{\partial q} = 0 \quad \text{or} \quad \frac{\partial C(\dots, i, \acute{a})}{\partial q} = - \frac{\frac{\partial V}{\partial q}}{\frac{\partial V}{\partial p}} \frac{t}{1+t} w, \quad (17)$$

and

$$\frac{\partial C(\dots, i, \acute{a})}{\partial e} = -\phi'(e). \quad (18)$$

The former first-order condition indicates that agents set their marginal cost of producing higher quality equal to the marginal revenue they earn by producing higher quality. The latter first-order condition indicates that agents set the reduction in costs from higher effort equal to the reduction in money utility from higher effort. Let $q^*(t, w, i, \acute{a})$ and $e^*(t, w, i, \acute{a})$ be the agent's solution to the problem given by (16).

Figure 3 is like figure 2, but it also shows the budget-quality curve generated by moving households and it shows iso-rent curves for agents whose costs are better than $\acute{a} = \tilde{a}$. The budget-quality curve (denoted $b_i(q)$) begins at point A, and thereafter follows equation (13). Since the curve is produced by the arbitrage of moving households, it looks like an indifference curve, except that it is scaled by t and w . It begins at point A because incumbents choose t so that the agent with $\acute{a} = \tilde{a}$ is on his zero-rent iso-rent curve. Agents who get better cost conditions (lower $\acute{a}$) have flatter iso-rent curves and choose to offer higher quality. This result comes from the first-order condition given by equation (17), since $\partial C / \partial q$ is increasing in $\acute{a}$ and utility is concave in q .²² In figure 3, agents with cost conditions $\tilde{a}$, $\acute{a}'''$, $\acute{a}''$, $\acute{a}'$, and $\acute{a}$ (where $\tilde{a} > \acute{a}''' > \acute{a}'' > \acute{a}' > \acute{a}$) chose increasingly high levels of quality. Only agents with $\acute{a} = \tilde{a}$ are earning approximately zero rent. For instance, if an agent with $\acute{a} = \tilde{a}$ can earn zero rents at point A, then an agent with $\acute{a} = \acute{a}$ could earn positive rent there. So, the agent with $\acute{a} = \acute{a}$ must be earning even higher, positive rent at point B.

There are some parts of the budget-quality curve where no agents will actually produce. These are parts where even the lowest cost agent cannot get sufficient compensation to cover his increases costs of producing higher quality. For instance, in figure 3, no jurisdiction will consume a budget-quality package that

²² Each agent faces a budget-quality curve that is a function of both the preferences and incomes of the type of households he serves. If the households place a high marginal value on quality, then $S(\partial V / \partial q / (\partial V / \partial p))$ will be high. If the households care a lot about quality for people of their incomes, then t will be high. An agent who faces households who do not value quality much will produce little quality even if his cost conditions are good.

is above point B, where the iso-rent curve of the $\hat{a}=\hat{a}$ agent is tangent to the budget-quality curve.

Non-Moving Households, Part 2

Do non-moving households want to choose a tax rate different from t_i^* given that there is a good probability that they will not end up in an $\hat{a}=\tilde{a}$ jurisdiction? The answer to the question is "possibly," but the fact that moving households make property prices reflect differences in agents' productivity eliminates much of the need for tax rate changes. In other words, the housing market produces automatic incentives for agents with good cost conditions to produce higher quality rather than take rent by pretending to have worse cost conditions. Thus, changes in the tax rate need not be the means by which households give agents incentives. In fact, much of the task of rewarding agents is taken out of incumbents' hands so that they do not have to rely on political processes to accurately reflect their desires.

t_i^* is the factor by which increases in indirect utility are scaled down into financial incentives for the agent. If an agent contemplates raising quality, he can expect total property wealth to rise according to equation (13) and can expect his budget to rise by the tax rate times the increase in property wealth. Although households who are actually in $\hat{a}=\tilde{a}$ jurisdictions will be best off with tax rate t_i^* , it does not necessarily give the optimal marginal incentive to agents. Figure 4 illustrates an example. It shows iso-rent curves just for $\hat{a}=\tilde{a}$ and $\hat{a}=\hat{a}'$. Incumbent households in jurisdictions with $\hat{a}=\hat{a}'$ consume at point C. The double crossing of the iso-rent curve and the indifference curve of incumbent households (enclosing the shaded area) makes it clear that the incumbents would have preferred a progressive tax schedule, where tax rates were higher than t_i^* for property prices higher than $\bar{p}^w$.

Note the contrast between moving and incumbent households. Moving households are all on the indifference curve given by ID_i^0 , so they only care that $t_i=t_i^*$ in the worst-off jurisdiction. Incumbents can have capital gains from appreciation of their property, so they might want a more complicated tax schedule. They do not, however, have the information about the functional forms of agents' rent functions that would allow

them to compute such a schedule. In short, incumbent households will be wary of choosing a tax rate less than t_i^* for property prices up to $\bar{p}^w$ because of the possibility of a quit. They know that t_i^* will be optimal if the district has $\acute{a}=\tilde{a}$, and (without better information than they have) they cannot determine the right amount of progressivity or regressivity in the tax schedule for property prices higher than $\bar{p}^w$. They are likely to stick with t_i^* , partly for administrative simplicity.²³ If, however, jurisdictions were to experiment with progressive or regressive tax rates for property prices higher than $\bar{p}^w$, all of the discussion that follows would go through. I return to related issues below.

Summary

In summary, conventional local property tax finance in a market with a sufficient number of jurisdictions gives agents incentives to produce higher quality. Moving households' behavior drives property prices which, in turn, generate the budget-quality curve. Although only a fraction of households are moving at any given time, their observations of quality are universalized so that every household gives its agent financial inducements to raise quality. Agents' behavior generates reliable information about the cost function for jurisdictions with $\acute{a}=\tilde{a}$, and agents in those jurisdictions will be held to approximately zero rent. This zero rent condition limits rent for other agents because it pins down the budget-quality curve offered to agents. The budget-quality curve can offer positive rent to agents who have better cost conditions and choose to offer higher quality. However, their rent is limited because the budget-quality curve scales the increase in households' utility by the tax rate, ensuring that households keep most of the gains associated with fortuitous, idiosyncratic cost conditions.

VI. What A Social Planner Could Achieve with A Generous Amount of Information

How good is the outcome achieved by a local property tax finance? We can assess it by determining

²³ A more complicated schedule would require households to keep track of $\bar{p}^w$ by following the property market.

how much information we would have to give a social planner to enable him to achieve a similar outcome with centralized finance. I show that a social planner could achieve a similar outcome if he were to have *perfect*, *free* information on quality, on every household's type, on all outcomes (objective and subjective), on how much the agent was spending, and on the exact functional form of the rent and utility functions. The social planner could achieve the same outcome as local property tax finance if he were to have all of the above information *except* the exact functional form of the rent function. In fact, the only pieces of information I do not consider giving to the social planner are the agents' efforts and jurisdictions' idiosyncratic cost conditions.²⁴ In other words, the social planner is given, for free, all the information that he might attempt to *estimate* (for a cost) plus some information he could not plausibly obtain at any cost.

I define centralized finance in a non-degenerate way. The state collects and allocates revenue among districts on the basis of any verifiable characteristics *except* for property value-based measures of demand for the district. I uphold a meaningful distinction between local property tax finance and centralized finance and do not allow the state to be a mere pass-through for local finance. Moreover, to keep the discussion focused on productivity issues, I will not consider social planners' rules that redistribute income among household types.²⁵ Thus, the agent's social welfare maximization is separable in household types.

Since the social planner does not directly observe each jurisdiction's α , he cannot dictate to agents the effort, quality, and costs they should evince. Instead, he will offer agents a "menu" of budget-quality packages (that is, a budget-quality curve) from which they can choose. His solution will minimize rent among the incentive-compatible budget-quality curves, subject to the constraints that agents must be willing to participate (a zero-rent constraint) and that the social planner cannot force households to live in one jurisdiction if they

²⁴ As noted previously, this is one condition, not two.

²⁵ As noted in the introduction, allocative efficiency and productivity can logically be considered separately. Of course, if redistribution is desirable for allocative efficiency, it might be carried out by means (for instance, lump-sum redistribution of income) that are effectively prior to the problem under consideration.

want to move to another (a no-moving constraint). A curve is incentive compatible if it gives an agent who has a better $\hat{a}$ enough rent for offering higher quality that he prefers to offer higher quality rather than get whatever rent he could get by pretending to have a worse $\hat{a}$.

Figure 5 illustrates how an incentive compatible curve is mapped. The social planner starts at point A because agents in districts with $\hat{a}=\bar{a}$ cannot be induced to produce for negative rent and because, among the points on the zero iso-rent curve for the $\hat{a}=\bar{a}$ agent, point A maximizes household utility. An agent with $\hat{a}=\hat{a}'''$ could earn positive rent by producing at point A. Therefore, he can only be induced to provide higher quality if he is offered a package, like that marked with a star, that is on an iso-rent curve that is higher than his iso-rent curve that runs through point A. (He need only earn slightly more rent than he could earn at point A. The distance between the iso-rent curves is exaggerated for clarity.) Now consider an agent with $\hat{a}=\hat{a}''$. Not only could he earn positive rent by producing at point A, he could earn even greater rent by producing at the point marked by a star. Thus, he can only be induced to provide higher quality if he is offered a package, like that marked with a diamond, that is on a higher iso-rent curve than his iso-rent curve that runs through the point marked by a star. The mapping continues with agents who have $\hat{a}=\hat{a}''$ and $\hat{a}=\hat{a}$. A curve that contains all the incentive compatible packages is an incentive compatible menu for agents. It will be concave until the iso-rent curve of the lowest cost agent is reached. After that, it follows the curve of the lowest cost agent.

There are many budget-quality curves that are incentive compatible--after all, the points marked with a star and a diamond were arbitrarily chosen. But, different incentive-compatible curves allow agents to earn different amounts of rent. How does the social planner find the incentive-compatible curve that minimizes rent? He solves a truth-telling problem that produces an incentive-compatibility constraint that reveals how much an agent's rent must increase as his reported cost conditions fall so that he has an incentive to report his true cost conditions and make the effort appropriate for those cost conditions (rather than report worse cost conditions and get the boost to utility associated with making less effort). The social planner then minimizes

rent subject to the incentive compatibility constraint and the agents' participation constraint.²⁶

Suppose the social planner gives each agent a budget b on the basis of the cost conditions he reports. This defines a function $b(\tilde{a})$ where $\tilde{a}$ are the cost conditions the agent *reports*. The social planner can require an agent to produce quality $q(\tilde{a})$ (which is observable and verifiable to the social planner given the generous amount of information he has) if he takes budget $b(\tilde{a})$.²⁷ What functions $b(\tilde{a})$ and $q(\tilde{a})$ will make agents reveal their true $\hat{a}$? The agents maximize their objective function over possible reports:

$$\max_{\hat{a}} R = b(\tilde{a}) - C(q(\tilde{a}), e, i, \hat{a}) + \varnothing(e) . \quad (19)$$

Notice that, while $b(\tilde{a})$ is a function only of reported cost conditions, C is a function of both reported and true cost conditions. Effort is determined for any given combination of true and reported costs because there is only one effort that minimizes $C + \varnothing$ for any given $\tilde{a}$ and $\hat{a}$. That is,

$$\min_e C(q(\tilde{a}), e, i, \hat{a}) + \varnothing(e) \quad (20)$$

gives us:

$$\varnothing'(e) = \frac{\partial C(q(\tilde{a}), e, i, \hat{a})}{\partial e} \quad (21)$$

which implicitly defines a function $e_i(\tilde{a}, \hat{a})$.²⁸

Rewriting the problem with e as a function of $\hat{a}$ and $\tilde{a}$, we have:

$$\max_{\hat{a}} R = b(\tilde{a}) - C(q(\tilde{a}), e(\tilde{a}, \hat{a}), i, \hat{a}) + \varnothing(e(\tilde{a}, \hat{a})) . \quad (22)$$

We want to choose $b(\tilde{a})$ and $q(\tilde{a})$ functions such that the first-order condition is fulfilled when $\tilde{a} = \hat{a}$:

$$b'(\tilde{a}) - \frac{\partial C}{\partial q} q'(\tilde{a}) - \frac{\partial C}{\partial e} \frac{\partial e}{\partial \tilde{a}} - \varnothing'(e) \frac{\partial e}{\partial \tilde{a}} = 0 \quad \text{if } \tilde{a} = \hat{a} . \quad (23)$$

Equation (23) is the incentive compatibility constraint.

²⁶ I allow the social planner to handle the no-moving constraint separately. See below.

²⁷ Equivalently, we could require agents to have costs equal to $C(\tilde{a})$. Regardless of whether we specify $[b(\tilde{a}), q(\tilde{a})]$ combinations or $[b(\tilde{a}), C(\tilde{a})]$ combinations, effort will be a deterministic function of the agent's true cost conditions and reported cost conditions--and this is what we need in order to proceed.

²⁸ The function is type-specific because costs depend on the type of household being served.

At this point, two observations allow the problem to be simplified. First, if all households of a type must be indifferent among the packages offered to them, then maximizing their social welfare reduces to maximizing the utility of the households who would necessarily be the worst off under any incentive compatible menu--that is, the households who happen to end up in the jurisdiction with $\hat{a}=\bar{a}$. Second, because the incentive compatibility constraint guarantees that all other agents earn greater rents than unlucky agents who get $\hat{a}=\bar{a}$, the participation constraint reduces to a zero rent condition for the $\hat{a}=\bar{a}$ agents. In summary, the problem can be written as:

$$\min_{(b(\hat{a}), q(\hat{a}))} \int_{\hat{a}}^{\bar{a}} [b(\hat{a}) - C(q(\hat{a}), e(\hat{a}, \hat{a}), i, \hat{a}) - \theta(e(\hat{a}, \hat{a}))] dF(\hat{a}) \quad (24)$$

subject to:

$$b'(\hat{a}) - \frac{\partial C}{\partial q} q'(\hat{a}) - \frac{\partial C}{\partial e} \frac{\partial e}{\partial \hat{a}} - \theta'(e) \frac{\partial e}{\partial \hat{a}} = 0 \quad \text{if } \hat{a}=\hat{a},$$

and

$$R = b(\hat{a}) - C(q(\hat{a}), e(\hat{a}, \bar{a}), i, \bar{a}) - \theta(e(\hat{a}, \bar{a})) = 0 \quad \text{if } \hat{a}=\hat{a}=\bar{a}. \quad (25)$$

The solution to this problem turns out to be similar to the Tiebout-market solution. A closed-form solution cannot be given unless we specify simple utility and cost functions, but we can characterize the solution sufficiently for our purposes without such simplification.²⁹ The rent-minimizing, incentive-compatible curve begins at point A, and proceeds along the agents' iso-rent maps, always offering an agent just a bit more rent than he could obtain by producing at the point offered to the previous agent. Households of a given type are, however, indifferent among the points on this curve, so the social planner must offer packages like (b^*, q^*, f^*) , in which the agent receives a budget of b^* , must produce quality q^* , and the household pays a fee of f^* . Except in jurisdictions with $\hat{a}=\bar{a}$, the households pay fees that are greater than the budgets the agents receive.³⁰ The social planner accrues a kitty equal to the differences between the fees households pay and the budgets agents

²⁹ Laffont and Tirole show closed-form solutions for a somewhat similar problem. They use quadratic utility functions and cost functions of the form: $C=(\hat{a}+q-e) \cdot \text{quantity}$.

³⁰ In some cases, households in jurisdictions with $\hat{a}=\hat{a}$ will also pay a fee equal to their agent's budget.

receive. He can distribute this kitty to households in any way so long as the distribution does not make households want to move and provided the distribution is not endogenous to how his centralized rule works out. If he were to make the distribution endogenous, the rule would break down.³¹

VII. A Comparison of the Local Property Tax Solution and the Social Planner's Solution

Figure 6 superimposes the budget-quality curve from local property tax finance on the social planner's solution, so that the two solutions can be compared. The rent-minimizing incentive-compatible curve may lie above, below, or on the budget-quality curve. Since the curves may cross, I illustrate all three cases with a single rent-minimizing incentive compatible curve. If the budget-quality curve lies on the rent-minimizing incentive-compatible curve (except at the end where the incentive-compatible curve follows the iso-rent curve of the $\hat{a}=\hat{a}$ agent), then productivity is the same under local property tax finance and the social planner.³² If and where the budget-quality curve lies below the rent-minimizing incentive-compatible curve, the packages on the budget-quality curve are not actually produced by an agent. In other words, the financial incentives for raising quality are not large enough to cover the costs. If and where the budget-quality curve lies above the rent-minimizing incentive-compatible curve, households are limiting but not minimizing rent. Some of the money that the social planner puts into household hands (through the distribution of the kitty) is in agents' hands under local property tax finance.

In short, productivity difference between the local property tax solution and the social planner's solution are due to the fact we have assumed that the social planner knows the exact form of the agents' rent functions. If households had this information, they would use a complex tax schedule, with a tax rate of t_i^* for

³¹ In other words, he must commit to a distribution of the kitty *before* implementing his rule.

³² The only difference between the solutions would be the difference between the distribution of luck (under local property tax finance) and the distribution of the kitty (under the social planner). We should not make too much of this differences because luck, under local property tax finance, is uncorrelated with household type.

property prices up to $\bar{p}^w$ and different tax rates (either higher or lower, depending on whether the rent-minimizing incentive-compatible curve is above or below the budget-quality curve) for property prices above $\bar{p}^w$. Conversely, if the social planner did not know the exact functional form of the agents' rent functions, then he would do about as well as the households do in choosing a factor for scaling down utility (the factor would function like the tax rate).

Overall, the local property tax finance solution is surprisingly similar to the social planner's solution. The similarity is striking because the social planner was given more information *for free* than he could possibly get even if he were to engage in costly information gathering. In other words, what seems like a ridiculously optimistic outcome under centralized finance is similar to what is achievable under conventional local property tax finance in a residential market with many jurisdictions. Local property tax finance is able to manage productivity as well as it does because it makes some information verifiable and partials out other information that the social planner needs.

The Tiebout residential market makes quality verifiable because it produces a measure of quality for cost--in the form of property prices. Households take account of the outcomes they observe, whether objective or subjective, when considering how much they value a property relative to alternatives in other jurisdictions. Local property tax-based finance exploits the information contained in property prices by making each agent's budget a function of the change in property prices that he can induce by offering higher quality for cost. In contrast, centralized finance rules can only reward quality by using costly information on objective outcomes (thereby neglecting subjective outcomes) and trying to partial out household type.

A nice feature of local property tax-based finance is that it removes the main disadvantage of incumbency. If incumbent property owners commit themselves to rewarding their agent wholly on the basis of *moving* households' valuations, they spread the benefits conveyed by moving households, effectively making costless mobility universal, and removing agents' incentives to exploit them.

Also, under local property tax finance, households use their private information to identify other households of their own type. Their recognition of same-type households allows household type to be partialled out of the productivity problem. No one needs to be able to classify all households or estimate the direct effect of household type on outcomes. Parents only need to make simple comparisons using the people they are most likely to know. They only need observe that parents like them have better outcomes, on average, in school district X than in school district Y. They need not quantify the degree of dissimilarity between themselves and parents of other types; they need not estimate the direct effect of type on outcomes in order to partial type out. In contrast, social planners with centralized finance *do* need to quantify dissimilarity among types and *do* need to estimate the effect of type on outcomes.

XIII. Caveats

Because local property tax-based finance manages productivity by exploiting individuals' comparisons among jurisdictions, it is not good at controlling productivity problems that are universal, or common to all jurisdictions. For instance, if regulations that lower productivity are imposed on all jurisdictions by state or federal governments, local property tax finance will not limit the productivity problem since no household will be able to move away from the regulation.

Since local property tax-based finance manages productivity by exploiting individual behavior, it is not good at raising productivity when peer interactions create sub-optimal productivity that could only be remedied by elaborate coordination by a social planner. Such coordination would require that the social planner not only learn all the information he has in section VI, but also find out how the peer effects work and arrange for side payments that convince households to voluntarily rearrange themselves. Although peer effects receive considerable attention in the school finance literature, the informational demands associated with taking advantage of them are so extraordinary that it is not surprising that conventional local public finance does not

handle them well.

The discussion so far might make one conclude that ongoing politics (after the initial choice of the tax rate) are largely unnecessary. In fact, ongoing politics are necessary because they are the means by which new verifiable information is incorporated. For instance, if a state were to fund the teaching of all special education students from general revenues, districts would want to reduce local tax rates appropriately. Politics should be thought of as the mechanism of switching between different "incentive contracts" operated by local property tax finance.

IX. Conclusions

This paper attempts to provide a realistic model of the productivity of local public goods providers. The key insight is that, when the Tiebout process functions with local property tax finance, it can generate verifiable demand information that can be used to manage the productivity of local public goods providers. Local property tax finance also takes advantage of households' ability to recognize other households similar to themselves. Because cost conditions are not observable, local property tax finance does not *eliminate* agents' rent. But, it does reduce rent significantly--attaining a solution similar to the second-best solution that could be achieved by a social planner who has an amount of information so generous as to be implausible. In addition, local property tax finance is largely automatic, cheap to administer, and removes much of the disadvantage of moving costs (which would otherwise allow agents to exploit incumbents). The most surprising result, overall, is how much conventional local property tax finance embodies (probably fortuitously) key features of mechanisms designed to manage the productivity of regulated goods providers.

Centralized finance of local public goods does not have a similar means of automatically employing information from individual households' actions. It can base rewards only on costs, student outcomes, and household characteristics that are verifiable. Even when information is verifiable, it is costly to gather and can

be politically difficult to use. Centralized finance is likely to allow producers to extract rents, either by minimizing unverifiable quality or by exaggerating unverifiable cost conditions.



The model is not meant to be an ending point for discussion, but rather a starting point for thinking about the productivity implications of flaws in the Tiebout process (too few jurisdictions, inter-jurisdictional spillovers, and so on) and of reforms such as vouchers and privatization. It is most useful for its exposure of the sources of productivity problems, its explanation of why the conventional system is able to reduce but not eliminate rents, and its description of the information required for maximum productivity.



I am grateful to two anonymous referees, Dennis Epple, Martin Feldstein, Claudia Goldin, Roger Gordon, James Hines, Lawrence Katz, Derek Neal, Thomas Nechyba, James Poterba, Richard Romano, David Sappington, Jean Tirole, Rick Van der Ploeg, and seminar participants at Harvard University, the National Bureau of Economic Research, the Trans-Atlantic Public Economics Seminar, and the University of Florida for conversations and comments about this paper. All errors are my own.

References

- Benabou, R., 1993. Workings of a City: Location, Education, and Production. *The Quarterly Journal of Economics*, 108(3), 619-52.
- Black, S., 1998. Taking Teachers' Measure. *American School Board Journal*, 185(2), 39-42.
- Brennan, G. and J. Buchanan, 1978. Tax Instruments as Constraints on the Disposition of Public Revenues. *Journal of Public Economics*, 9(3), 301-18.
- , 1977. Towards a Tax Constitution for Leviation. *Journal of Public Economics*, 8(13), 255-73.
- Eberts, R., E. Schwartz, and J. Stone, 1990. School Reform, School Size, and Student Achievement. *Economic Review Federal Reserve Bank of Cleveland*, 26(2), 2-26.
- Epplé, D., R. Filimon, and T. Romer, 1984. Equilibrium Among Local Jurisdictions: Toward an Integrated Treatment of Voting and Residential Choice. *Journal of Public Economics*, 24(3), 281-308.
- , 1983. Existence of Voting and Housing Equilibrium in a System of Communities with Property Taxes. *Regional Science and Urban Economics*, 23(5), 585-610.
- Epplé, D. and R. Romano, 1998. Competition between Private and Public Schools, Vouchers, and Peer Group Effects. *American Economic Review*, 88(1), 33-62.
- Epplé, D. and G. Platt, 1998. Equilibrium Among Jurisdictions When Households Differ by Preferences and Income. *Journal of Urban Economics*, 43(1), 23-51.
- Epplé, D. and A. Zelenitz, 1981. The Implications of Competition Among Jurisdictions: Does Tiebout Need Politics?" *Journal of Political Economy*, 89(6), 1197-1217.
- Glaeser, E., 1996. The Incentive Effects of Property Taxes on Local Governments. *Public Choice*, 89(1-2), 93-111.
- Hamilton, B., 1976. Capitalization of Intrajurisdictional Differences in Local Tax Prices. *American Economic Review*, 66(5), 743-53.
- Fernandez, R. and R. Rogerson, 1996. Income Distribution, Communities, and the Quality of Public Education. *Quarterly Journal of Economics*, 111(1), 135-64.
- Hoxby, C., forthcoming 1999. Does Competition Among Public Schools Benefit Students and Taxpayers? *American Economic Review*.
- , 1994 "Do Private Schools Provide Competition for Public Schools?" NBER Working Paper No. 4978.
- Inman, R., 1987. The 'New' Political Economy. Chapter 12 in A. Auerbach and M. Feldstein, eds. *Handbook of Public Economics*. North-Holland, Amsterdam.

- Laffont, J. and J. Tirole, 1993. *A Theory of Incentives in Procurement and Regulation*. The MIT Press, Cambridge, Massachusetts.
- Lopez-de-Silanes, F., Shleifer, A., and R. Vishny, 1997. Privatization in the United States. *Rand Journal of Economics*, 28(3), 447-71.
- Nechyba, T., 1997. Existence of Equilibrium and Stratification in Local and Hierarchical Tiebout Economics with Property Taxes and Voting. *Economic Theory*, 10(2), 277-304.
- Niskanen, W., 1971. *Bureaucracy and Representative Government*. Aldine-Atherton Press, Chicago.
- Peltzman, S., 1996. The Political Economy of Public Education: Non-College Bound Students. *Journal of Law and Economics*, 39(1), 73-120.
- , 1993. The Political Economy of the Decline of American Public Education. *Journal of Law and Economics*, 36(1), 331-70.
- Rauscher, M., 1998. Leviathan and Competition Among Jurisdictions: The Case of Benefit Taxation. *Journal of Urban Economics*, 44(1), 59-67.
- Romer, T. and H. Rosenthal, 1978. Political Resource Allocation, Controlled Agendas, and the Status Quo. *Public Choice*, 33(4), 27-43.
- Rubinfeld, D., 1987. The Economics of the Local Public Sector. Chapter 11 in A. Auerbach and M. Feldstein, eds. *Handbook of Public Economics*. North-Holland, Amsterdam.
- Stiglitz J., 1977. "The Theory of Local Public Goods," in M. Feldstein and R. Inman, eds., *The Economics of Public Services*. London: MacMillan, 274-333.
- Tiebout, C., 1956. A Pure Theory of Local Public Expenditures. *Journal of Political Economy*, 64, 416-24.

**Figure 1****Figure 2**






Figure 5




Figure 6