Business Review

Federal Reserve Bank of Philadelphia

September•October 1993 ISSN 0007-7011

Do Americans Save Too Little?

B. Douglas Bernheim & John Karl Scholz

Highways and Education: The Road to Productivity?

Gerald A. Carlino

Business Review

The BUSINESS REVIEW is published by the Department of Research six times a year. It is edited by Sarah Burke. Artwork is designed and produced by Dianne Hallowell under the direction of Ronald B. Williams. The views expressed here are not necessarily those of this Reserve Bank or of the Federal Reserve System.

SUBSCRIPTIONS. Single-copy subscriptions for individuals are available without charge. Institutional subscribers may order up to 5 copies.

BACK ISSUES. Back issues are available free of charge, but quantities are limited: educators may order up to 50 copies by submitting requests on institutional letterhead; other orders are limited to 1 copy per request. Microform copies are available for purchase from University Microfilms, 300 N. Zeeb Road, Ann Arbor, MI 48106.

Please send subscription orders, back orders, changes of address, and requests to reprint to Publications, Federal Reserve Bank of Philadelphia, Department of Research and Statistics, Ten Independence Mall, Philadelphia, PA 19106-1574, or telephone (215) 574-6428. Please direct editorial communications to the same address, or telephone (215) 574-3805.

SEPTEMBER/OCTOBER 1993

Do Americans Save Too Little?

B. Douglas Bernheim & John Karl Scholz**

Since the mid-1980s, low rates of national saving in the United States have generated an enormous amount of concern among both economists and policymakers. Proposals to address these concerns fall into two broad categories: policies designed to increase public saving and policies intended to promote private saving. The former is synonymous with deficit reduction, while the latter includes tax incentives, pension policy, and strategies for discouraging the use of private debt. Some economists argue that deficit reduction is the most reliable and efficacious method of increasing national saving (Summers, 1985), while others maintain that restoring adequate rates of private saving is essential (Bernheim, 1991). To evaluate the merits of strategies that target private saving, we must resolve two issues. First, aside from the obvious fact that private saving is one component of national saving, is there reason to be concerned about the rate of private saving? Second, are there any effective and reliable methods of promoting private saving?

*B. Douglas Bernheim is the John L. Weinberg Professor of Economics and Business Policy, Princeton University. John Karl Scholz is an assistant professor of economics, University of Wisconsin, Madison. When this article was written, Bernheim was a visiting scholar in the Research Department of the Philadelphia Fed. The authors gratefully acknowledge the work of Robert Avery and Arthur Kennickell, who developed a clean copy of the 1983-86 *Survey of Consumer Finances* and provided extensive documentation.

THE ADEQUACY OF HOUSEHOLD SAVING

According to common wisdom, Americans consume too much and save too little. This impression is largely traceable to widely publicized statistics on aggregate personal saving. International comparisons reveal that U.S. households save significantly less than their foreign counterparts. Between 1980 and 1991, Americans saved 6.4 percent of disposable personal income, compared with 9.8 percent for OECD Europe and 15.7 percent for Japan (Organization for Economic Cooperation and Development, 1992). And since the mid-1980s, the rate of household saving in the U.S. has been well below its historical average (Figure 1).

Although these statistics raise legitimate concerns, they do not provide definitive evidence of a problem. As measured, personal saving excludes capital gains. Thus, in principle, households can accumulate wealth at a rapid rate even when their measured rates of saving are low. Rates of personal saving can also vary across both time and countries for reasons unrelated to the adequacy of saving considered from the perspective of individual households.[1] To understand this second point, consider the following hypothetical example. Envision two countries, A and B, that are identical in all respects except that the elderly make up a larger fraction of the population in A than in B. Since households tend to accumulate wealth prior to retirement and spend wealth thereafter, we would expect to observe a higher rate of aggregate personal saving in country B. Indeed, in an economy with no growth in either population or productivity, dissaving by retirees could completely offset saving by workers: in principle, regardless of how well individual households prepared for retirement, we might observe virtually no aggregate personal saving. Thus, ultimately, we can judge the adequacy of personal saving only by examining microeconomic data on the behavior of individual households.

Generally, the available evidence suggests that American workers have prepared poorly for retirement. Diamond (1977) found that, during the 1960s, 40 percent of couples and more than 50 percent of unmarried individuals reported that after retirement they received no money income from assets. At age 60, nearly 30

FIGURE 1
Rate of Personal Saving, National Income Accounts

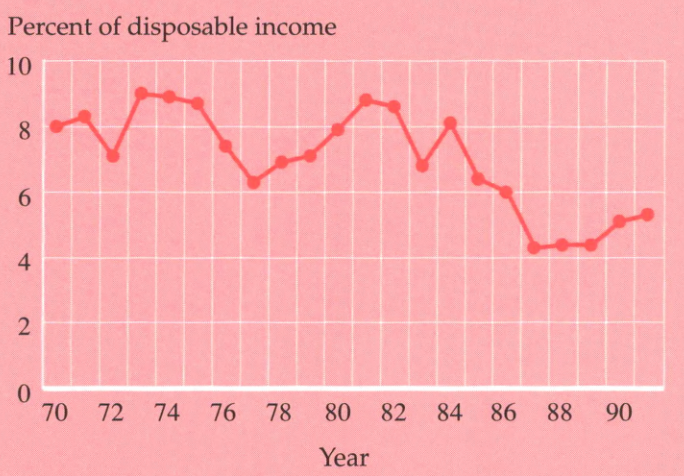

[1]Indeed, Meyer, 1992, argues that demographic differences account for roughly one-third of the gap in personal saving relative to GNP between Germany and the U.S. during the 1980s and roughly two-thirds of the gap between Japan and the U.S.

percent of middle-class individuals lacked sufficient wealth to replace two years' worth of income. Similarly, Hamermesh (1984) concluded that, during the 1970s, most elderly individuals had not accumulated sufficient resources to sustain their accustomed standards of living. Indeed, consumption shortly after retirement exceeded the highest sustainable level of consumption by an average of 14 percent. Hamermesh also found that within a few years of retirement most retirees were forced to reduce their expenditures substantially.[2]

Asset Accumulation Profiles. More recent evidence on the adequacy of saving appears in Bernheim and Scholz (1992a). Using an elaborate model of household decision-making, we simulated asset accumulation profiles (trajectories) that households should follow (given the assumptions of the model) to prepare adequately for retirement.[3] We then compared these simulated profiles with ones estimated from recent surveys of households' actual saving behavior. (For a more detailed description of the model, see *Explanation of the Model*.)

The simulation model describes only the accumulation of assets for retirement. There are, of course, many reasons to save. Households should take precautions against the possibility of illness, layoff, disability, death, and other risks for which they are imperfectly insured. In addition, most households accumulate resources to pay for large expenses such as college tuition or the purchase of an automobile. For some individuals, saving is motivated in part by the desire to leave a substantial bequest upon death. Unfortunately, when examining the data, we cannot determine whether particular assets were accumulated for retirement or for some other purpose. Consequently, the comparison between estimated trajectories and simulated trajectories may provide an overoptimistic picture of the adequacy of household saving.

We show graphic depictions (Figures 2 and 3) of a simulation for a household with the following characteristics: age 27 (as of 1991), two years of college education, married, two workers with total current earnings of $60,540, and the primary earner covered by a private pension plan. This household's optimal trajectory of consumption and after-tax earned income (including pensions and Social Security) is shown in constant 1991 dollars (Figure 2, page 7).[4] Note that after-tax earnings rise steeply early in life. Earnings growth continues at a reduced level until the individual reaches age 55, at which point it begins to fall. After retirement, earned income consists of Social Security and private pension benefits. Since pensions are not perfectly indexed for inflation, real benefits decline gradually over time.

As a direct consequence of the household's rapid earnings growth early in life, it saves nothing for retirement prior to age 30. Between ages 30 and 80, the consumption trajectory is relatively flat. This flat trajectory reflects the household's preference for a stable standard of living. However, during the 30s and 40s, consumption is elevated relative to the 60s and 70s. This pattern results from changes in household composition: between the ages of 30 and 50, the typical household incurs significant child-rearing costs. Consumption declines rapidly after age 80 until, at age 101, it matches after-tax retirement benefits. Falling survival probabilities cause this end-of-life decline. Since there is a relatively low probability of reaching age 90,

[2]Other economists have reached somewhat more optimistic conclusions. See Kotlikoff, Spivak, and Summers, 1982.

[3]Development of this model was sponsored by Merrill Lynch & Co., Inc., and is described in Bernheim, 1992b.

[4]We use the word "trajectory" to describe the manner in which an economic variable, such as consumption, income, or wealth, evolves as the household ages.

Explanation of the Model

Our simulation model reflects a "life-cycle" approach to the average household's financial decision-making process. It takes into account the fact that predictable changes in household earnings resulting from age and stage of career may not match up very well with consumption needs. For example, the financial needs of most households are usually highest during the child-rearing years, while household earnings usually reach their highest point after children have left home. The household varies its rate of saving in order to achieve a better match between the ability to spend and the need to spend. It saves least in years when spending needs are high and more in years when spending needs decline.[a] The model forecasts households' future income and derives the optimal consumption (and thus saving) trajectories consistent with those income forecasts.

Our life-cycle calculations account for a variety of current and future household characteristics, including age, income level, pension coverage, education, marital status, gender (if unmarried), and household composition (the numbers of children and dependent adults).[b] The model also projects and adjusts for future macroeconomic conditions that ought to affect savings behavior, including interest rates, inflation rates, and baseline wage growth. In addition, the model provides a realistic treatment of income taxes, payroll taxes, and social security benefits.

To conduct simulations, one must also choose values for several "preference parameters." For example, the model includes a parameter commonly known as the "pure rate of time preference," which expresses the value that a household places on future consumption relative to current consumption.[c] The value of this particular parameter has a profound effect on the simulation results. When the pure rate of time preference is sufficiently low, it is optimal for the household to save nothing. For this reason, the absence of saving is not necessarily the result of irrationality. Rather, it may simply reflect impatience.

We have calibrated our model (that is, chosen values for the preference parameters) so that the simulations produce a standard of living during retirement that is roughly comparable to the standard of living enjoyed prior to retirement.[d] Consequently, it is appropriate to interpret our results as follows: if households fall significantly short of simulated asset accumulation targets, they will ordinarily be forced to accept serious reductions in their standards of living after retirement.

---

[a]When spending needs are sufficiently high relative to income, a household may wish to liquidate or borrow against accumulated assets. Once assets are exhausted, it may be optimal for the household to borrow against future income. However, for most households, it is extremely difficult to obtain sizable unsecured loans. Our model therefore imposes a "liquidity constraint," which ensures that the household's net wealth remains positive.

[b]Our calculations reflect the fact that larger households benefit from significant economies of scale. Research on household scale economies indicates that two adults in a household can obtain the same standard of living as one adult living alone with added expenditures of slightly more than 40 percent. Research also shows that the financial impact of adding one adult to a household is roughly equivalent to adding 2.5 children. See Cutler and Katz, 1992.

[c]Other important preference parameters include a minimum subsistence level for consumption and a parameter known as the "intertemporal elasticity of substitution," which measures the extent to which the household's willingness to trade off current consumption for future consumption is affected by the level of current consumption relative to future consumption.

[d]Specifically, we use a pure rate of time preference equal to the product of 0.99 and one-year gender-specific survival probabilities (taken from standard life tables). The minimum consumption level is set equal to $10,000 (measured in 1991 dollars), and is adjusted for family size. A value of 0.25 is used for the intertemporal elasticity of substitution.

the household would prefer to accept a lower standard of living at age 90 and later (if it survived that long) in favor of a higher standard of living earlier in life.

The associated optimal trajectory of retirement assets is also depicted (Figure 3). Assets accumulate at an increasing rate from age 30 to retirement, peak at retirement, then decline steadily until they are exhausted at age 100.

We then estimated actual asset trajectories using data from the *Survey of Consumer Finances* (SCF) for 1983 and 1986.[5] The Board of Governors of the Federal Reserve (in conjunction with other federal agencies) sponsored the SCF, recognized as one of the best available sources of data on household balance sheets.[6]

[5]Our measure of accumulated net worth includes stocks and mutual funds, bonds, checking and savings accounts, IRA and Keogh accounts, money market accounts, certificates of deposit, profit-sharing and thrift accounts, the dollar cash value of whole life insurance, and other financial assets, as well as equity in property (other than primary residences) and business assets, less credit card, consumer, and other debt. This measure excludes all assets and liabilities associated with homes and vehicles, since households appear to have a strong aversion to paying living expenses during retirement by drawing down the equity in their homes (see Venti and Wise, 1989). Also, it seems likely that few individuals save for retirement by accumulating wealth in the form of vehicles. Accumulated wealth for 1983 is expressed in 1986 dollars using the Consumer Price Index.

[6]See Avery and Elliehausen, 1988, and Avery and Kennickell, 1988, for a more complete discussion of the SCF.

FIGURE 2
Simulated After-Tax Income and Consumption Trajectories

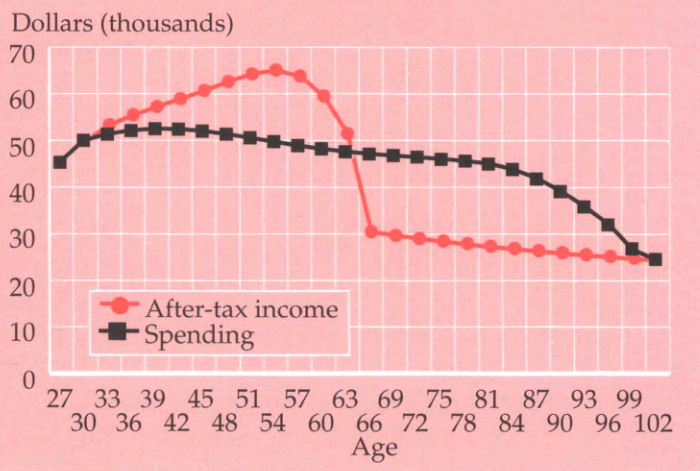

FIGURE 3
Simulated Wealth Trajectory

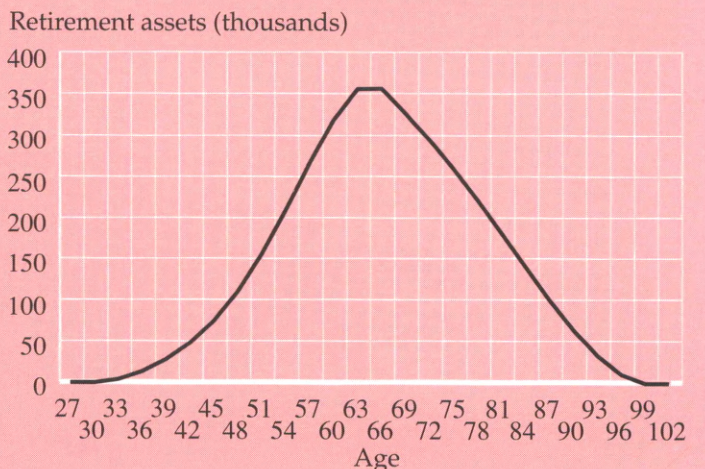

Our analysis allows us to compare actual and simulated optimal behavior. The results for households in which the primary worker has not completed college are shown in Figure 4. In this figure, "actual" refers to the estimated change in wealth (measured as a fraction of wage income) for the representative household within each age group (calculated using the SCF); "Sim/no pen" indicates the simulated change in wealth (again as a fraction of wage income) for a representative household without pension coverage for the primary earner; and "Sim/pen" denotes the simulated change in wealth for a representative household with pension coverage for the primary earner. Note that the simulated change in wealth rises steeply with age. This steep increase in assets results from two factors. First, during most of an individual's working life wages rise more rapidly than consumption (see Figure 2). Second, reinvested capital income rises as the household accumulates assets. In contrast, the estimated change in wealth does not vary significantly with age. By the time the household reaches middle age, simulated asset accumulation exceeds actual accumulation by a wide margin.[7] Overall, between 1983 and 1986, households without a college education saved far less than the simulation model predicts (Figure 4).

Results for households in which the primary earner completed college are depicted in Figure 5. The contrast between Figures 4 and 5 is remarkable. In cases where the household head completed college, both simulated and estimated changes in wealth rise steeply with age. Moreover, simulated asset accumulation tracks actual asset accumulation remarkably well. Taken at face value, Figure 5 suggests that highly educated households saved adequately for retirement between 1983 and 1986.

Although it is tempting to conclude that inadequate saving is largely confined to those without a college education, this conclusion must be tempered by two considerations. First, as is apparent from Figure 1, personal saving declined sharply after the 1983-86 period on which the estimates are based. Using a sample of relatively young individuals (ages 25 through 44) surveyed in early 1992, Bernheim (1992a) found much more pervasive evidence of inadequate saving. Second, the model probably understates the amount of wealth that each household ought to accumulate. The most obvious reason for this discrepancy is that the

FIGURE 4
Rates of Asset Accumulation
No College Degree

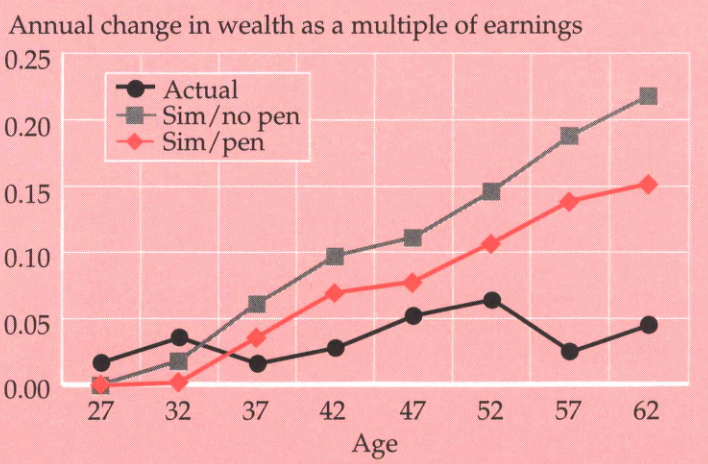

[7] Although estimated asset accumulation is actually higher at ages 27 and 32, this is of little consequence; recall that the data reflect saving for a variety of purposes aside from retirement.

simulations envision retirement planning as the sole motive for saving.[8]

To the extent that many households prepare poorly for retirement, there is cause to be concerned about the rate of personal saving, per se. Historically, pension policy and tax policy have been the two most important tools for stimulating personal saving. We will discuss evidence on the efficacy of each of these strategies in turn.

PENSION POLICY

In recent years, the accumulation of assets in private pension plans has accounted for a substantial fraction of personal saving (Bernheim and Shoven, 1988). This observation raises the possibility that policies affecting private pensions may have powerful effects on aggregate personal saving. Whether these effects would actually materialize depends on the way workers would respond to an expansion of private pension coverage. Economic theory suggests that such an expansion would simply crowd out other forms of personal saving: once workers realize that their employers are, in effect, saving for them, workers will save less themselves. The simulation results presented in the previous section illustrate this principle. However, previous studies of personal saving have generally failed to find evidence to support the notion that private pensions significantly reduce other forms of personal saving.[9] Depending on whether we credit the theoretical analysis or the empirical studies, we can reach dramatically different conclusions about the effect of pension policy on aggregate personal saving.

The analysis described in the preceding section raises an intriguing possibility: if the behavior of those with a college education (and higher average incomes) conforms to the predictions of standard economic theories, while the behavior of those without a college education (who have lower average incomes) does not, perhaps private pensions do displace personal saving among the college educated, but not among the rest of the population. In that case, pension policy could be an effective tool for stimulating total personal saving, so long as it is primarily used to pro-

[8]In addition, it is quite likely that the model overstates mortality probabilities (since it does not make any allowance for the fact that these probabilities are projected to decline in the future), understates the importance of health and long-term care costs for the elderly, and fails to consider the effects of mounting economic pressures that may force Congress and employers to scale back existing retirement benefits.

[9]See, for example, the review in Shefrin and Thaler, 1988, particularly pages 622-24.

FIGURE 5
Rates of Asset Accumulation
College Degree

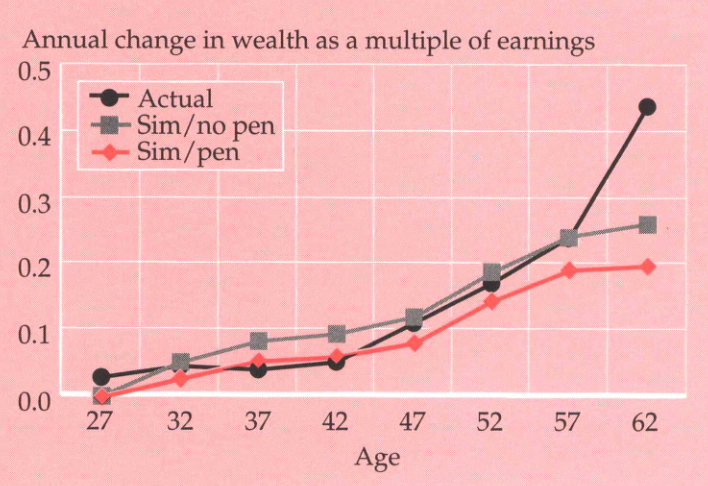

vide incentives for expanded coverage among lower income, generally less educated, workers.

To investigate this idea, we estimated equations that explained the median value of household wealth as a function of age, total household earnings, private pension coverage, and educational attainment. We then used these equations to project asset accumulation profiles.

Results for the median household in which the primary earner has not completed college are presented in Figure 6. Note that pension eligibility has little or no effect on the actual path of household wealth accumulation. From a statistical perspective, the estimated equation supports the notion that, at every age, less educated households with private pensions accumulate wealth at the same rate as those without private pensions.

Results for households in which the primary earner has completed college are displayed in Figure 7. Statistically, the data decisively reject the premise that the rate of asset accumulation is unrelated to pension eligibility. Note that those eligible for pensions accumulate resources at a significantly slower rate than those without pensions. Remarkably, at age 62, the gap between the assets of these two groups is almost identical in magnitude to the predicted gap that emerges from our computations. These patterns are strongly consistent with the view that private pensions displace other personal saving for college-educated households.[10] These results suggest that other studies may have failed to find a significant saving displacement effect simply because they did not distinguish between households on the basis of education (or permanent income).

The contrast between Figures 6 and 7 points to a clear and important conclusion for pension policy: private pensions displace personal wealth accumulation only when the head of the household is college-educated. This observation aligns with the evidence on the adequacy of personal saving described in the first section of this article. Indeed, our evidence broadly supports a more general conclusion: college-educated households behave in the manner predicted by standard economic theories of saving, while less well-educated households do not. Past and current policies have been more successful at stimulating the expansion of pension coverage among college-educated

[10]It is unlikely that the observed relationship between pension coverage and saving results from spurious factors, since such factors would presumably also have produced the same patterns for less educated households.

FIGURE 6
Estimated Wealth Trajectories
No College Degree

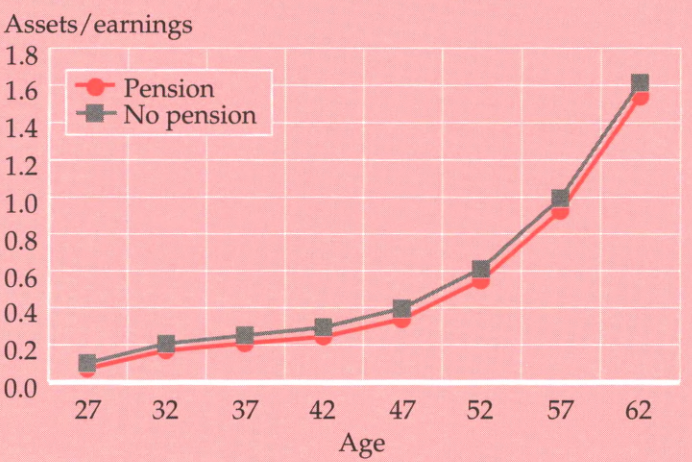

workers than among those with less education. Analysis of the SCF data reveals that 75.2 percent of college-educated husbands are covered by private pensions. In contrast, only 55.7 percent of husbands who lack a college education are covered by private pensions. In other words, the current system is quite effective at providing pensions to those individuals who reduce other saving in response and much less effective at providing coverage to those individuals for whom pensions would represent incremental saving.

TAX POLICY

The most commonly discussed strategies for stimulating personal saving entail reductions in the taxation of capital income. Economic theory suggests that households will respond to a higher after-tax rate of return on savings by increasing future consumption relative to current consumption. However, theory does not necessarily predict that current saving will rise. (The reason is that a higher rate of return will make wealth grow more rapidly, enabling greater future consumption, even if the household were to save a bit less out of its current income.) Indeed, empirical estimates of the sensitivity of saving to the after-tax rate of return (called the interest elasticity of saving) vary widely (Boskin, 1978; Summers, 1981; and Hall, 1988).

Individual Retirement Accounts. Most current proposals to provide tax incentives for saving are patterned after individual retirement accounts (IRAs). IRAs were established as part of the 1974 Employee Retirement Income Security Act to give workers not covered by employer-provided pension plans added incentives to accumulate resources for retirement. In 1981, IRA eligibility was extended to all taxpayers. Subsequently, the Tax Reform Act of 1986 curtailed the tax-deductibility of IRA contributions for high income households. The existence of an income cap for IRAs raises an important question: does the sensitivity of saving to the after-tax rate of return vary systematically across income classes? The answer to this question makes it possible to determine whether the current system targets the most responsive groups.

Simulations based on the model described in this article suggest that higher income individuals will be much more responsive than lower income individuals to changes in the after-tax rate of return. Averaging across individuals with pensions and individuals without pensions, the simulations imply that saving by 35-year-old, college-educated households would increase by 10.2 percent in response to a permanent one-percentage-point in-

FIGURE 7
Estimated Wealth Trajectories
College Degree

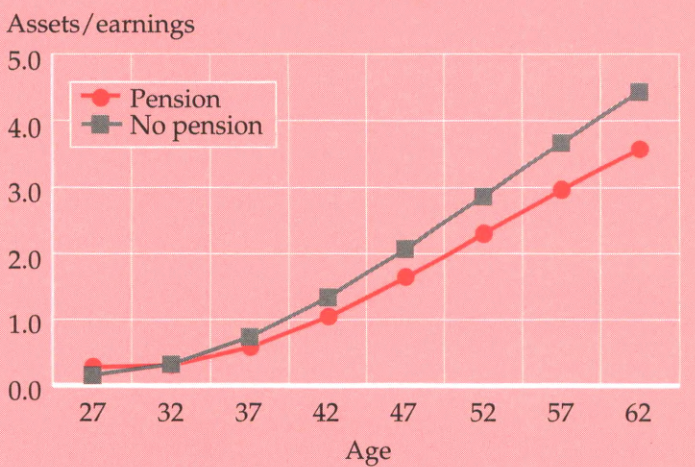

crease in the before-tax rate of return, while the saving of 35-year-old, high-school-educated households would *fall* by 4.5 percent. Consequently, policies that provide tax incentives for saving exclusively to lower income households exclude those individuals most likely to increase saving in response to tax incentives; indeed, such policies could actually reduce aggregate personal saving.

This positive relationship between income and the interest elasticity of saving results from a natural economic consideration, rather than from some peculiar feature of the simulation model. It is natural to assume that when planning for the future, most households are concerned first and foremost with saving enough to assure themselves of some minimum standard of living. As lifetime resources increase, households have more discretion to allocate resources in a manner that increases consumption above and beyond this minimum standard both today and in the future.

For low income households, saving to achieve some minimum future consumption is probably far more important than saving to fund incremental consumption. Saving to provide for minimum consumption is, in effect, saving for a fixed target. An individual who saves to achieve some target will reduce saving in response to an increase in the rate of return (Bernheim and Shoven, 1988). Thus, because target saving dominates the simulated behavior of these households, they exhibit a low or negative interest elasticity of saving. For high income households, however, saving to fund incremental consumption is probably far more important than saving to achieve the minimum consumption target. Incremental saving dominates the simulated behavior of these households. Thus we observe a high interest elasticity of saving among higher income, well-educated households. Discretionary saving to finance consumption over and above the target responds positively to an increase in the rate of return.

Of course, in the preceding sections, we observed that the behavior of less educated (generally lower income) households may not conform to standard economic theories. Although this finding reduces our faith in the applicability of our simulation results, it does not reverse our conclusions concerning the interest elasticity of saving. The notion that households will respond to a change in the after-tax rate of return is predicated on the assumption that households rationally anticipate and plan for future economic contingencies. To the extent that this assumption proves incorrect, there is no particular reason to believe that lower income households will respond to a change in the after-tax rate of return in the first place.

Tax Policy Initiatives. Two prominent current policy initiatives would reverse the direction of the 1986 reforms and improve tax incentives for saving to households in higher income brackets. Family saving accounts (FSAs), proposed by the Bush administration, would allow single individuals with adjusted gross incomes (AGI) below $60,000 and married couples with AGI below $120,000 to make contributions of up to $2500 to qualified accounts. The FSA proposal is an example of a "back-loaded" system: contributions are nondeductible, but accumulated funds are not taxed upon withdrawal. An alternative proposal, the Bentsen-Roth "super-IRA," would allow individuals to contribute up to $2000 to either a traditional or a back-loaded IRA.[11]

[11]On August 3, 1992, the Senate Finance Committee approved H.R. 11, the Revenue Bill of 1992. Like the Bentsen-Roth super-IRA, this bill would restore the deductibility of IRA contributions for all taxpayers and establish new back-loaded IRAs. Contributions to back-loaded IRAs could be withdrawn without penalty after five years. The bill would also allow taxpayers to make penalty-free early withdrawals from IRAs for the purchase of a first house, for higher education expenses, for medical expenses, and for long spells of unemployment.

Unfortunately, there are sound conceptual reasons to doubt the effectiveness of extending eligibility for IRA-style accounts to higher income households. First, contributions are capped. Under the current system, a single taxpayer, for example, can make no more than $2000 in tax-deductible contributions. For an individual taxpayer who would have saved more than $2000 in the absence of IRAs, the availability of an IRA does not affect the costs or benefits that might result from an additional dollar of saving and, therefore, provides no incentive on the margin for the taxpayer to increase saving. In such cases, the IRA constitutes a "giveaway" of public funds (it reduces federal tax receipts but does not promote more saving). In addition, the IRA may actually induce the taxpayer to increase consumption, since it increases his or her total after-tax resources. For both of these reasons, the IRA would contribute to a lower rate of national saving. These concerns are of little significance for low income households, since few of them would save more than $2000 in the absence of the program. It is far more likely that high income households would save more than the contribution limit. Thus, IRA-style proposals may be a particularly ineffective vehicle for providing tax incentives for saving to high income households.

A second reason for doubting the effectiveness of IRA-style accounts for high-income households is that even if such a taxpayer would not (in the absence of IRAs) have saved more than the IRA contribution limit in a given year, he or she could take full advantage of the IRA deduction either by financing contributions with previously accumulated assets or by borrowing. Indeed, the 1991 *Tax Guide for College Teachers* devotes a full page to the issue "What If You're Short of Cash to Fund Your IRA?" (pp. 229-30). The Guide describes an IRS private letter ruling that allows households to finance their IRAs by borrowing. Contributions funded either by shifting existing assets or by borrowing do not increase household saving. Instead, by reducing federal tax receipts, they add to the federal budget deficit and depress national saving. Once again, it is more likely that high income households (who possess greater wealth, financial sophistication, and access to credit markets) would engage in borrowing or asset shifting and thus defeat the purpose of the program.

Empirical evidence on the efficacy of IRAs is mixed. Gale and Scholz (1992) find little evidence that IRAs stimulated household saving between 1983 and 1986. Venti and Wise (1986, 1987, 1990, 1991) and Feenberg and Skinner (1989) suggest that most IRA contributions during this period represent net increases in household saving. Joines and Manegold (1991) conclude that the effects of IRAs on household saving are unlikely to be as large as the estimates of Venti and Wise and may be as small as the estimates of Gale and Scholz.

An alternative proposal to promote household saving, based on "premium saving accounts" (PSAs), is described in Bernheim and Scholz (1992b). A PSA system would require each taxpayer to save—in total—some fixed amount (the floor) before becoming eligible to make contributions to a tax-favored account. The taxpayer would be eligible to contribute each additional dollar of saving to the tax-favored account, up to some limit (the ceiling). These floors and ceilings would rise with AGI and certain types of capital income. As with IRAs, capital income accrued on balances held in PSAs would be exempt from taxation.[12]

The use of both floors and ceilings would

[12]With this essential structure, a PSA system could be either front-loaded or back-loaded. Penalties could be established to lock funds into tax-favored accounts for relatively short periods (e.g., seven years) or until some age close to retirement (perhaps age 59 1/2). Accounts could be established for specific purposes (e.g., retirement, purchase of a house, college education), or the accounts could be unrestricted.

create "windows" of program eligibility. Consider, for example, a married couple with an AGI of $80,000. They might face a floor of $8000 and a ceiling of $12,000. Should they save less than $8000 in the corresponding tax year, they would not be eligible to make any contributions to a tax-favored account. If, on the other hand, they saved $9500, they would be eligible for favorable tax treatment on $1500. If they saved more than $12,000, they would be eligible to make the maximum contribution of $4000 (the difference between $8000 and $12,000).

The most important distinctive feature of a PSA system is that floors and ceilings would vary with AGI. Eligibility windows could be positioned to maximize, within each income class, the number of households receiving tax breaks on the marginal dollar of saving. Doing so would maximize the incentive to save more. Higher-income taxpayers would not be deprived of tax incentives for saving; rather, they would simply be required to save much larger fractions of their incomes before becoming eligible for PSAs. It would also be much more difficult for households to take advantage of tax-favored PSA accounts by shifting assets or by borrowing because eligibility would be based on total saving. An individual cannot increase his total saving by shifting assets from one account to another or by borrowing to invest.[13]

To implement a PSA system, one needs to measure a household's total saving. Bernheim and Scholz (1992b) propose the following measure:[14]

Net purchases of assets (i.e., total purchases minus total sales) for assets on which investors receive capital gains and losses

plus

The January 1 to January 1 change in cash account balances (e.g., bank accounts),

minus

The January 1 to January 1 change in total debt (mortgages, consumer credit, etc.).

In effect, saving is defined as the incremental resources that an individual sets aside in any year over and above reinvested capital gains.[15,16]

Now we'll evaluate the effects of three distinct strategies for promoting household saving: an IRA-like program with an AGI cap (hereafter referred to as the "standard IRA" system), an IRA-like program without an AGI cap (hereafter referred to as the "universal IRA" system), and a PSA system. We compare the cost-effectiveness of extending tax incentives for saving to higher-income taxpayers through universal IRAs and PSAs.

Sample schedules that define eligibility windows for each level of AGI for a PSA system are given in Table 1. Separate schedules are given for married couples and single individuals. The schedules are chosen to maximize the ben-

[13]The administrative feasibility of monitoring total saving for each taxpayer is discussed in Bernheim and Scholz, 1992b.

[14]Many economists would define saving as the change in the stock of wealth between two points in time. If one adopts this definition, saving is very hard to measure: one would need to assess the market value of all assets every year. The definition used in the text represents a compromise between economic logic and administrative feasibility.

[15]Note that it is possible to compute this measure of saving without assessing the value of unrealized capital assets, since, by definition, unrealized gains are fully reinvested.

[16]If this definition of saving is employed, it is also important to adjust each taxpayer's eligibility floors and ceilings upward by the amount of capital income other than capital gains. See Bernheim and Scholz, 1992b, for a detailed discussion of this issue.

eficial effects of the program within each population subgroup.[17] To facilitate comparisons with IRAs, we have adopted window widths of $2000 per year for single households, $2250 per year for married couples with one earner, and $4000 per year for married couples with two earners. For example, a dual-earner married couple with an AGI of $30,000 and no capital income would have a floor of $0 and a ceiling of $4000 (Table 1). In contrast, a couple with an AGI of $120,000 and dividend and interest income of $2000 would have a floor of $16,362 (.167 × $86,000 + $2000) and a ceiling of $20,362.

The standard and universal IRA systems differ from the PSA proposal in that they anchor the eligibility window at $0 for all income classes and make no adjustment for capital income. The standard IRA system phases out deductible contributions for married couples with incomes between $40,000 and $50,000 and for single taxpayers with incomes between

[17]Note that the floor rises with income at different rates for married couples (16.7 cents for each dollar of income over $34,000) and single individuals (34 cents for each dollar of income over $42,000). Since actual patterns of saving differ by marital status, different schedules must be used to maximize the beneficial effects of the program.

TABLE 1[a]

Deductible Contribution Formula

Married Couples

If your income is	Deductible Qualified Contribution Floor (Added to Capital Income)	Deductible Qualified Contribution Ceiling (Added to Floor)
Less than $34,000	0	$2250 or $4000
Greater than $34,000	.167×(Income-34,000)	$2250 or $4000

Single Households

If your income is	Deductible Qualified Contribution Floor (Added to Capital Income)	Deductible Qualified Contribution Ceiling (Added to Floor)
Less than $42,000	0	$2000
Greater than $42,000	.34×(Income-42,000)	$2000

[a]For the purpose of comparison with IRAs, married couples with one earner are allowed to contribute $2250 and married couples with two earners can contribute $4000. In the actual implementation of this proposal we see no compelling reason to make this distinction.

$25,000 and $35,000.[18] The universal IRA system allows all households to make deductible contributions.[19]

We compare these plans on the basis of three criteria. The first criterion is a measure of effectiveness. Specifically, for each plan, we estimate the number of households that would receive a higher after-tax rate of return on the incremental dollar of saving. We refer to these households as the IMPACT GROUP. Our second criterion is a measure of wasteful subsidization. Specifically, for each plan, we estimate the number of households that would make the maximum eligible contribution to a tax-favored account while continuing to receive the unsubsidized after-tax rate of return on the incremental dollar of investment. We refer to these households as the NO-IMPACT GROUP. Our third criterion is also a measure of wasteful subsidization: we calculate the budgetary cost of subsidizing the NO-IMPACT GROUP. We refer to this cost as the GIVEAWAY.

Our calculations are once again based on data obtained from the SCF for 1983 and 1986. The interested reader is referred to Bernheim and Scholz (1992b) for details.

Compare the effects of the policies on married couples as shown in Table 2. The top panel shows the size of the IMPACT GROUP. Overall, the PSA system provides real incentives to 2.4 million couples, roughly 90 percent more than the IRA with AGI restrictions and 30 percent more than the universal IRA. The difference is particularly pronounced in the top income quintile. By definition, the IRA with AGI caps ignores these households. Relative to the universal IRA, the PSA increases the number of couples receiving marginal incentives in the top income quintile by nearly 125 percent. Since, in this sample, over 60 percent of positive household saving is attributable to households in the top quintile of the income distribution, this improvement is particularly important.

The bottom two panels of Table 2 measure the NO-IMPACT GROUP and the cost of these ineffective subsidies. The calculations show, for example, that the PSA system would reduce the number of households in the NO-IMPACT GROUP by 1.75 million (28.2 percent) and would reduce federal expenditures on ineffective subsidies by $2.0 billion (34.0 percent), relative to the universal IRA. In terms of cost-effectiveness, the PSA system increases the ratio of the IMPACT GROUP to the GIVEAWAY by 96.5 percent overall, and by 287.2 percent (that is, by a factor of almost four) in the top income quintile. The IRA with AGI caps also effectively reduces ineffective subsidies and budgetary cost, but it achieves this reduction by excluding the very households most likely to respond to tax incentives.

Note the results for single individuals (Table 3). Under a PSA system, the size of the IMPACT GROUP would increase significantly relative to other proposals. The size of the IMPACT GROUP in the highest income quintile would more than triple. Moreover, both the size of the NO-IMPACT GROUP and the GIVEAWAY would fall relative to the universal IRA. The result is a 49.7 percent increase in overall cost-effectiveness (the ratio of the IMPACT GROUP to GIVEAWAY), and a 551.3 percent increase in cost-effectiveness for the top income quintile, relative to the universal IRA proposal.

Other Initiatives. Pension policies and tax policies do not exhaust the full range of strategies for stimulating personal saving. One particular class of policies not discussed here mer-

[18]It should be noted that the current IRA system differs from the standard IRA system considered in the text in that it phases out deductible contributions only for households that are covered by private pension plans. The current system is, therefore, a blend of a standard system and a universal IRA system.

[19]The IRA-like proposals we simulate are superior to actual IRA schemes because, in practice, IRA schemes are susceptible to tax arbitrage strategies involving borrowing and asset shifting, which our simulations do not capture.

TABLE 2[a]

A Comparison of Three Saving-Incentive Proposals, Married Couples

Simulated Effect	IRA w/AGI Cap	Universal IRA	PSA
IMPACT GROUP (in 1000s)			
Highest Income Quintile	0	102	228
Full Population	1256	1840	2388
NO-IMPACT GROUP (in 1000s)			
Highest Income Quintile	0	1416	817
Full Population	3578	6218	4467
ANNUAL GIVEAWAY (in $ millions)			
Highest Income Quintile	0	1950	1119
Full Population	2006	5861	3870
COST-EFFECTIVENESS (ratio of IMPACT group to GIVEAWAY)			
Highest Income Quintile	—	.0523	.2038
Full Population	.3510	.3139	.6171

[a]Simulations use data from the 1983-86 *Survey of Consumer Finances*. Saving and column headings are defined in the text. The PSA schedule is given in Table 1.

its further attention. An accumulating body of evidence, including that contained in this article, suggests that the behavior of many households (particularly those with lower incomes) is not well described by traditional economic theories. To some, saving decisions appear to be governed by such factors as habit, mental accounting, and self-control. Consequently, it may be possible to design more effective policies by educating the population or by exploiting the psychology of saving. The Japanese appear to have had considerable success with such a strategy during the postwar period (Horioka, 1988, and Bernheim, 1991). The development of a framework for analyzing policies of this sort is an important research priority. Bernheim (1993) provides a preliminary analysis of these issues.

CONCLUSION

The evidence presented in this article supports the view that many Americans, particularly those without a college education, save too little. Our analysis indicates that it should

TABLE 3[a]

A Comparison of Three Saving-Incentive Proposals, Single Taxpayers

Simulated Effect	IRA w/AGI Cap	Universal IRA	PSA
IMPACT GROUP			
(in 1000s)			
Highest Income Quintile	0	40	134
Full Population	454	603	694
NO-IMPACT GROUP			
(in 1000s)			
Highest Income Quintile	0	350	197
Full Population	1078	1405	1155
ANNUAL GIVEAWAY			
(in $ millions)			
Highest Income Quintile	0	292	151
Full Population	460	845	650
COST-EFFECTIVENESS			
Highest Income Quintile	—	.1370	.8874
Full Population	.9870	.7136	1.0677

[a]Simulations use data from the 1983-86 *Survey of Consumer Finances*. Saving and column headings are defined in the text. The PSA schedule is given in Table 1.

be possible to increase total personal saving among lower income households by encouraging the formation and expansion of private pension coverage for such families. It is doubtful that favorable tax treatment of capital income would stimulate significant additional saving by this group. Conversely, the expansion of private pensions would probably have little effect on saving by higher income households. However, these households are more likely to increase saving significantly in response to favorable tax treatment of capital income. These findings imply that the design of the current system, which links eligibility for IRAs to an AGI cap, and which provides higher income households with more complete pension coverage, ensures a minimal impact on personal saving.

Extending tax incentives for saving to higher income households is problematic. We have discussed two competing options: the universal IRA and the premium saving account (PSA). Our analysis reveals that the PSA system is a more cost-effective vehicle for providing incentives to those households most likely to respond to tax incentives.

REFERENCES

Avery, Robert B., and Gregory E. Elliehausen. "1983 Survey of Consumer Finances: Technical Manual and Codebook," mimeo, Board of Governors of the Federal Reserve System, August 1988.

Avery, Robert B., and Arthur B. Kennickell. "1986 Survey of Consumer Finances: Technical Manual and Codebook," mimeo, Board of Governors of the Federal Reserve System, November 1988.

Bernheim, B. Douglas. *The Vanishing Nest Egg: Reflections on Saving in America.* Twentieth Century Fund, 1991.

Bernheim, B. Douglas. "Is the Baby Boom Generation Preparing Adequately for Retirement? Summary Report," mimeo, Princeton University, August 1992a.

Bernheim, B. Douglas. "Is the Baby Boom Generation Preparing Adequately for Retirement? Technical Report," mimeo, Princeton University, August 1992b.

Bernheim, B. Douglas. "What Determines Personal Saving? The Roles of Information and Economic Literacy," in *Tax Policy for Economic Growth in the 1990s.* Washington, D.C.: American Council for Capital Formation (forthcoming 1993).

Bernheim, B. Douglas, and John Karl Scholz. "Private Saving and Public Policy," mimeo, Princeton University and the University of Wisconsin—Madison, September 1992a.

Bernheim, B. Douglas, and John Karl Scholz. "Premium Saving Accounts: A Proposal to Improve Tax Incentives for Saving," mimeo, Princeton University and the University of Wisconsin—Madison, September 1992b.

Bernheim, B. Douglas, and John B. Shoven. "Pension Funding and Saving," in Zvi Bodie, John B. Shoven, and David A. Wise, eds., *Pensions in the U.S. Economy.* Chicago: University of Chicago Press and NBER, 1988, pp. 85-111.

Boskin, Michael. "Taxation, Saving, and the Rate of Interest," *Journal of Political Economy* (April 1978), pp. 3-27.

Cutler, David M., and Lawrence F. Katz. "Rising Inequality? Changes in the Distribution of Income and Consumption in the 1980s," mimeo, Harvard University, January 1992.

Diamond, Peter A. "A Framework for Social Security Analysis," *Journal of Public Economics* (December 1977), pp. 275-98.

Feenberg, Daniel R., and Jonathan Skinner. "Sources of IRA Saving," in Lawrence Summers, ed., *Tax Policy and the Economy*, Cambridge, MA: MIT Press, 1989, pp. 25-46.

Gale, William G., and John Karl Scholz. "IRAs and Household Saving," mimeo, UCLA and University of Wisconsin—Madison, 1992.

Hall, Robert. "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96, 1988, pp. 337-57.

REFERENCES (continued)

Hamermesh, Daniel S. "Consumption During Retirement: The Missing Link in the Life-Cycle," *Review of Economics and Statistics* (February 1984), pp. 1-7.

Hatsopoulos, George N., Paul R. Krugman, and James M. Poterba. *Overconsumption: The Challenge to U.S. Policy.* Washington, D.C.: American Business Conference, 1989.

Horioka, Charles Y. "Why Is Japan's Private Saving Rate So High?" in R. Sato and T. Negishi, eds., *Recent Developments in Japanese Economics.* Tokyo: Harcourt Brace Jovanovich Japan/Academic Press, 1988.

Joines, Douglas H., and James G. Manegold. "IRA and Saving: Evidence from a Panel of Taxpayers," mimeo, U.S.C., 1991.

Kotlikoff, Laurence J., Avia Spivak, and Lawrence H. Summers. "The Adequacy of Saving," *American Economic Review* (December 1982), pp. 1056-69.

Meyer, Stephen A. "Saving and Demographics: Some International Comparisons," this *Business Review,* March/April 1992, pp. 13-23.

Organization for Economic Cooperation and Development, *Economic Outlook,* 1992.

Shefrin, Hersh M., and Richard H. Thaler. "The Behavioral Life Cycle Hypothesis," *Economic Inquiry* 26 (October 1988), pp. 609-43.

Summers, Lawrence. "Capital Taxation in a Life Cycle Growth Model," *American Economic Review* (September 1981).

Summers, Lawrence. "Issues in National Savings Policy," National Bureau of Economic Research Working Paper 1710, Cambridge, MA (September 1985).

Venti, Steven F., and David A. Wise. "Tax-Deferred Accounts, Constrained Choice and Estimation of Individual Saving," *Review of Economic Studies,* LIII, 1986, pp. 579-601.

Venti, Steven F., and David A. Wise. "IRAs and Saving," in Martin Feldstein, ed., *The Effects of Taxation on Capital Accumulation.* Chicago: University of Chicago Press and NBER, 1987, pp. 7-48.

Venti, Steven F., and David A. Wise. "Aging, Moving, and Housing Wealth," in David A. Wise, ed., *The Economics of Aging.* Chicago: University of Chicago Press and NBER, 1989.

Venti, Steven F., and David A. Wise. "Have IRAs Increased U.S. Saving?: Evidence from Consumer Expenditure Surveys," *Quarterly Journal of Economics* 105 (August 1990), pp. 661-98.

Venti, Steven F., and David A. Wise. "The Saving Effect of Tax-Deferred Retirement Accounts: Evidence from SIPP," in B. Douglas Bernheim and John B. Shoven, eds., *National Saving and Economic Performance.* Chicago: University of Chicago Press and NBER, 1991, pp. 103-28.

Highways and Education: The Road to Productivity?

*Gerald A. Carlino**

From 1948 to 1969, output per hour worked grew at an average rate of 2.5 percent per year. From 1969 to 1987, growth of labor productivity slowed to 1.1 percent per year. Economists and policymakers have acknowledged that the slowdown in productivity growth is one of the major economic problems facing the United States because sluggish productivity growth means slower growth in our standard of living. The decline in investment in public infrastructure and the decline in educational quality may have played a role in this slowdown. Growth of real government spending on nonmilitary public infrastructure declined from an annual rate of 4.1 percent between 1948-69 to only 1.6 percent during 1969-87. There is also some indication that educational quality may have slipped over time as witnessed by the fact that Scholastic Aptitude Test (SAT) scores have been declining since the mid-1960s.[1]

The current Administration would like to increase national productivity by, among other things, increasing investment in public infrastructure and by creating job training programs to improve the quality of the work force. Would

* Gerald A. Carlino is an economic adviser in the Regional and Urban Section of the Philadelphia Fed's Research Department.

[1]The data reported in this paragraph are taken from Alicia H. Munnell, "Why Has Productivity Growth Declined? Productivity and Public Investment," *New England Economic Review*, January/February 1990a, pp. 3-22.

programs such as these improve productivity and ultimately the level of output?

Differences across states in investment in public infrastructure and education provide insight into the likely effects of national spending in these areas. A number of recent studies have looked at the impact of public infrastructure and educational attainment on output at the state and local levels. Studies have found that increases in highway density and educational attainment improve a region's productivity and boost output. A recent study by Carlino and Voith found that a 10 percent increase in educational attainment of a state's residents boosts its output by 8 percent, and a 10 percent increase in highway density increases state output by 1.4 percent.[2]

REASONS PRODUCTIVITY DIFFERS ACROSS STATES

Productivity measures the ratio of output to inputs such as land, labor, and capital. If two regions used the same quantities of inputs, output would be greater in the more productive region. One region might have higher productivity than another because the quality of inputs is higher. Regional productivity depends not only on the number of machines used to produce an output but also on their age, technical quality, and degree of utilization. Regional productivity may also depend on the scale at which production takes place within a region's firms. As firms increase their size, they can sometimes increase productivity by having their workers specialize in particular tasks or by using their capital equipment more efficiently. These *internal* factors may vary from one region to another and therefore may influence regional productivity.[3] While these internal factors are an important source of productivity differentials across regions, this article focuses on public infrastructure and the quality of the region's work force, factors that are *external* to the firm but which influence productivity in a market or region. Before we look at how much public infrastructure and work force quality matter for productivity, we need to understand other external factors that affect productivity, such as a region's industry mix and the degree of urbanization, so that we can control for their effects.

Industry Mix. Regional differences in productivity arise partly because individual regions often specialize in the mix of goods or services they produce. For instance, the growing of wheat and corn tends to be concentrated in the Plains states. Because many of the states in the Northeast and Midwest have historically specialized in the production of manufactured goods, this broad geographic area is commonly referred to as the "industrial belt" or "industrial core." Since some industries are more productive than others, regions with a relatively large concentration of the more productive industries will have greater overall productivity than regions with a concentration of the less productive industries.[4]

Urbanization Economies. Just as a region's industry mix can influence its productivity, the

[2]Gerald A. Carlino and Richard Voith, "Accounting for Differences in Aggregate State Productivity," *Regional Science and Urban Economics*, 2, December 1992, pp. 597-617.

[3]These internal decisions by firms may be influenced by external factors. For example, the size of a region's market (external factor) may influence a regional firm's scale of operation (internal factor).

[4]Baumol, Blackman, and Wolff looked at national productivity growth by industry during the 1947-86 period. They found that productivity growth does differ by industry. They also reported that the traditional high-productivity-growth industries continued to perform well during the 1947-86 period, implying long-term differences in the level of productivity across industries. See William J. Baumol, Sue Anne Batey Blackman, and Edward N. Wolff, *Productivity and American Leadership: The Long View* (Cambridge, MA: The MIT Press, 1989).

percentage of a region's firms that are located in metropolitan areas also affects its productivity. Metropolitan areas offer their firms access to a common pool of trained labor, so that firms not only share the cost of training new workers, but any firm can vary its work force without incurring lost productivity during training periods or by carrying idle workers. Metropolitan locations also help firms by providing wholesaling facilities that reduce the level of inventories any one firm needs to keep on hand and by providing access to accounting, data processing, legal, financial, and other specialized business services. Firms located in nonmetropolitan areas would need to employ people who provide these specialized business services on a full-time basis or else spend considerable time and money bringing them from a distance when they are needed. By locating in a metropolitan area firms can contract for these on an as-needed basis.

Economists refer to the advantages offered by metropolitan areas as urbanization economies. These urbanization economies should increase the productivity of urban firms. Thus, other things being equal, the more urbanized regions should have greater productivity than less urbanized regions. In other words, with fewer inputs metropolitan firms can produce the same level of goods and services as nonmetropolitan firms.

Urbanization economies can increase firms' productivity only up to a point. Urbanization brings not only greater productivity but also greater problems, such as congestion, that eventually balance or outweigh the efficiency gains from urbanization. At some point, increases in the number of people and firms residing in a metropolitan area clog its roads and transportation network and raise the average time and cost of transporting goods and commuting either to work or to leisure activities. In addition, as a metropolitan area grows, its boundaries may spread out, which increases both the time and distance of the average commute. When urban size becomes a hindrance rather than a help, firms experience urbanization diseconomies. Urbanization economies are balanced by these diseconomies, suggesting that there may be some optimal degree of urbanization.

Individual firms that have incentives to exploit urbanization economies are guided by the "invisible hand" of the marketplace to locate in metropolitan areas. Local policymakers can lend a hand to lessen the negative consequences of congestion by providing public infrastructure, such as highways, airports, and mass transit facilities, that link a region's labor and product markets with one another and with those of other regions.

Public Infrastructure. Some economists believe that an increase in the capital stock of the public sector leads directly to increases in private sector output because public infrastructure is an essential input in the production of private output.[5] For example, driver productivity increases when a good highway system allows truck drivers to avoid circuitous back roads and to bring supplies to a firm and goods to market more quickly. Similar arguments can be made for the public provision of police and fire protection, water supply facilities, airports, and mass transit. An increase in the public capital stock, like an increase in any factor of production, directly increases private sector output.[6]

Of course, some public sector spending may actually substitute for private sector spending. This would be the case if close substitutes for

[5]For a useful survey of the recent literature, see John A. Tatom, "Should Government Spending on Capital Goods Be Raised?" *Review,* Federal Reserve Bank of St. Louis, March/April 1991, pp. 1-15; and Randall Eberts, "Public Infrastructure and Regional Economic Development," *Economic Review,* The Federal Reserve Bank of Cleveland (First Quarter 1990a), pp.15-27.

[6]Munnell (1990a); see footnote 1 for complete citation.

publicly provided services are available from the private sector.[7] Public finance theory tells us, however, that most public sector spending should be for goods and services that would be either not provided or underprovided if left to the private sector. For example, private companies could build roads and bridges and charge tolls for using them. But private provision may not be efficient. Although there is a large initial fixed cost associated with construction of bridges and highways, once constructed, the additional cost of one more vehicle on uncongested roads is nearly zero. In this case, economic efficiency requires setting a zero price for use of uncongested roads. Thus, while it is possible to exclude those unwilling to pay for the use of infrastructure, such exclusion often is inefficient.[8] In such cases, the public sector should provide infrastructure.

Labor-Force Characteristics. Policymakers in state and local government in the U.S. have a great deal of influence on the quality of the work force because their policies affect the cost and quality of the public education system. Studies have shown that higher educational attainment of a region's labor force is an important contributor to higher regional productivity.[9] These investments in human capital may lead to increased regional productivity because education introduces a region's workers to new techniques and skills. Since educational attainment differs across regions, these differences can lead to variations in regional productivity.

THE EVIDENCE

Studies on regional productivity have tended to limit their focus to specific aspects of regional productivity. A number of studies since the mid-1970s have looked at the impact of urbanization economies on manufacturing productivity at the regional level. These studies have shown that manufacturing productivity in general increases with metropolitan population size (a proxy for urbanization economies), at least over the observed ranges of metropolitan sizes.[10] Another group of regional productivity studies has examined the role of public infrastructure in regional production, and most studies find that greater investment in public capital does raise regional productivity.[11]

[7]Studies have found that labor and public capital are complements in production, while there appears to be some degree of substitutability between public capital and private capital. See Jose da Silva Costa, Richard W. Ellson, and Randolph C. Martin, "Public Capital, Regional Output and Development: Some Empirical Evidence," *Journal of Regional Science*, 27 (1987), pp. 419-37; and Alicia H. Munnell, "How Does Public Infrastructure Affect Regional Economic Performance?" *New England Economic Review*, September/October 1990b, pp. 11-33. Munnell finds that highways and streets appear to be substitutes for private capital and speculates that well-maintained roads reduce wear and tear on commercial vehicles, lowering private sector maintenance and replacement of these vehicles.

[8]See Eberts (1990a; see footnote 5 for complete citation), for a discussion of the public goods aspects of public inputs.

[9]See, for example, Gerald A. Carlino and Edwin S. Mills, "The Determinants of County Growth," *Journal of Regional Science*, 27 (1992), pp. 39-54.

[10]For a survey of this literature, see Ronald Moomaw, "Spatial Productivity Variations in Manufacturing: A Critical Survey of Cross-Sectional Analysis," *International Regional Science Review*, 8 (1983), pp. 1-22.

[11]See Randall Eberts, "Estimating the Contribution of Urban Public Infrastructure to Regional Economic Growth," Working Paper 9004, Federal Reserve Bank of Cleveland (May 1990b). While Eberts concentrates on the influence of public capital on *manufacturing output*, an article by Alicia Munnell and one by Teresa Garcia-Mila and Therese J. McGuire extend the analysis of public infrastructure to *aggregate output* at the state level. See Munnell (1990b; footnote 7 has complete citation); and Teresa Garcia-Mila and Therese J. McGuire, "The Contribution of Publicly Provided Inputs to States' Economies," *Regional Science and Urban Economics*, 22 (1992), pp.229-41. Both studies find that public infrastructure has positive effects on aggregate productivity at the state level.

The examination of each of these factors in isolation can result in misleading conclusions. For example, the contribution of public infrastructure to regional productivity may be overstated if the other factors thought to influence regional productivity are not taken into consideration. The clustering of firms in metropolitan areas creates urbanization economies, which, in turn, increases a region's overall productivity and output. More output leads to increased tax revenue for state and local governments. Some of the increased tax revenue may be used to supply public infrastructure. Perhaps it is urbanization economies that largely contribute to regional productivity, and public infrastructure contributes to a much lesser extent or not at all. Since increased urbanization economies lead to more output, which, in turn, leads to more public infrastructure, studies that look at the role of public infrastructure on regional productivity but fail to control for urbanization economies run the risk of overstating the relative importance of public capital.[12]

The Carlino and Voith study provides a more comprehensive view of the factors affecting state productivity by considering the relative importance of industry mix, urbanization economies, public infrastructure, and labor quality on aggregate production at the state level during the 1967-86 period (see Appendix, page 30).[13]

Industry Mix. Carlino and Voith measured industry mix by the share of state output attributable to each of the nine major industry groupings.[14] By including these industry-mix variables, their study controlled for industrial structure differences across states, which helped to isolate the effects of the other variables thought to have independent effects on state productivity. Carlino and Voith found that state productivity varies a lot, running from about 50 percent above the national average in Delaware to about 35 percent below average in Wyoming. They also found that controlling for industry mix alone explains about 26 percent of the variation (see *Industry Mix Is an Important Component of a Region's Aggregate Productivity*).

Urbanization Economies. The Carlino and Voith study used the percent of a state's population that is metropolitan to capture the effects of urbanization economies. The percent of the population living in metropolitan areas varied widely across states in 1984; for example, it is as low as 14.7 percent in Wyoming and as high as 100 percent in New Jersey.[15] The positive effects

[12]An unresolved issue is whether public capital precedes private capital formation or vice-versa. There is evidence that the formation of public capital and private capital is a simultaneous process. See Eberts (1990a; footnote 5 has complete citation).

[13]Carlino and Voith (1992; see footnote 2) used multiple regression analysis to examine the relative importance of industry mix, labor-force quality, urbanization economies, and infrastructure on state aggregate productivity. One problem with analyzing the results from a multiple regression analysis is that the variables are generally measured in different units. For example, educational attainment is measured in *years*, and public infrastructure is measured in terms of *highway density*. To facilitate the comparison of the effects of different variables, we must standardize our findings. A common approach couches relationships in percentage terms—the percent change in one variable associated with the percent change in another. This unitless measure is known as an elasticity. The elasticity for state output tells us the percent change in state output given a percentage change in any of the explanatory variables, while holding all other explanatory variables constant.

[14]These groupings are agriculture; mining; construction; manufacturing; transportation, communication, and public utilities; trade (wholesale and retail combined); finance, insurance, and real estate (FIRE); services; and government. Since the industry shares of state output sum to one, it is necessary to drop the percentage share of one of the industries. Although agriculture is the excluded industry in the Carlino and Voith study, the study could just as easily have excluded any one of the other industries.

[15]Every county in New Jersey is part of a metropolitan area even though large parts of some counties are rural.

Industry Mix Is an Important Component of a Region's Aggregate Productivity

The estimates of total factor productivity from the Carlino/Voith study can be used to compare aggregate productivity across states by looking at the ratio of productivity in a state relative to productivity averaged across all states. If productivity in a state is equal to the national average, the ratio would equal one. If the state is more productive than the average state, the ratio would be greater than one. And the ratio is less than one if the state is less productive than the average state.

State productivity varies from about 50 percent above the national average (48-state average) in Delaware to about 35 percent below the national average in Wyoming (see Table). Even with the exclusion of Delaware, there is a 58 percent differential between Rhode Island, the second most productive state, and Wyoming, the least productive state. But controlling for industry mix alters the picture substantially.

Industry Mix. Total productivity was recalculated for each state, controlling for industry mix differences across states by assigning the national industry mix to each state. Controlling for industrial structure reduces the differential in total productivity across states by 26 percent. The differential in state productivity runs from about 43 percent above the national average (compared with 50 percent above average before standardization) to 19 percent below the national average (compared with 35 percent below before standardization). Of the 16 states in the top one-third of the productivity distribution before standardization, 13 states remain in the top one-third after standardization. Indiana, Maine, and Massachusetts, which were in the top one-third before standardization, moved to the middle third after standardization. Three states, Louisiana, Oklahoma, and New Mexico, were in the bottom one-third before standardization but moved to the top one-third after standardization.

Wyoming is an interesting example of how industry mix can affect a state's productivity in that it moves from being 35 percent below the U.S. before standardization to just about at the national average after controlling for industrial structure. A relatively large portion of total employment in Wyoming is in the extractive industries, especially oil and gas. Mining employment in Wyoming accounted for 22 percent of total employment in 1980, compared with only one percent nationally. Wyoming also tends to be much less manufacturing oriented. In 1980, only 6 percent of total employment in Wyoming was accounted for by manufacturing, compared with 28 percent nationally. One recent study shows that while productivity in the mining industry fell dramatically during the period 1947-86, it improved slightly in manufacturing.[a]

[a]William J. Baumol, Sue Anne Batey Blackman, and Edward N. Wolff, *Productivity and American Leadership: The Long View* (Cambridge, MA: MIT Press, 1989).

Aggregate Productivity Differences Across States[a]

		Total	Controlling for Industry Mix				Total	Controlling for Industry Mix	
1	Delaware	1.5002	1	1.4338	25	Nevada	1.0210	42	0.9058
2	Rhode Island	1.2282	2	1.1886	26	Arkansas	1.0178	19	1.0149
3	South Carolina	1.2081	3	1.883	27	Maryland	1.0177	31	0.9660
4	Georgia	1.1833	12	1.0603	28	Arizona	0.9953	36	0.9350
5	Connecticut	1.1685	6	1.1210	29	Oregon	0.9867	39	0.9144
6	New Hampshire	1.1456	10	1.0723	30	Virginia	0.9751	26	0.9908
7	Tennessee	1.1230	13	1.0595	31	Iowa	0.9687	25	0.9922
8	Vermont	1.1156	14	1.0591	32	Mississippi	0.9653	20	1.0127
9	Indiana	1.1061	17	1.0481	33	Colorado	0.9364	40	0.9102
10	West Virginia	1.1007	5	1.1609	34	Florida	0.9343	46	0.8332
11	Missouri	1.0949	29	0.9780	35	Washington	0.9289	44	0.8700
12	North Carolina	1.0913	11	1.0603	36	Kansas	0.9274	32	0.9642
13	Alabama	1.0868	9	1.0730	37	Nebraska	0.9250	38	0.9184
14	Massachusetts	1.0864	24	0.9997	38	Idaho	0.9138	41	0.9089
15	Michigan	1.0832	16	1.0506	39	Oklahoma	0.9061	7	1.1194
16	Maine	1.0830	21	1.0075	40	Utah	0.9022	43	0.8881
17	Illinois	1.0788	27	0.9870	41	Texas	0.8433	28	0.9816
18	Ohio	1.0662	22	1.0005	42	North Dakota	0.8409	35	0.9426
19	Wisconsin	1.0515	18	1.0324	43	California	0.8285	48	0.8090
20	New York	1.0450	34	0.9528	44	Louisiana	0.8236	4	1.1647
21	Pennsylvania	1.0386	33	0.9533	45	South Dakota	0.8143	45	0.8626
22	New Jersey	1.0381	37	0.9350	46	New Mexico	0.7989	15	1.0531
23	Minnesota	1.0337	30	0.9765	47	Montana	0.7744	47	0.8296
24	Kentucky	1.027	8	1.0880	48	Wyoming	0.6457	23	0.9998

[a]Index represents ratio of aggregate productivity in each state to the national average.

of increased urbanization make up one side of the urban size ledger. The negative effects of congestion brought on by increased urbanization make up the other. Thus, Carlino and Voith allowed for the fact that increasing the degree of urbanization would increase productivity up to a point, after which productivity would decrease.[16] Both forces influence productivity: increased urbanization encourages growth, and increased congestion discourages it. Carlino and Voith found that the positive effects of urbanization economies are greatest when roughly half of a state's population is metropolitan.[17]

Infrastructure. A state can mitigate the effects of congestion by building and maintaining streets and highways. The Carlino and Voith study employed highway density (highway miles per square mile of land area in a state) as a proxy for state infrastructure, partly because of the relative importance of highways and partly because data for the other categories of public capital are generally not available.[18] The study found that state productivity responds to the availability of a highway network. A 10 percent increase in a state's highway density leads, on average, to a 1.4 percent increase in total output. The Carlino and Voith study corroborates the findings reported in several recent studies in terms of the importance of infrastructure spending on state output. One study, by Garcia-Mila and McGuire, employed annual expenditures on highways by state and local governments during 1969-83 as a measure of public sector capital. The study found that a 10 percent increase in highway spending results in a 0.7 to 1.7 percent increase in aggregate state output.[19] A study by Munnell, using a broader measure of infrastructure than the one employed by Carlino and Voith, found that a 10 percent increase in infrastructure led to a 1.5 percent increase in aggregate state output during the 1970-86 period.[20] The similarity of the findings among the three studies supports the concept of public infrastructure spending as a public policy instrument for fostering productivity growth at the state level.[21]

---

[16]To capture the effects of congestion Carlino and Voith took the percent of a state's population that is metropolitan and squared it. This follows William Baumol's reasoning that if each resident of a metropolitan area imposes external costs on every other, and if the magnitude of the cost borne by each resident is roughly proportional to a metropolitan area's population size, then since these costs are borne by each of R residents involved, the total external cost will increase not with R but with R^2. See William J. Baumol, "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis," *American Economic Review*, Vol.57 (1967), pp. 415-26.

[17]Of course, factors other than percent of a state's population that is metropolitan can influence the urbanization economies states offer. For example, urbanization economies may spill over state boundaries so that states that are not highly urbanized may benefit from urbanization economies if they are near highly urbanized states.

[18]In 1988 nonmilitary infrastructure amounted to $2 trillion, compared with $4.4 trillion in private capital. Most of this infrastructure consists of assets owned by state and local governments. The largest single item is highways and streets, which account for 39 percent of total state and local wealth. See Munnell (1990b; footnote 7 has complete citation).

[19]Garcia-Mila and McGuire (1992; footnote 11 has complete citation).

[20]Munnell (1990b; footnote 7 has complete citation). Munnell found that an additional dollar of public infrastructure spending yielded the same increase in aggregate state output as an additional dollar spent on private capital. Munnell used the stock of state and local public capital, which includes highways and streets, water and sewer systems, buildings (schools, hospitals, etc.), and equipment. The results of this study are somewhat controversial. See John A. Tatom, "Public Capital and Private Sector Performance," *Review*, Federal Reserve Bank of St. Louis, May/June 1991, pp. 3-15; and Alicia H. Munnell, "Infrastructure Investment and Economic Growth," *Journal of Economic Perspectives*, 16, Fall 1992, pp. 189-198.

Labor-Force Characteristics. Differences in labor-force composition—education, experience, degree of unionization—across states can result in differences in aggregate productivity. The Carlino and Voith study uses educational attainment, defined as the percent of a state's population that is 25 years old and over with 12 or more years of schooling, as its measure of labor-force quality. The percent of a state's 25-and-over population with at least a high school diploma varies widely across the United States; for example, in 1980 it was as low as 53 percent in Kentucky and as high as 80 percent in Arizona. Carlino and Voith's results indicate that a 10 percent increase in educational attainment leads, on average, to an 8 percent increase in aggregate output.[22] This finding suggests that education is an important public policy instrument for promoting productivity growth at the state level.[23]

CONCLUSION

The research summarized in this article supports the view that increased infrastructure spending and greater educational attainment do improve productivity and ultimately the level of output. Further research should help determine the relative effects of additional spending on infrastructure and education. But the findings so far suggest that state governments should pay close attention to investment in public capital and to the level of educational attainment of their workers.

[21]The magnitude of the effect of public infrastructure on state level output is about half as large as that found for the national economy. For example, Aschauer found that a 10 percent increase in the stock of public capital led to a 3.9 percent increase in national output. See David A. Aschauer, "Is Public Expenditure Productive?" *Journal of Monetary Economics*, 23, March 1989, pp. 177-200. When one state adds to its stock of public infrastructure, this increased investment most likely has a beneficial effect on the output of neighboring states. For example, the opening of Interstate 476 in Pennsylvania in 1992 not only made Pennsylvania's workers more productive, but it may have improved the productivity of workers in Delaware and New Jersey as well. For a general critique of Aschauer's findings, see Laura Rubin, "Productivity and the Public Capital Stock: Another Look," Working Paper No. 118, Board of Governors of the Federal Reserve System, May 1991.

[22]Of course, more productive workers may place a higher value on educational attainment. To some extent, therefore, productivity and educational attainment may be a simultaneous process.

[23]Factors other than those discussed here could affect state productivity, including state policies and regulations, the degree of unionization, research and development spending, and technical progress. While these factors may determine differences in state productivity, few, if any, data are available to determine the relative importance of these omitted variables.

APPENDIX

A state's output of goods and services depends on the quantities of inputs, such as capital and labor, and on the productivity of those inputs. The relationship among output, inputs, and productivity is given in the following production function:

$$Q = AF(K, L)$$

Accordingly, the amount of real output, Q, that a state can produce during some period, such as a year, depends on the size of its capital stock, K, and the number of hours worked, L. The symbol F is a function, or equation, relating output to capital and labor inputs. The symbol A measures the overall effectiveness with which a state uses its capital and labor resources. The symbol A is therefore referred to as a measure of total factor productivity. If two states used the same levels of capital and labor, the more productive state would have a larger A term and would therefore produce more output than the state with a lower A term.

While some studies have treated the various productivity factors as inputs in the production function, the Carlino/Voith study treated them as affecting the efficiency parameter, A. Specifically, the value of A depends on industry mix, urbanization economies, public capital, and the quality of labor. This means that the various productivity factors augment private sector use of labor and capital. In this case, an increase in the level of public capital increases the efficiency of both private capital and labor.

The Empirical Model. Empirical analysis of state productivity has had to deal with an important data problem, namely, data on the stock of capital at the state level are not available. Fortunately, a production function technique has been developed that permits the estimation of productivity without the need for data on the capital stock.[a] The technique involves estimating a wage equation. It is assumed that workers are paid according to their productivity (that is, there is perfect competition in and across local labor markets), and therefore wages and the demand for labor reflect the differentials in productivity across states. Under these conditions, the following wage equation is derived from the aggregate production function:

$$\ln W_{it} = \beta_0 + \Sigma_{j=1}^{8}\beta_j S_{jit} + \beta_g P_i + \beta_{10} P_i^2 + \beta_{11} I_i + \beta_{12} E_i + \beta_{13} T_t + \beta_{14} Z_t + \beta_{15} U_i + \theta \ln Q_{it} + \gamma \ln L_{it}$$

where

W_{it} = Annual aggregate real wage bill divided by number of employees in state i for time t.

S_{jit} = The real output share of the j-th one-digit industry (mining; construction; manufacturing; transportation, communication, and public utilities; wholesale and retail trade; finance, insurance, & real estate; services; and government) in state i for time t.

P_i = For each year, the percent of state i's population living in metropolitan areas in 1970 or 1980 (whichever is closest) based on 1983 metropolitan area definitions.

I_i = Total primary Federal-Aid Highway System miles per square mile of land area in state i for 1980.

E_i = Educational attainment (percent of the population 25 years old and over with 12 or more years) in i in 1980.

[a]See Gerald A.Carlino, "Increasing Returns to Scale in Metropolitan Manufacturing," *Journal of Regional Science*, 19, 1979, pp. 363-73.

APPENDIX (continued)

T_t = Technical progress, represented by a time index.
Z_t = Dummy variable to capture the effects of the energy shock years; $Z_t = 1$ if t = 1973 to 1978; and 0 otherwise.
U_i = Union membership as a percent of employees in nonagricultural establishments in i for 1970.
Q_{it} = Real gross state product in state i at time t.
L_{it} = Aggregate employment in state i at time t.

The findings reported in the text of this article are based on a random-effects estimation of a pooled cross-section time series model for the 48 contiguous states for the period 1967-86 (providing 960 observations).[b] While a wage equation was estimated, we obtained the effects of industry mix, urbanization economies, public infrastructure, and labor force quality on output indirectly by transforming the appropriate estimated coefficients of the wage equation.[c]

---

[b]The estimated coefficients for industry mix, urbanization economies, public capital, and labor quality capture the direct effect of these variables on labor productivity. There may also be important indirect effects that are not captured by the estimates. For example, states with high educational attainment may also attract the more productive industries.

[c]Let α_k represent the output effect of the k-th productivity variable. Then the output effect is calculated indirectly as $\alpha_k = \beta_i/\rho$, where $\rho = \theta - 1$. For details see Gerald A. Carlino and Richard Voith, "Accounting for Differences in Aggregate State Productivity," *Regional Science and Urban Economics*, 22, 1992, pp. 597-617.

FEDERAL
RESERVE BANK OF
PHILADELPHIA

BUSINESS REVIEW Ten Independence Mall, Philadelphia, PA 19106-1574