

Redistributional Consequences of Educational Reform

Paul M. Romer

To an economist, the fundamental problem in elementary and secondary school education seems simple. Throughout the world, government funding for education has been unnecessarily equated with government operation of the schools. According to this interpretation, the solution is equally simple: maintain government funding but privatize the schools. Specifically, let a variety of private organizations operate schools, let parents choose among them, and let the government compensate the schools on the basis of the number and type of children that they attract and how well their students perform. On the face of it, this approach would continue to make education freely available to all children and still let a nation capture the gains associated with competition and free entry in the provision of educational services.

If this economic analysis captured the essence of the problem in education, it would surely be a simple matter to put together a broad political coalition that would support the privatization of government-run school systems. The efficiency gains from privatization would be large. Under this strategy, they seemingly can be achieved with no reduction in the commitment that a society makes to help the disadvantaged. There would, of course, be a small group of

self-interested individuals who have an economic interest in the preservation of the existing state-run system. They would resist any change but would be overwhelmingly outnumbered by parents and other concerned citizens. Yet, despite four decades of debate, proposals for voucher-style privatization schemes have made remarkably little progress in the developed world in general, and in the United States in particular. There must be something that the economic analysis misses.

To a politician and a voter, the fundamental problem in elementary and secondary school education is more complex. The gap in the economic analysis lies in its implicit assumption that it will be possible to maintain the same level of government support for the education of children from poor families under an alternative system based on private provision of educational services. It assumes that the level of public support for education and the mechanism for delivering the educational services can be chosen independently. The politician and the voter recognize, however, that the level of public funding for education is the outcome of a political process and that changes in the delivery system can induce changes in this level of funding.

As a result, the challenge that reformers face in designing a proposal for reforming education goes beyond the one assumed in the standard economic analysis. To attract broad support, a reform proposal must simultaneously increase efficiency in the delivery of educational services and maintain the level of redistribution that emerges from the political processes determining the level of support for education.

WHY SHOULD SOCIETY REDISTRIBUTE INCOME, AND WHY USE SCHOOLS TO DO IT?

Because redistribution lies at the heart of the current impasse over school reform, it is useful to step back and review the logic behind the two premises shared by the economic and

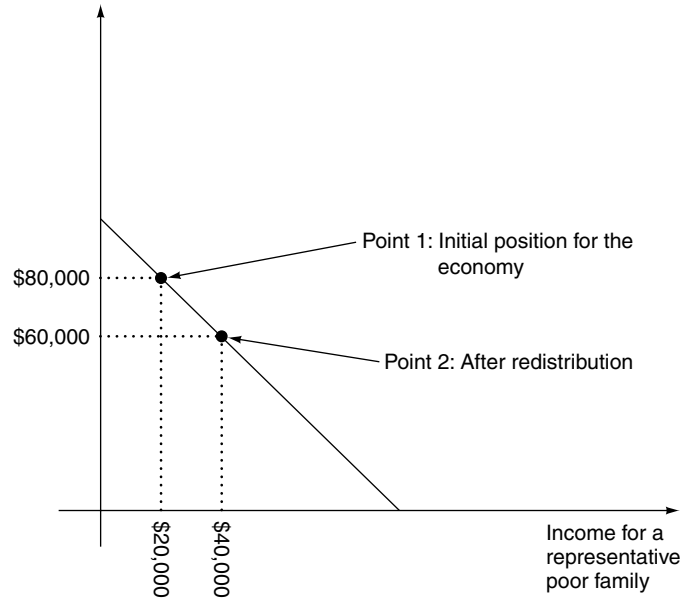


FIGURE 1. Redistribution of Income

political analyses of school reform: (1) The government should redistribute income from the rich to the poor, (2) The school system should be one of the vehicles that the government uses to implement a program of redistribution.

Figure 1 illustrates the usual framing of discussions about the redistribution of income. Point 1 illustrates income levels for representative rich and poor families. By assumption, these incomes are determined by the operation of a competitive market system.

The most fundamental result in economics tells us that, to a first approximation, the market system is efficient in the sense that it maximizes total income available to a society. This result is often referred to as the “invisible hand” theorem because of Adam Smith’s famous assertion that in a competitive market, each person “. . . intends only his own gain, and he is in this, as in many other cases, led by an invisible hand to promote an end which was no part of

his intention. . . . By pursuing his own interest he frequently promotes that of the society more effectually than when he really intends to promote it.”

In a hypothetical society made up of equal numbers of rich and poor families, total income is signified in the Figure by the position of the downward sloping line. Along this line, the sum of the income for the rich and the poor families is constant. The points on the line represent all the possible divisions of total income between these two types of families. According to the invisible hand theorem, if a government assigns and enforces ownership rights for the productive resources in an economy and lets people trade these resources in a competitive market, market forces ensure that total income is maximized. The line will be shifted out as far to the right as possible.

In political debates, this powerful result is sometimes used to support a limited role for the government, a role that extends only to establishing and enforcing property rights. This use of the theorem fails to take account of its full implications. In the process of maximizing total income and pushing the downward sloping line out to the right, the market selects a division of income among the various market participants. That is, at the same time that it determines a position for the line, it also picks a point on the line, a point such as Point 1 in Figure 1. The invisible hand theorem is silent about the merits of this particular division of total income.

There is a clear sense in which different positions of the downward sloping line can be ranked from worse to better. More total income for society is unambiguously better than less. A market system supported by a limited government has an advantage because it maximizes total income. But in contrast to the different positions of the entire line, points on a specific line cannot be ranked. Moving along a line makes one family better-off but another worse-off, and there is no scientific way to determine whether any movement along the line is good or bad. Economics as a science has absolutely nothing to say about what would be an efficient, desirable,

or appropriate distribution of income. In particular, it does not provide any basis for arguing that the particular division of income that results from market competition has any special moral, ethical, or philosophical justification.

Although economics does not give us any basis for saying that an outcome such as Point 2 in the Figure is better or worse than a less-equal outcome such as Point 1, individual people do have preferences over the distribution of income. Most people seem to prefer outcomes like Point 2, in which income is more equally distributed, to outcomes like Point 1, in which it is less equally distributed. This preference is particularly clear when a person expresses preferences over the incomes of others. Of course, people tend to prefer more income for themselves, but holding their own income constant, they prefer more equal divisions of income among others. Faced with a choice, people also seem to be willing to agree to a tradeoff between these two goals. They will sacrifice some of their income if, in so doing, they can raise standards of living for the most disadvantaged. We see evidence of this willingness both in political support for redistribution and in individual decisions about charitable giving.

There is nothing paradoxical about the claim that economics as a science has nothing to say about what constitutes an appropriate division of income but that people do have preferences over the division of income. This is analogous to the claim that economics does not have anything to say about whether people prefer drinking wine or salt water with their meals, but that people do have meaningful preferences about which is better. Economic theory does not prescribe what tastes people should have. It observes that the best way to achieve high levels of satisfaction is to give people the opportunity to make their own choices. In the same way, economics does not prescribe any particular division of income but suggests that people in a society may be better-off if they are free to make a decision about what that division should be.

The difficulty associated with a change in the distribution of income is that it is a choice that must be made collectively.

Both you and I care about the distribution of income between two other families. If this distribution is changed somehow, it affects both of us. As a result, societies must use a nonmarket mechanism for aggregating all the different preferences that people like you and me have over the distribution of income and selecting a distribution of income that will prevail.

The typical way in which this collective decision-making process operates is for people to vote. For example, they can decide by majority rule whether to adopt a system of taxes paid by the rich and transfers given to the poor. As before, imagine that Figure 1 represents income levels in a society with equal numbers of rich and poor families. If they were faced with a yes-or-no decision, a majority of citizens might vote in favor of a system that takes \$20,000 from each rich family and gives a transfer of \$20,000 to each poor family. If so, this society can move from the unequal distribution of income represented by Point 1 toward the more equal distribution represented by Point 2. Some people may have preferred more redistribution. Others may have preferred less. But many people, perhaps even everyone, may prefer a world described by Point 2 instead of one described by Point 1.

The simple argument in favor of minimal government advocates the use of the market on grounds of efficiency. A more nuanced economic argument in favor of the market outlines a two-step procedure that expands the role of the government. First, the government should maximize total income by adopting a market system. Second, it should use a system of taxes and transfers to achieve a distribution of income that is preferred by a large fraction of the votes to the distribution selected by the market. Without this second step, the members of a society would not, in general, be able to achieve a distribution of income that they prefer to the one determined by the market.

A government can be essential for achieving this outcome. If there were no collective element to the decision-making process, people could adjust the distribution of income by unilateral action—for example, by giving to

charity. People do give to charity, but because of what economists call the “free-rider problem,” unilateral charitable giving is unlikely to achieve the distribution of income that most people would prefer. This is true even if there is unanimity between all members of society over the outcome that would be preferred. Suppose, for example, that everyone—even the rich families that would have to give up \$20,000 of their income—would prefer Point 2 in Figure 1 to Point 1. That is, each rich family is willing to give up \$20,000 to be able to increase the income of each poor family from \$20,000 to \$40,000. But suppose that there are 1,000 rich families and 1,000 poor families. Suppose that the other 999 rich families have all given \$20,000 to charities that redistribute the income to poor families, so that poor families have income after transfers equal to $\$20,000 + [(999/1000) \times \$20,000] = \$39,980$. Suppose you are making decisions for the remaining rich family. If you give up \$20,000 in income, you can raise the income of each poor family from \$39,980 to \$40,000. Although you were willing to give up \$20,000 to raise the income of every poor family from \$20,000 to \$40,000, you might not be willing to give up \$20,000 to raise the income of each poor family from \$39,980 to \$40,000. Because all the other rich families are giving a portion of their income to raise the income of the poor, you may be tempted to shirk, or “free ride” on, their contributions and not to make any of your own. But, of course, the same possibility will occur to other rich families. If each rich family is free to choose whether to contribute, many of them may choose not to contribute and to leave the problem to others. In the end, this could lead to the absence of effective redistribution, an outcome that none of the rich families desired.

The key to giving a society the ability to achieve a point like Point 2 that everyone prefers is to give the rich families some way to write a binding contract among themselves, one that can avoid the free-rider problem. The essence of this contract is that, once the desired level of redistribution has

been set, everyone is obligated to contribute. In effect, each rich family is obliged to make its contribution with the knowledge that all other rich families are similarly bound. This is precisely what a mandatory tax-and-transfer system can achieve. No voluntary system of charitable donations can replicate this kind of binding agreement.

The outline of a two-step model of economic organization—establish a market system and establish a tax and transfer system to achieve the desired level of redistribution—leaves unresolved a number of ambiguities about how a society reaches a decision concerning the amount of redistribution to undertake. It also says nothing about how it is that members of a society come to have preferences over abstractions such as the distribution of income and about the definition of the group within which the distribution of income is measured. These ambiguities are central to the debate about school reform, and they will resurface shortly. But before turning to them, I will explore two further points about the connection between efficiency and the distribution of income.

THE STATIC TRADEOFF BETWEEN EFFICIENCY AND EQUITY

Figure 2 gives an expanded look at the possible effects of a system designed to redistribute income. In Figure 2, the shift from Point 1 to Point 2 implicitly assumed that redistribution could take place without any loss of efficiency. In terms of the description in that Figure, this meant that it was possible to move along the line that holds total income constant. In general, this kind of movement may not be possible. Many forms of redistribution entail a loss of efficiency.

To give a simple illustration of these efficiency losses, a tax that causes a reduction in income of \$20,000 for a rich family may generate only \$10,000 in revenue for the government. The rich family may give up a net of \$10,000 in resources, complying with the tax laws or implementing tax avoidance

strategies, and then pay an additional \$10,000 to the government. If the \$10,000 in actual revenue that the government collects is given to the poor family as a cash transfer, the family's income goes up by the full \$10,000. Yet the income of the rich family (after paying its taxes and its compliance and avoidance costs) goes down by the full \$20,000. This causes a loss of efficiency, a reduction in total income for society as a whole. In Figure 2, this is represented by Point 3, which lies below the line passing through Point 1. In the jargon of economics, we say that the taxes impose "distortions," "deadweight losses," or "efficiency costs."

A less obvious point, but one that may be at least as important for policy purposes, comes not on the tax collection side but on the transfer side. Suppose that by offering the \$10,000 transfer to the poor family, the government inadvertently encourages members of the family to undertake actions that also reduce total income in society. For example, suppose that the government pays this transfer to anyone who drops out of high school and gives birth to a child. It is possible that, in so doing, the government could dissuade some people from continuing their schooling, people who could have completed more schooling and gone on to become members of rich families. Because these people do not acquire as much education as they would have but for the transfer, they earn less. Total income for society is reduced. Hence, there can be efficiency costs associated with both the collection of the tax revenue and the transfer of income. Both the tax and the transfer can change behavior in ways that reduce total income.

Private charitable giving has drawn increased political support because of a growing perception that some government spending programs, such as cash transfers for unwed mothers, may impose higher costs for society than comparable levels of spending by private charitable organizations. When people advocate a shift away from government programs and toward charitable mechanisms for redistributing income, they make the same mistake as the

defenders of the existing public school system. They equate the financing mechanism with the delivery system. Just as it is possible to use government funding to pay private providers of educational services, it is possible to use government funding to pay private providers of social services. It is not necessary that we shift to a voluntary system of funding to get the benefits of private provision of services. This is important because voluntary systems of funding will always suffer from the free-rider problem described in the previous section.

Point 4 in Figure 2 illustrates the effect of an opposing force that could in principle be strong enough to offset the efficiency losses associated with spending on redistribution. If this effect is large enough, it can overturn the usual tradeoff between efficiency and equity. Point 4 suggests there are some forms of redistribution that can increase efficiency. Suppose once again that the government collects \$10,000 in income from the rich family and does so at a cost of \$20,000 in lost consumption opportunities for the rich family. The process of collecting the taxes still causes \$10,000 in efficiency losses.

But suppose that the government spends the \$10,000 in revenue on additional education for a worker from the poor family, and that this extra education increases the market wage of this worker by \$30,000. In this case, total income for the society can go up, even after taking account of the distortions imposed on the rich family by the tax. As suggested by Point 4, total income for the two families can increase to \$110,000, up from \$100,000. The \$10,000 efficiency cost of the tax on the rich family is more than offset by the \$20,000 net efficiency gain associated with the increased investment in education for the member of the poor family.

The numbers here are not intended to be realistic, and they skip over the issue of how to compare a one-time expenditure on education with a lifetime of higher earnings. If you think about the analogy with investment in financial assets, you can see that the right way to do this would be to

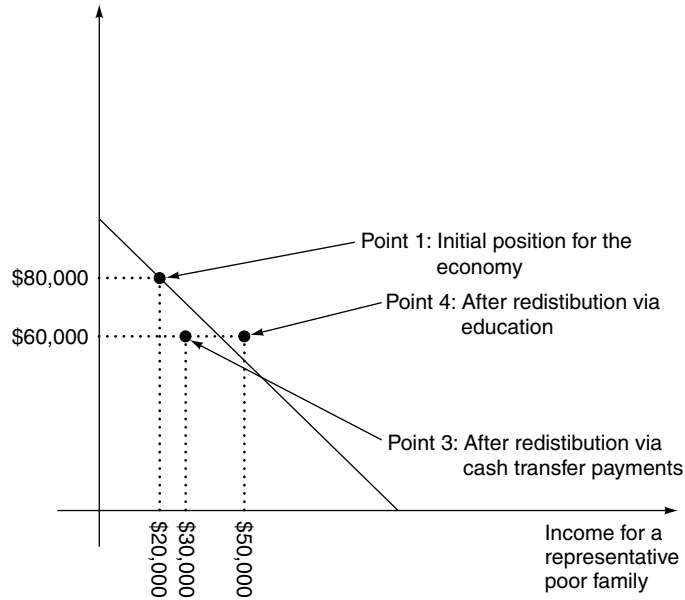


FIGURE 2. Possible Effects of System to Redistribute Income

calculate a rate of return on the investment in education. All that matters for the discussion here is that the rate of return on investments in education may be very high, so high that society as a whole is better off if the government increases its support for education.

This kind of outcome, one in which the government raises total income for society by subsidizing education, can arise only if the invisible hand theorem fails. Recall that this theorem says that the market will maximize total income all by itself, with no intervention by the government. The logic of the theorem, applied to this situation, is as follows. If it is possible for a poor worker to earn benefits from education that exceed the costs of acquiring the education, then self-interested poor workers should arrange to receive the extra education. If the market worked perfectly, the poor worker would be able to borrow money, spend it on education, use the higher wages that result to pay off the loan, and still have extra income left over to improve the quality of life.

There are two well-known reasons why the market mechanism might not lead to sufficient investment in education and might thereby cause the invisible hand theorem to fail. The first problem is that a poor worker might not be able to borrow to finance education for herself (or her children). Lenders might not be able to collect on loans made to finance education because there is nothing that they can repossess if the borrower defaults. As a result, they are unwilling to lend. It is because of this first problem that the government offers guaranteed student loans to children from poor families who want to attend college.

A second and more difficult problem, one that is much more important for an analysis of elementary and secondary schooling, is that the educational investments need to be made on behalf of a child. School-age children may not be either legally or intellectually competent to make an informed decision about the costs and benefits of an investment that will pay returns for decades into the future. Unfortunately, some parents may not be willing or able to make the investments on their children's behalf that are required to achieve full efficiency. In effect, the parents may be in the same position as the bank. It would be efficient for them to finance an investment in education for their children and then to have their children repay them later in life. But like the bank, the parents may be incapable of collecting on investments such as these that pay off much later in life.¹ For this reason, governments have not only financed educational expenses but have also made attendance at primary and secondary schools mandatory. Governments require by law that parents send their children to school.

THE DYNAMIC TRADEOFF BETWEEN EFFICIENCY AND EQUITY

Section 1 presented the justification for the first of the two premises from which all discussions of educational reform proceed. Governments redistribute income because voters care

about inequality. Section 2 addresses the second premise. Governments use subsidies for education as a mechanism for undertaking the redistribution because it is a form of transfer to poor families that is likely to be less harmful to efficiency than other transfer mechanisms, such as cash transfers. Subsidies for education may even improve efficiency.

Each of the justifications for these premises is based on a static analysis, one assuming that the underlying structure of the economy does not change. The next step is to think about how changes in the economy affect the analysis. As it turns out, changes in our economy reinforce both presumptions: that the government should redistribute income and that support for education is an efficient way to achieve more equal economic outcomes in a society.

Figure 3 illustrates the situation that now confronts the developed countries of the world. Points 1 and 5 illustrate the positions of the economy before and after a change in its structure. In moving from Point 1 to Point 5, the sum of income for the two representative families goes up from \$100,000 to \$140,000, but the distribution becomes more unequal. It shifts from a split of \$80,000 and \$20,000 to a split of \$130,000 and \$10,000.

This kind of change in the structure of an economy presents society with a more important tradeoff between efficiency and inequality than the one that is assumed in the static analysis from Figure 2. Citizens in the world illustrated in Figure 3 can increase both efficiency and inequality by letting this change in the structure take place, or they can reduce efficiency and inequality by blocking the change. If these are the only choices that are available to an economy, the tradeoff between efficiency and redistribution is stark. Out of a natural concern for their own well-being and that of their children, parents from a poor family will naturally try to block this kind of change. So might rich, altruistic families.

One would hope that economic progress did not present society with this kind of dilemma, but unfortunately two of

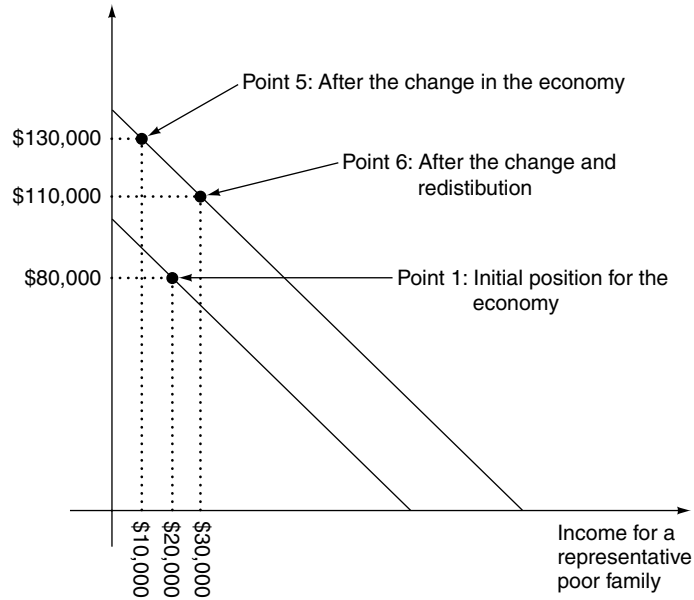


FIGURE 3. Situation in the Developed Countries

the most important driving forces in the modern economy have precisely this character. Both technological change and increased trade between developed and developing nations are likely to increase inequality at the same time that they increase efficiency.

Evidence from labor markets suggests that technological change increases the demand for highly educated workers and reduces the demand for less educated workers and has done so since the early years of the twentieth century.² To appreciate how this kind of process can arise, picture the effects of technological change in telephony. When telephone systems were first installed, the phone company employed an operator, with little schooling or technical training, who put plugs into a switchboard to connect calls. Then engineers working for the phone system developed electromechanical switches that could connect lines in response to the electrical

pulses sent by the dialing phone. This technological advance raised efficiency and reduced the cost of phone service, but it reduced the demand for relatively less-skilled operators who formerly made the connections. It also increased the demand for the higher-skilled workers who could design and maintain the electromechanical switches.

More recently, computer switches have replaced these electromechanical switches. New technologies such as voice recognition are further reducing demand for the less-skilled workers who currently provide services such as directory assistance. These technologies are raising the demand for highly skilled workers who can install and maintain these complicated computer switches.

These changes in the technology for providing phone service are representative of changes that have taken place throughout the economy. In manufacturing, and especially in agriculture, technological change has steadily reduced the number of less-skilled workers needed to produce a given amount of output.

If the relative supply of different types of workers had stayed constant over the course of the twentieth century, the demand shifts caused by technological change would have steadily increased the wages of the more-educated workers and reduced the wages of the less-educated workers. In this sense, technological change seems to present society with the dilemma illustrated in Figure 3—but only *seems*, because there is a way out. Fortunately, the relative supplies did not stay constant. The various levels of government in the United States made massive investments in elementary, secondary, and tertiary education and dramatically reduced the fraction of the labor force with low levels of education. As a result of these investments, the economy in the United States moved not to a position such as Point 5, but instead to a new position such as Point 6 in Figure 3. By itself, technological change would have moved the economy toward a position such as Point 5, but the combination of technological change

and large government-financed investments in education moved the economy to a position such as Point 6.

Of course, one can argue that in principle, people could have made the same kind of investment in education even if the government had not financed it. This is the parallel to the argument noted above that poor people should always be able to borrow and finance the efficient level of education. However, the historical correlation of high levels of school enrollment with high levels of government support for schooling suggests that government funding does lead to substantially higher levels of educational attainment, and hence to less income inequality than would prevail without government support.³

The available evidence points to a continuation of the trend from the twentieth century. Technological change will continue to raise productivity and increase total income. But acting alone, it will tend to do so in ways that lead to concentrated income gains for the best-educated families and quite possibly to losses for families with the lowest levels of education. This sets up what Katz and Goldin (1996) called *a race between technological change and government investment in education*. If the government can keep increasing the average level of educational achievement in a society rapidly enough to keep up with technological change, it can ensure that economic growth leads to gains for all segments of the population. Even the workers who remain less educated will face rising wages over time because the investment in education means that there will be fewer less-educated workers such as themselves competing for the jobs that they can do. In effect, the government can raise wages for workers with less education by creating a shortage of these workers.

If the government falls behind in this race, income inequality will widen. During the 1970s and 1980s, government investment in education did not seem to keep up.⁴ Wages for workers with low levels of education remained

constant or fell at the same time that wages for workers with high levels of education continued to grow.

The challenge posed by a process of technological change that increases the demand for highly educated workers is larger than many people realize. It will not be enough to return our school system back to the level of achievement of an average or even good school at some point in the past. During this century the United States will need to increase overall educational attainment dramatically just to keep up with the ongoing process of technological change. If it fails to do so, a growing segment of the population will suffer a reduction in their economic opportunities. As a result, many people will naturally support proposals designed to slow progress and limit trade.

The other change in the economy that will have the same kind of effect as technological change will be increased trade between the rich and poor nations of the world. Technological change will continue to reduce transportation and communication costs. Poor countries will increasingly recognize that it is in their interest to trade with the rich countries of the world. Because the poor countries have relatively abundant supplies of less-skilled labor, supplies that are very large compared to the size of the workforce in the developed world, the effects of this increased trade will be to put additional downward pressure on wages for less-skilled workers in the developed countries.

The effects on the distribution of income will be even larger if workers with low levels of skill are free to migrate to countries such as the United States. In effect, the education-based strategy followed in the United States for more than a century has been to create a shortage of workers with low levels of skill. This raises the wage for a less-skilled worker in the United States relative to what it would otherwise have been. As long as there is a large pool of less-skilled workers in the rest of the world, free international migration of workers can completely undercut a strategy of driving up

wages for less-skilled workers by making this type of worker relatively scarce.

In all likelihood, international migration will be tightly regulated in years to come. But one would hope that the processes of trade and globalization could proceed. The efficiency gains that remain to be exploited are too large to forgo. If these gains are not realized, firms in the United States will increasingly find that they can shift production activities that require only high school-educated workers to parts of the world where high school-educated workers earn far less than they do here. If these firms do not do so, they will find that they are driven out of business by new firms that spring up in these areas.

If, by the end of this century, most workers in the United States have high levels of education, this process need not lead to increased wage inequality within its borders. Assuming that barriers to immigration persist, the relatively few workers in the United States who have low levels of skill will be well compensated for doing the less-skilled jobs that are difficult to move abroad. But if the United States falls behind in the race to raise education levels and continues to have large numbers of workers with only a high school education, these workers could see their wages fall toward the level that prevails in the rest of the world as trade expands.

These changes—technological progress and increased international trade—mean that countries such as the United States will face a very different kind of tradeoff between efficiency and income inequality than the static tradeoff illustrated in Figure 1. If large numbers of citizens find that these twin forces are reducing their income and that no offsetting policy response moderates these effects, they will use the political system to slow down these forces. Even affluent families may support measures designed to limit technological change and trade if they think that this is the only way to preserve existing levels of income for poor families. Resistance to free trade, which simmers just below the surface in

domestic political debate, could quickly become a powerful political force. Legislation designed to protect the income of less-skilled workers could easily act as a brake on the adoption of new technology.

In Figure 3, if a poor family is offered a choice between Points 1 and 5, they will surely opt for Point 1, even though it means that the economy fails to take advantage of large potential efficiency gains. A society may be able to maintain a political consensus in favor of both technological progress and free trade only if it can find ways to transfer some of the gains that these forces can generate to workers from poor families.

THE REQUIRED INVESTMENT IN EDUCATION

U.S. citizens can continue using education as a way to create a relatively homogeneous society, one that does not suffer from the extremes of income inequality that we see in other countries of the world. If this happens, it should be possible to maintain a political consensus in favor of continued technological change and increased international trade. But this path will not be easy. The education race is going to become more challenging in the coming century, and there are clear indications that our educational institutions are not performing well.

To show just how hard this challenge is, suppose that the near future resembles the recent past. In 1940, only 25 percent of the population in the United States had completed high school. Sixty years later, in 2000, a high school degree is considered the minimum acceptable level of education for a new entrant in the job market. Over 80 percent of the adults in the United States now have a high school degree.

In 2000, the degree held by the top 25 percent of the population was a tertiary degree—either a bachelor's degree from a college or university or an associate degree from a junior college or vocational school. If past trends continue,

in sixty years a tertiary degree or its equivalent may become the new minimum standard for entry into the job market. By then we may need to aim for universal completion of a tertiary degree just as we now aim for universal completion of high school. Our primary and secondary schools will therefore have to prepare virtually all students for advanced study. To do this for students from the bottom half of the achievement distribution, our school system will have to do a much better job.

The available indicators suggest that the United States at present does a particularly bad job of educating students from the least-advantaged backgrounds. The Organization for Economic Cooperation and Development (OECD) has used a standardized set of materials to test the prose, document, and quantitative literacy skills of a representative sample of adults in most of its member countries. The average score in the United States is comparable to that in most other countries but there are many more people here who have very low levels of basic literacy skills, more than other nations at a comparable level of development. For example, reading abilities at the 5th percentile in the distribution are markedly lower in the United States than in any country in Western Europe (Figure 2.1).⁵

When the data are analyzed by cohort, it also becomes clear that most other countries have been able to increase mean literacy scores from one generation to the next. For example, people in Canada who are between the ages of 46 and 65 and who were graduated from high school between thirty and fifty years ago perform, on average, at about the same level on all three measures of literacy as people in the United States who were graduated from high school at the same time. Over time, schools there seem to have been doing a better job. People in Canada who were graduated from high school more recently score at higher levels than the people from older cohorts.

In the United States, there is no such evidence of improvement. People who have graduated from high school

more recently do not do any better on the literacy tests than the 46–65-year-olds who graduated from high school thirty to sixty years ago. In this regard, it is the United States, not Canada, that is unusual. Other countries such as Poland and Hungary, which have lower levels of literacy for older cohorts than the United States, and countries such as the Netherlands and Sweden, which have higher levels of literacy for older cohorts than the United States, show big improvements in the abilities of the people who were graduated from high school more recently.

ALTERNATIVE MECHANISMS FOR REDISTRIBUTING INCOME

It will be hard to keep up in the education race as this century advances. It will also be hard to implement true educational reform in the context of an education system that is intended to reduce inherited inequality. If there were a viable alternative that would preserve something like the current distribution of income, one that would make sure that all families share in the coming gains from technological change and increased trade, it would make sense to adopt it. If we could relieve education of the burden of being our primary program for reducing inequality, it would be much easier to reform the educational system and raise the level of its performance. Unfortunately, there are no good alternatives for redistributing income.

As the discussion in Section 2 has already suggested, direct cash transfers have proven to be an unworkable way to deal with growing income inequality because of the changes in behavior and the high efficiency costs they induce. The policy experiment with more generous cash support levels that was undertaken in the United States and Europe in the last three decades has convinced voters and policymakers that these transfers reduce employment, encourage antisocial behavior, and create an intergenerational cycle of dependency that is costly for society as a whole. The only

exception to this general rule may be the case of cash transfers to the elderly, which are not perceived to have high social costs.

Governments could replace unconditional cash transfers to young workers with subsidies for low-wage employment. In the United States, the imposition of time limits on eligibility for welfare payments, together with a more generous Earned Income Tax Credit, has had the effect of shifting our tax and transfer system in this direction.

In contrast to government support for education, which takes decades to have its full effect, a wage subsidy for low-income workers has the advantage that it can raise take-home pay quickly. But over a longer time horizon, the relative merit of the two strategies shifts. Over time, wage subsidies will have trouble keeping up with the cumulative effect of the forces that will impinge on the economy.

Suppose that, during the next century, technological change and increased trade will lead to an average rate of growth of real gross domestic product per capita of about 1.5 percent per year. This means that after correcting for inflation, total income from all sources, divided by the population, will increase by about 1.5 percent per year. To keep the argument here conservative, this rate is less than the 1.8 percent real rate of growth sustained in the twentieth century. Because wages tend to be a constant fraction of GDP or total output for a society, this means that, on average, total wages will also increase by 1.5 percent per year. Suppose that in the absence of any policy that offsets the effects of technological change and trade, this average increase would take the form of no wage growth for high school-educated workers and rapid wage growth for highly educated workers. (In the 1970s and 1980s, wages for high school-educated workers actually fell, so this is not an unreasonable assumption.) Through the power of compound interest, GDP per capita would double in fewer than fifty years. This means that, by 2050, the wage subsidy for high school-educated workers would have to be more than 100 percent to be able to give the high school-educated

workers the same doubling of after-tax income that the more highly educated workers would enjoy. This would mean that the marginal tax rate faced by all workers would have to be very high just to finance the transfers required to keep the relative distribution of wages constant.

The problem that arises here is common to all wage subsidy or negative income tax proposals. To have a large effect on the distribution of income, they have to rely on very high marginal tax rates. Because the distortions associated with income taxes grow very rapidly as the marginal tax rate increases, this type of tax and transfer system would impose extremely large efficiency losses. It is highly unlikely that it would prove to be politically acceptable. Moreover, with the passage of time, the efficiency costs would grow and political support would fall as the required marginal tax rate grew.

In the long run, the steadily increasing marginal tax rates make a wage subsidy a very inefficient way to deal with pressures on the distribution of income that will arise from increased trade and more technological change. The flip side of this observation is that a subsidy for education becomes even more efficient as time passes. The rates of return to education are already very high and the level of educational attainment very low for children from disadvantaged backgrounds. The returns to more education for them seem to be far higher than the cost of the investment in the education. Hence, even now, government attempts at redistributing income that take the form of increased support for education could raise efficiency rather than lower it. And with the passage of time, the same forces that are increasing income inequality—trade and technology—will further raise the return to investments in education. Education becomes an even more efficient mechanism for raising income for people in the bottom half of the distribution of income.

As a result, on purely economic grounds the most sensible policy going forward is one that continues our current policy of using the educational system to give children from

disadvantaged backgrounds better economic opportunities. The immediate challenge this poses is to find a way to reform the educational system and raise its efficiency without undermining the commitment to helping the disadvantaged. The proposal outlined in the beginning—maintaining government funding for education but privatizing the provision of educational services—seems at first to be the appropriate response to this challenge. Why, then, have voters been so hesitant to adopt this kind of reform?

MEDICARE AND MEDICAID

To understand the risks that some thoughtful people perceive from the adoption of a privatized system for delivering educational services, it helps to invoke an analogy with health care. Medicare, the health-care financing system that applies to the elderly, resembles the current public school system. Everyone over age 65 is eligible to participate and virtually everyone does. In contrast, people under the age of 65 are covered by a split or “two-tiered” system consisting of private medical insurance for the fortunate and the means-tested Medicaid program for the poor. The shift from the existing public school system to a voucher-based system could imply a shift from something like the universal Medicare program that we use for the elderly to something like the private insurance–Medicaid hybrid that we use for everyone else. Judging from the experience with these two programs, this shift would imply a substantial reduction in the commitment to helping the disadvantaged.

Medicare and Medicaid both cover the costs of medical care for the poor. The striking thing about the two programs is that Medicaid does so at levels of funding that are much lower than those offered by the Medicare program. In 1998 in the United States as a whole, the average reimbursement for a service provided under the Medicaid program was equal to 64 percent of the reimbursement for exactly the same service offered by the Medicare program.⁶ Over time, this gap has been

widening. Because funding decisions for Medicaid are made by the states, reimbursement rates also vary from state to state. In some states with large populations of patients on Medicaid, such as New York and California, reimbursement for a specific service such as an appendectomy under the universal Medicaid program for the elderly is less than half the reimbursement under the Medicare targeted at the poor.

The question this raises is why voters would select a health-care financing system that offers dramatically lower reimbursement rates to a poor person who is less than 65 years old than it offers for the same person when she is more than 65 years old. From the point of view of redistribution, the gap is perverse. The Medicaid program is means-tested and offered only to the most disadvantaged members of society, people who typically have more serious health problems than the rest of the population. Medicare is offered to everyone above the age of 65. A preference for redistribution toward the poor would argue for higher government reimbursement levels for services targeted at the poor, not lower levels.

When people in the Roosevelt administration began work on a comprehensive social welfare system for the United States, they had a slogan that guided their efforts: "A program for the poor is a poor program."⁷ The first program that they implemented was the Social Security system for funding old-age pensions, but they anticipated that the full range of government programs would eventually grow by stages to include coverage of health-care costs, first for the elderly and then ultimately for the entire population. They and their successors were successful in creating universal service programs for Social Security and Medicare but were not able to institute a health program that covered the entire population. They settled for a health program "for the poor." They believed that despite the disadvantage of higher cost, universal service programs would provide more redistribution than programs that targeted a narrow group of poor or disadvantaged citizens. The experience with the Medicare and Medicaid programs

suggests that their intuition about the politics of redistribution was correct.

A simple model that would seem at first to rationalize this observation is one based on the idea that under a program with universal coverage, rich and poor families have to consume the same level of service. Rich families will then have an incentive to lobby for a higher level of service than they would if they did not receive the service themselves. Because the service levels for all families are the same, the rich families are thereby forced to undertake more distribution than they would have selected if they could have unbundled the choice of their own service level from the level that they would provide to the poor.

This forced-redistribution model explains why it is that a universal service plan leads to more redistribution, but it does not explain why the underlying forced-redistribution plan is politically popular. If such a program were in place, it would force more redistribution than voters would otherwise select. But if voters are also given the option of shutting down the forced-redistribution plan and substituting a plan that lets them unbundle their own choices from the options that they offer to the poor, they should eagerly vote to do so. They would get rid of the universal service plan and replace it with a two-tiered, or multi-tiered, plan that has less redistribution.

This leaves a paradox. If the forced-redistribution model does not apply, this must mean that a majority of voters approve of the level of redistribution that is built into a universal service program such as Medicare. But if this is true, why would their altruism diminish when they consider the well-being of the young poor served by the Medicaid program? A preference by voters for universal service programs seems particularly hard to understand when we consider the fact that universal service programs are inherently inefficient ways to help the disadvantaged. A targeted program can avoid the efficiency cost of having a program that takes taxes away from rich families and then turns around and gives them services that they could have pur-

chased on their own. Because there is no net redistribution to the rich, it would be more efficient for them to buy their services directly.

If a broad majority of people support the Medicare program because of the high level of redistribution that it provides, why is it not possible to transfer this political support to a targeted program for the poor? Could it provide the same level of redistribution at lower efficiency costs? For example, in the context of health care, why is it that people vote to sustain the Medicare program with its high level of redistribution but are not willing to fund the Medicaid program at comparable levels?

If voters have a “self-control problem,” this behavior could be rationalized. Voters could argue that they want to provide a higher level of redistribution. They could also recognize that, if they were to make decisions about redistribution in the future in the context of choices over funding levels for a program that targets the poor, they will not support the level of redistribution that they want to select now. They could recognize that when they vote in the future about funding levels for a program that provides service to themselves, they will support higher levels of services, hence higher levels of redistribution. So as a means of tying their own hands, the voters create or maintain the universal service program.

At first, this kind of inconsistency in a voter’s behavior seems hard to accept. If a thoughtful voter is in favor of redistribution now, why won’t she be in favor of redistribution in the future? But this kind of inconsistency becomes easier to understand if the degree to which all voters support redistribution is a function of the amount and quality of the information that they take in about the circumstances of recipients of government aid. Suppose that someone is personally acquainted with a child who is suffering from malnutrition and sees this child face-to-face on a regular basis. This person will be very likely to give direct assistance to the child and to vote in favor of government programs that will help this child and other children who are in simi-

lar circumstances. But if, in contrast, the same observer knows in the abstract that there are children in other countries of the world who are suffering to a similar degree, the observer will typically not be willing to give as much aid to feed one of these “faceless” children.

The emotional immediacy that comes from direct face-to-face contact with people who are suffering has a very different impact on human behavior than abstract knowledge about the existence of suffering. This differential sensitivity is probably an inherent and unavoidable feature of human psychology, one that leads to striking differences in how different people see the world. To an aid worker who is distributing food to malnourished children in a country suffering from famine, it may seem monstrous that families living in the rich countries of the world spend more on food for their pets than they contribute to food for starving children. It may seem even more monstrous that these families are not even willing to support government foreign aid programs that would collect less tax from them than the amount they spend on pet food. If you consider the circumstances of the hypothetical child and the pet from the same emotional distance, it seems wrong for such a family to put the welfare of the pet ahead of the welfare of the child. But if you imagine yourself in the circumstances of the rich family, it is easy to see how normal people can behave as they do.

Any moderately educated person knows at some abstract level that there are children throughout the world who suffer terribly from deprivation that could easily be alleviated with small amounts of expenditure. We are not surprised by the fact that these people continue to lead ordinary lives, acting as if they did not know this fact, taking no step to reduce the suffering. However, because we understand the strong emotional effect that more immediate contact has on feelings of altruism toward others, we would expect that this person would respond very differently if a starving child

lived next door or attended the same classes at school as their children do.

According to this kind of model, more direct personal experience with the circumstances of others leads to a stronger empathetic response to their situation. Under this model, most voters are willing to support a more generous Medicare payment system for the elderly because they have more direct personal experience with this program. Many voters in the United States have either experienced the Medicare system or know someone who has—a parent or a relative. Far fewer have the same kind of first- or second-hand experience with the grimmer Medicaid program. In the same sense that they know that there are children in parts of the world who will die from malnutrition, voters may know that there are people in the United States who receive very low levels of medical care under the Medicaid program. But in each case, this abstract knowledge has the same, greatly attenuated effect on their emotions and therefore on how they vote.

A representative voter with these kinds of experiences might therefore vote to increase spending on Medicare patients but not on Medicaid patients. She might also vote to preserve the universal coverage of the Medicare system. She realizes that were it to be converted to a means-tested program such as Medicaid, future voters like her would not provide as much support for the program, and the people that she knows and cares about would suffer.

This model does not presume any form of irrationality on the part of the voter. She may understand full well that, if she were to spend time with starving children in a foreign country, she would also vote to raise foreign aid to help these children. But in her current position, one that lets her maintain a large emotional distance from the foreign children, she has no reason to change how much of their world she experiences and how she votes on their behalf. In the same way, she has little direct experience with the welfare of poor people who

use the Medicaid program and has no reason to vote to increase support for this program.

In contrast, she has detailed knowledge of the Medicare system through the experiences of the people she knows who have used this system. Because she knows what it is like for her affluent aunt to pay for her own prescription drugs, she can empathize with the burden that payments for prescription drugs impose on people like her aunt who are covered by the Medicare system. Hence she will vote to maintain the Medicare system and may even vote to expand it to cover prescription drugs, even as she votes to constrain costs in the much-less-generous Medicaid program.

Under this model, the amount of empathy any voter feels for someone else depends on how much the voter knows about the circumstances of that other person and how the voter came to know it. The more detailed and immediate information the voter has, the more readily she can imagine and feel what it would be like for the other person. When she feels more empathy, the voter will be more likely to support redistribution programs that help a disadvantaged person whose circumstances she understands.

VOTING OVER VOUCHERS

At present, many voters have direct exposure to the problems of the public school system because their children, or the children of people they know, attend these schools. In this sense, public education resembles Medicare. It is nearly a universal service program. Currently, 89 percent of primary and secondary school students in the United States attend public schools. As a result, voters as a group have more direct information about the public school system and care about the well-being of children in this system.

Voters who care about the inadequacies of the existing school system might reasonably fear that a move toward a privatized system for providing education will cause the representative voter to have less information about the educa-

tional circumstances of poor children. Most voters now have very little feel for what it is like for a poor child to be cared for under the Medicaid program. Their children and the children of most people they know are part of a very different program.

In the same way, a voucher-based system could lead to a larger separation between the educational experiences of the poor and the experiences of most other children. Under most proposals, affluent parents would be free to make additional payments for the schooling that their children receive from a private provider. As a result, affluent families will naturally tend to supplement any voucher payment and spend more on the education of their children than poor families. Schools will naturally cater to different income groups and offer higher levels of educational service to children from richer families. Over time, as these families lose direct contact with the school experience of the poor, the low quality of schools that are available to poor children in the United States will become an increasingly abstract problem, one like the problem of starving children in some foreign country. “Yes, it’s a shame, but what can anyone do?”

These kinds of concerns are reinforced by the details of the specific voucher initiatives that have been offered to voters. These proposals typically offer a voucher payment that is substantially below what the government now spends on the average public school student. Judging from this evidence, it seems reasonable to project an evolution toward a system that offers very low levels of funding for the education of the poor and that offers higher quality education for the children of families that can afford to pay more.

As the nation evolves down this path, it could easily reach an equilibrium where the majority of middle- and upper-income families send their children to a private school that they pay for with a combination of a voucher worth as little as \$2,000 or \$3,000 per child and with a sig-

nificant tuition payment of their own. As they lose contact with the rest of the school system, they will have no reason to support increases in funding for the public school system or increases in the size of the vouchers. Over time, children from very poor families will eventually face an unattractive choice between attending a private school that has to pay for its facilities and teachers on the basis of the meager voucher payments alone; or attending a public school system that suffers from a similarly low level of funding, potentially a much lower level of funding per student than the one that prevails now.

It is entirely possible that at any given level of funding, private schools for poor children will do a better job of educating poor children than the existing public school system. It is also possible that increased competition from private schools will encourage the public schools to do a better job of educating children. But at some point, reductions in funding will start to undercut the efficiencies that can be achieved through reform. As funding levels fall, the quality of schooling must ultimately fall with them.

CONCLUSION

For thoughtful voters, it would be a gamble to initiate a process that leads to a privatized provision of educational services. Many would agree that holding expenditure constant, privatization and competition would lead to significant gains in the productivity of our schools. The offsetting risk is that funding levels for the education of the poorest children might erode significantly. The gamble is whether the productivity gains will be big enough to offset the likely reductions in spending on the education of the poor. There is no question that under the existing system, children from poor families do not receive a good education in this country. However, it is possible that their educational opportunities could become even worse under a move to privatized provision of educational services. The negative effects of

spending cuts could dominate the positive effects of competition and choice.

It is quite likely that children from middle- and upper-income families would benefit from a switch to the privatized system. They would get the advantages of reform, and their families could offset the reduction in funding by making their own payments for the education of their children. Parents from these families would see the beneficial effects of the change in the experiences of their children and of children that they know. They would have little direct information about the educational circumstances of the poor. As a result, the switch could lead to a new, politically stable equilibrium in which our system provides less redistribution through our educational system. It could be this prospect that makes many voters hesitate about making a fundamental change in our educational system.

Some thoughtful proponents of a privatized school system have at least begun to think about how to achieve the benefits of competition and choice in ways that do not pose the same political risks to long-term funding levels that simple voucher systems could pose. The most obvious modification along these lines is a voucher system that targets only the most disadvantaged students. Under this approach, there is still a risk that the level of support offered through the voucher may be far too low to support the entry of entirely new providers of educational services. Even if the program begins with good intentions, it could easily end up resembling the Medicaid program. However, if this proposal did not have any adverse effect on the public school system, poor families could always return to the public schools if funding levels fell in the voucher-based system. In effect, it would be like giving the poor the option to be covered under either Medicaid or Medicare.

But even these targeted voucher systems pose a risk to existing public school systems. Now, it seems to be impossible to get a majority of voters to vote in a referendum for a switch to a voucher-based system that helps large numbers

of middle- and upper-income families escape from the public schools. If a limited voucher system were in place, it might be possible to win a series of much-lower-profile legislative votes or administrative decisions, each of which expands the number of families that are eligible to collect the vouchers. Both sides in the political fight over targeted voucher programs for the poor seem to understand that what is being contested is not simply a program for poor children, but the initiation of a dynamic process that will ultimately move far beyond its initial focus.

It may seem excessively pessimistic to worry that political institutions in the United States will lead to an outcome in which children born into poor families receive schooling that in relative terms, and perhaps even in absolute terms, is even worse than it is now. However, one must take full account of the fact that virtually all the decent, ordinary citizens of the United States already are reconciled to the fact that there are millions of children throughout the world who receive abysmal levels of schooling, nutrition, and health care, and millions of children in the United States who fail to receive a level of medical care that is taken for granted in most developed countries of the world.

As a matter of pure logic, it is not clear why affluent voters should treat the disadvantaged children who live in the United States any differently from the disadvantaged children who live elsewhere in the world. In the abstract, it is not clear why they should support relative levels of spending for the education of rich and poor children in the United States that are more equal than relative levels of spending for medical care. Our current institutions force a level of interaction and exposure to a public school system that has important elements of commonality between the rich and poor in our own country. These institutions may force more equality in educational attainment than would prevail in their absence. Changes in our institutions could have a major impact on our commitment to education as a means of redistributing income within our nation.

Citizens of the United States who care about disadvantaged children living here have reason to fear that a move toward privatized education will lead in the future to less reliance on education as the mechanism for reducing income inequality. They recognize that the existing educational system is deeply flawed and serves students from disadvantaged backgrounds very badly. They are, nevertheless, correct that there is a substantial risk that existing proposals for privatization could significantly reduce the level of public funding for education at the same time that they generate true efficiency gains.

To persuade a majority of voters to embrace change, proponents of privatization should look much more carefully at the long-term effects that their proposals could have on levels of public funding for education. They should think more creatively about institutional mechanisms that could both capture the gains from privatization and preserve a broad political commitment to the use of education as a means of making progress toward our goal of giving every child from every family the chance to thrive in the economy of the future.

NOTES

1. Edward Lazear, "Intergenerational Externalities," *Canadian Journal of Economics* 16 (May 1983): 212–28.
2. Claudia Goldin and Lawrence F. Katz, "Technology, Skill, and the Wage Structure: Insights from the Past," *American Economic Review* 86, 2(1996): 252–57.
3. Claudia Goldin and Lawrence F. Katz, "The Origins of State-Level Differences in the Public Provision of Higher Education: 1890–1940," *American Economic Review* 88, 2(1998): 303–8. Claudia Goldin and Lawrence F. Katz, "The Origins of Technology-Skill Complementarity," *Quarterly Journal of Economics* 113, 3(1998): 693–732.
4. Lawrence Katz and Kevin Murphy, "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics* 107(1, 1992): 35–78.
5. OECD (Organization for Economic Cooperation and Development), "Literacy in the Age of Information. Final Report of the International Adult Literacy Survey," Paris 2000.

6. Stephen Norton, "Recent Trends in Medicaid Physician Fees, 1993–98," Urban Institute Discussion Paper 99-12, September 1999.
7. Paul Romer, "Preferences, Promises, and the Politics of Entitlement," in *Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-Term Care in America*, ed. Victor R. Fuchs (Chicago: University of Chicago Press, 1995).