13

Redistribution toward Low Incomes in Richer Countries

Emmanuel Saez

During the twentieth century, most developed countries have adopted extensive government-managed income support programs for low-income families and individuals. For example, the United Stated launched income support programs such as Social Security for the old, unemployment compensation for those who lose their jobs, and Aid for Families with Dependent Children (AFDC) for low-income families during the Roosevelt administration in the 1930s.[1] Today, most developed countries devote significant means and effort to redistribution to those with low incomes. However, both the level of generosity and the structure of the programs vary substantially from country to country. Northern European and Commonwealth countries have developed substantially more generous programs than Southern European countries and the United States. Levels of generosity have steadily increased in Europe while there has been a cutback in the United States since the 1980s. The United States targets aid to specific groups, such as the disabled, single-parent families, and the old, and often imposes tight limits on the duration of benefits,[2] whereas many European countries have developed almost universal welfare programs covering most individuals with low incomes,[3] and the duration of welfare payments and unemployment insurance is often much longer.

Redistributive programs toward the poor generate substantial controversy among policy makers and economists. At the center of the controversy is an equity–efficiency trade-off. On the one hand, governments value redistribution and want to transfer resources from the middle- and high-income earners toward low-income individuals. On the other hand, such transfers are generally costly in terms of economic efficiency. First, raising taxes to finance

the income transfer programs may reduce labor supply or savings incentives among the middle- and high-income earners who have to pay the extra taxes. Second, transfer programs may also reduce labor supply incentives of the low-income recipients. As a result, these adverse labor supply effects may substantially raise the cost of improving the living standards of the poor. The equity–efficiency trade-off is reflected in the political debate. Liberals emphasize the redistributive benefits of transfer programs and their important role in raising the welfare of the most needy individuals and families. Conservatives emphasize the efficiency costs, especially at the low-income end, blaming the welfare system for creating dependence and loss of economic self-sufficiency (see, e.g., Murray 1984, for such a point of view on the U.S. experience).

The problem of redistribution is tackled in two steps in economics research. The first step is a positive analysis in which economists develop models of individual behavior to understand how individuals or families respond to various transfer programs along various margins, such as labor supply, education, and human capital investment choices, or family and fertility decisions. The central part of the positive analysis is the empirical estimation of the models of individual behavior in order to assess the quantitative magnitudes of behavioral responses. In the United States and the United Kingdom, there is an enormous literature trying to estimate the size of the behavioral responses to most government transfer programs (see, e.g., Bane and Ellwood 1994; Blundell and MaCurdy 1999; and Krueger and Meyer 2002 for recent surveys). The literature in other countries is smaller but growing quickly. The second step is the normative analysis or optimal policy analysis. Using models developed and estimated in the positive analysis, the normative analysis investigates the structure and size of the transfer program that should be implemented to maximize social welfare. The social welfare criterion used by the government defines the redistributive tastes of the government. Presumably, a liberal government would use a more redistributive criterion than a conservative government. The normative analysis is crucial for policy-making because it shows how programs should be set or reformed in order to best attain the goals of the policy maker. In particular, the normative analysis allows separate assessments of how changes in the redistributive tastes of the government and changes in the size of the behavioral responses to taxes and transfers affect the optimal redistributive program.[4]

The discussion in this essay is organized as follows. The second section briefly describes the positive and empirical analysis (for which numerous good surveys are available) and focuses mostly on the normative analysis. The third section starts with a discussion of the optimal structure of cash transfer programs, depending on the nature and size of labor supply responses. We then extend the analysis of optimal transfer programs along various new dimensions. First, we discuss whether programs should be universal or targeted to specific groups, such as single mothers or the disabled. Second, we review whether in-kind transfers, such as food or shelter, or

workfare programs, in which low-income individuals are required to perform some tasks for the government in order to receive public assistance, may be preferable to cash transfers. Last, we briefly analyze the issue of time limits. Is it preferable to have unlimited programs or to impose tight limits on the duration of benefits? The analysis of the third section is based on the standard assumption in economics that individuals do what is in their best interest. However, there are strong reasons to suspect that, in many cases, individuals may not realize how beneficial a training program could be, or how detrimental for skills a long unemployment spell could be. The fourth section discusses how the results of the third section might be affected in those situations where individuals may not be able to make the best choices for themselves. Finally, the fifth section provides a conclusion and an educated (as well as personal) view on what should be done for redistribution toward low incomes.

It is important to note that while some of the problems described in this essay have been investigated in depth and relatively robust answers have been established, a number of situations reviewed here have received much less attention. In these cases, the discussion presented here is more speculative and should be taken as an encouragement to research rather than as a collection of solidly established results.

OPTIMAL TRANSFER PROGRAMS

Market economies generate substantial levels of income inequality. Because earning and work abilities are very unequally distributed among individuals, without government intervention many individuals would end up with small or even no incomes. The existence of poverty in a developed economy is generally considered a bad market outcome that ought to be corrected to some extent. Surveys carried out in Western countries show that a very large majority favors some level of redistribution (see, for example, Alesina and La Ferrara 2001). Therefore, a government representing its constituents would like to transfer resources from those with high earnings abilities to the disadvantaged who have low earnings abilities or skills.

It is central to note that if earnings abilities were immutable and observable by the government, the government could base transfers directly on earning ability. Such transfers would be independent of investment in human skill or work effort choices, and thus would not create negative incentives.[5] However, earnings abilities cannot be observed directly, and can be inferred only indirectly, through market outcomes such as actual earnings. Thus, the government has to base redistribution on market outcomes, earnings being the most obvious one. Therefore, we first consider the benchmark case where transfers are based on income only, and we discuss below whether using other observable characteristics, such as family or disability status, can be useful to improve redistribution.

Cash Transfer Programs

Most welfare programs are means-tested in the sense that benefits are reduced when earnings increase. For example, TANF, as well as disability benefits and Supplemental Security Income for the old in the United States, or the French *revenu minimum d'insertion* are designed as means-tested programs. Such a transfer schedule relating pretransfer income (horizontal axis) to after-transfer income is depicted in Figure 13.1a. Absent any transfer program, pretransfer and after-transfer incomes would be identical, and thus the budget constraint would be the 45-degree line (the dashed line with slope 1 on the figure). The budget constraint with the transfer program is the solid line on Figure 13.1a: a guaranteed income is given to those with no earnings, and benefits are phased out as earnings increase until the break-even point, at which benefits are lost altogether. The phasing-out effectively reduces the slope of the budget constraint: for each extra dollar of earnings, the after-tax income increases by less than a dollar due to the reduction in benefits. In actual programs, the phasing-out rate is in general high, often equal to one in which case benefits are lost one-for-one as earnings increase. This

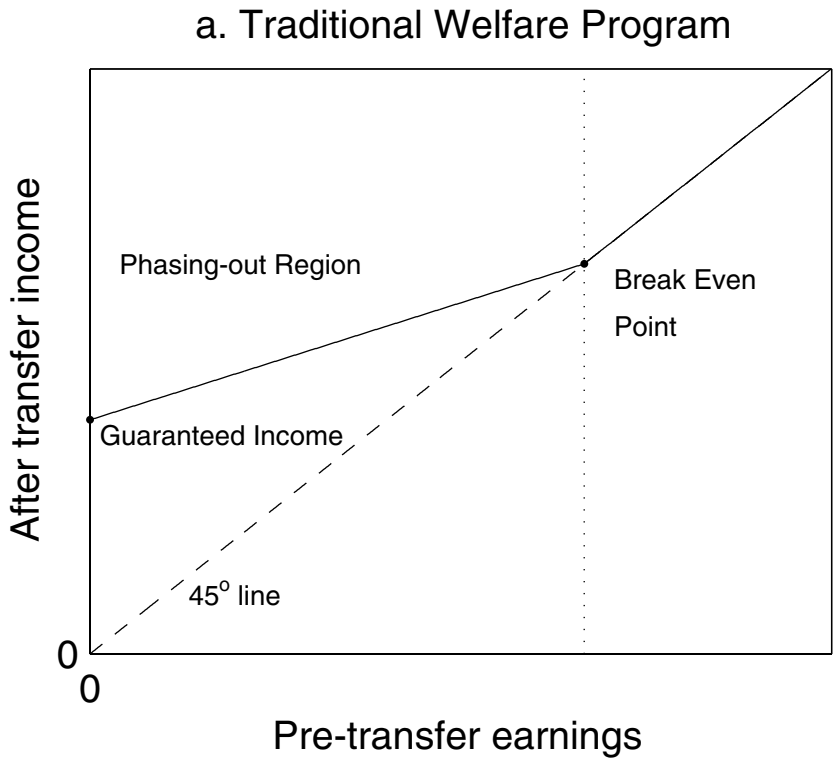

Figure 13.1a Need caption for figure 13.1a

−1
0
+1

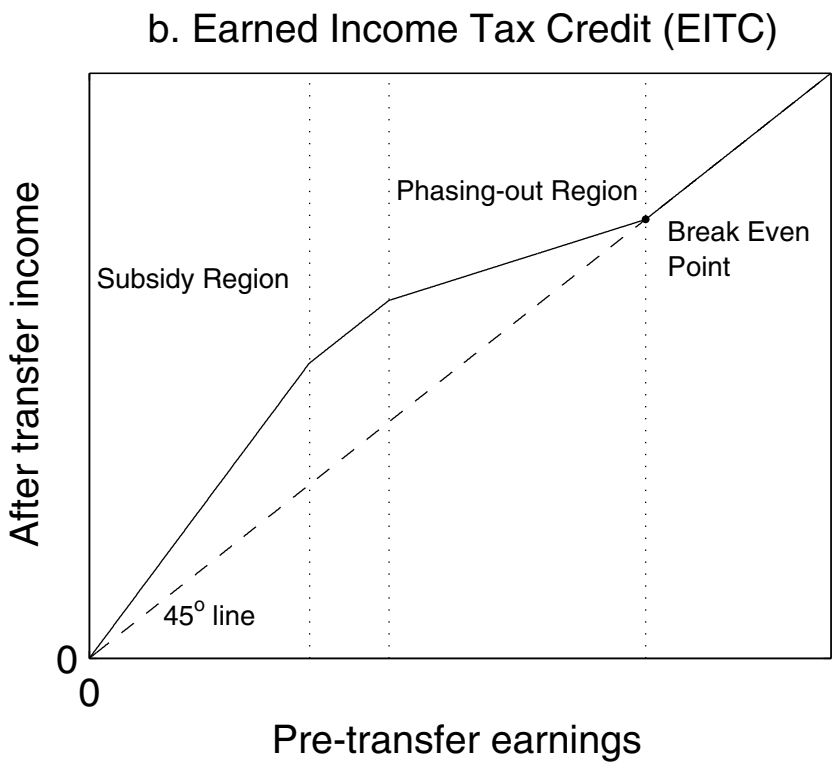

Figure 13.1b Need caption for figure 13.1b

type of program provides the largest benefits to those who have the lowest earnings, and hence are the most in need of income support. However, these redistributive virtues come at a potentially high efficiency cost. The introduction of such a program clearly induces recipients to work less, because the benefits provide extra income (income effect) and because recipients get to keep a much lower share of their earnings (substitution effect due to the phasing-out rate). A number of empirical studies surveyed in Krueger and Meyer (2002) have shown that these programs indeed have negative effects on labor supply, although the reductions in hours of work are typically fairly small for those who are in the labor force.

The normative analysis of the optimal shape of the transfer program (size of the guaranteed benefit and the phasing-out rate) was first investigated by Mirrlees (1971). He considered a simple model in which individuals adjust their labor supply along the intensive margin (i.e., individuals adjust the intensity of work on the job, measured, for example, by the number of hours worked, when taxes or transfers change). In that situation, Mirrlees showed that the optimal transfer program is characterized by a guaranteed benefit for those with no earnings (which depends positively on the strength of

government redistributive tastes) and a positive phasing-out rate of the benefit as earnings increase. Numerical simulations performed in Saez (2001), calibrated using empirical labor supply estimates, show that the guaranteed benefit may be as high as 40% of average earnings even for moderate redistributive tastes and that the phasing-out rate is very high, typically around 70–80% (as depicted, for example, in Figure 13.1a). Such a schedule is desirable because it targets benefits to the most needy individuals in the economy and concentrates the labor supply disincentives to individuals with low earnings ability. These reductions in labor supply incentives at the bottom are not very costly because the beneficiaries would not have had very high earnings even in the absence of the program. Therefore, the Mirrlees model provides a clear answer to an important welfare policy question. It is better to target the program to low-income earners with a high phasing-out rate rather than having a lower phasing-out rate that would reduce incentives to work for a much larger number of individuals.

This suggests that many existing programs with very high phasing-out rates may actually be close to the optimum predicted by the Mirrlees (1971) model. However, such programs have often been held responsible for the low working rates among welfare recipients in the United States (see, e.g., Murray 1984). This has led politicians to advocate programs that would make work sufficiently attractive to reduce the need for income support. In the early 1990s, the Earned Income Tax Credit (EITC) program in the United States was substantially increased, and is now the largest cash transfer program for the poor. The EITC schedule, shown in Figure 13.1b, is fundamentally different from a traditional means-tested program (Figure 13.1a). The EITC does not provide any income support for individuals with no earnings, but all earnings below a given threshold are partially matched by the government, creating a strong incentive to enter the labor force and work. As a result, the slope of the budget constraint in the phasing-in range (depicted in Figure 13.1b) is higher than 1: an extra dollar of earnings translates into more than a dollar in after-transfer income.[6] Empirical studies have shown that the expansion of the EITC in the United States successfully induced low-skilled single mothers to enter the labor force (see, e.g., Meyer and Rosenbaum 2001). The United Kingdom has introduced and expanded a similar program (Working Family Tax Credit). Many other European countries have started implementing such programs on a more modest scale or are contemplating introducing such programs (see Gradus and Julsing 2001).

The key feature missing in the Mirrlees (1971) model is the labor supply decision along the extensive margin, that is, the decision to enter the labor force. Empirical studies (see, e.g., Heckman 1993) have shown that the extensive margin response (choosing whether or not to enter the labor force) is much more elastic than the intensive response margin (choosing how many hours to work once one has decided to enter the labor force). The main reason why this is the case is the fixed costs of working: search costs of finding a job, transportation costs, child care expenses, and so on. Moreover,

most jobs, including part-time jobs, require a minimum and regular number of hours per week. As a result, very few people report working less than twenty hours per week. Saez (2002) shows that a subsidy for low-skilled workers is optimal when labor supply responses are concentrated along the extensive margin. The intuition is as follows: introducing a subsidy for low-skilled workers is good for redistributive purposes and also induces some individuals to enter the labor force, and thus allows the government to save on welfare money. In contrast, in a model with intensive margin responses, a subsidy for low-skilled workers would induce some higher-skill workers to work less in order to take advantage of the subsidy, and would thus increase the cost of the program. That is why such subsidies are not optimal in the Mirrlees (1971) model. Therefore, a government contemplating an increase of incentives for low-skilled workers must precisely weigh the positive effect on work participation and the negative intensive labor supply effect for higher-skilled workers. Saez (2002) presents simulations of this optimal transfer model using empirical estimates of the intensive and extensive elasticities of labor supply. Since the extensive elasticity appears to be much larger than the intensive elasticity, the simulations show that the optimal program should have lower guaranteed benefits (perhaps around 20% of the average earnings in the economy) but that the phasing-out rate should close to zero on the first $6,000 of earnings, so as to make work pay and not deter labor force participation. The benefits should then be phased out at substantial rates for earnings between $6,000 and $15,000. A high phasing-out rate in that earnings range creates only moderate reduction in labor supply because effort on the job (intensive margin) is not very sensitive to incentives.

The transfer programs we have described here are individually based and not family based. However, in practice, transfer programs are often family based.[7] The main reason for using the family is that welfare is better measured by family income than individual income. For example, the nonworking spouse of a high-income husband has no earnings but is not in need. However, the empirical literature has shown that labor supply of secondary earners is much more sensitive to incentives than labor supply of primary earners (see, e.g., Killingsworth and Heckman 1986). Therefore, basing transfers on family income can create perverse incentives for the secondary earner. For example, the EITC in the United States may deter the secondary earner from entering the labor force because the primary earner's income is enough to push family income into the phasing-out range, where the second earner's earnings are implicitly taxed (see Eissa and Hoynes 1998 for such an analysis). Therefore, even though carefully calibrated numerical simulations have not yet been done, it seems that incentives considerations outweigh redistributive considerations, and that transfer programs for low-income persons should be based to a large extent on individual income rather than family income.

Universal versus Targeted Transfers

In the previous subsection, we considered cash transfer programs based solely on earnings. However, as described above, basing transfers on earnings creates negative work incentives. As we discussed at the beginning of the section, it would be more efficient to base redistribution on characteristics that are immutable and related to work ability. For example, disabled people cannot work, and therefore targeting transfers specifically to disabled people should not create negative labor supply incentives for those who can work. Akerlof (1978) made the important theoretical point that tagging welfare programs to observable types such as the disabled can enhance redistributive efficiency. There are two important points to note on this issue.

First, tagging will enhance efficiency the most in situations where the characteristic used for targeting is less easily manipulable. In principle, disability status is not easily manipulable and thus should be an efficient way to target welfare. However, the empirical literature on the U.S. disability insurance system has shown that disability is measured with substantial error, and there is a controversial debate among researchers on whether those on disability insurance are really unable to work (see Bound 1989; Parsons 1991). If disability status is easily manipulable, then a special program targeted to disabled people will create efficiency costs and the gain relative to a universal program will not be great. A characteristic such as age is clearly not manipulable, and therefore, adopting special transfers for the elderly who are in need and who can no longer work, as done in most countries, may be desirable (see Kremer 1997 for a formal analysis).[8]

Second, the characteristic used to target welfare should be closely related to need for support. Therefore, in practice, targeted programs are always means-tested, potentially creating some efficiency costs such as those discussed above. Related to this point, targeting welfare to specific groups such as the disabled or the old may leave large numbers of those in need outside the welfare net, and thus may create unequal treatment of individuals in equal need of support. The U.S. transfer programs target specific groups such as the disabled, the old, and single mothers, and provide almost no support to able adults without children. On the other hand, a number of European transfer systems have a strong universal component.

In sum, optimal redistributive programs should do some targeting, especially using characteristics not easily manipulable and related to earnings abilities. However, it is clearly hopeless to design a program with no efficiency costs because all characteristics related to being in need of support are manipulable to some extent. An optimal program should also provide some support to those with very low incomes but no observable disability. While these theoretical considerations are well understood, it is an interesting and still open research question as to whether the optimal level of targeting should be closer to the U.S. situation or to the European situation.

Cash versus In-Kind Transfers

So far, we have considered only transfers taking the form of cash. While cash programs form the bulk of transfers, most countries have also adopted non-cash transfer programs such as health care provision, education and training programs, public housing, or food and shelter support for the homeless. If we assume that individuals make the best choices for themselves, then receiving cash is better than receiving an equivalent amount in-kind.[9] However, as shown by Nichols and Zeckhauser (1982), even for rational individuals it might be optimal for the government to provide in-kind transfers instead of cash because in-kind transfers might be valued differently by different people. For example, suppose that a shelter program is offered freely. Obviously, shelter is of no value for well-off individuals who want high-quality housing, and thus only individuals in need will take advantage of the program. As a result, and in contrast to a cash transfer that is universally desirable, an in-kind transfer allows the screening of individuals and hence endogenously targets redistribution toward the needy. Therefore, a formal analysis shows that introducing an in-kind transfer program can, in some cases, improve the redistributive power of the government and should be part of an optimal transfer structure. The analysis of workfare by Besley and Coate (1992) is based on the same idea. Workfare provides support but requires individuals to perform some time-consuming task in order to receive the support. Even if the task is completely unproductive, requiring it might be desirable because it allows the screening of recipients, since those in need presumably have a much lower opportunity cost of time.

While the theoretical advantage of in-kind transfers is well understood, it is still an open research question in regard to what extent those types of transfers should be used and how much improvement they would allow the government to make over and above pure cash transfer programs. We believe that this improvement would be small because the efficiency gains of screening come at the welfare cost of providing less desirable transfers. In any case, the theory clearly shows that in-kind transfers cannot completely replace cash transfers but should be used as a complement to cash transfers.

Time Limits for Benefits

Our discussion so far has considered static models with no time dimension. Introducing the time dimension in the optimal transfer problem raises important and interesting questions that have not yet been studied very extensively. In practice, the government can vary the duration of benefits. This issue is especially important in the case of unemployment insurance benefits for those who are temporarily unemployed and where the problem is dynamic by nature. Before the important welfare reform of 1996 in the United States, traditional welfare programs had no time limits. These programs were blamed for creating despondency and promoting a culture of welfare dependence because recipients had no incentives to find work and could rely on

welfare benefits for very long time periods (see, e.g., Murray 1984). Careful empirical analysis has shown that the culture of welfare dependency is much less pervasive than the conservative view suggests, and that most welfare recipients use it for short periods of time (see, e.g., Bane and Ellwood 1994). In many European countries, the long duration of unemployment benefits has been blamed for keeping unemployment rates at high levels.[10] The 1996 U.S. welfare reform imposed a tight five-year limit on the duration of welfare benefits over a lifetime. Indeed, the number of families on welfare in the United States declined from over 5 million in 1994 to about 2.2 million in the early 2000s. Empirical research (see, e.g., the extensive review by Grogger and Karoly 2005) has shown that most of the decline is due to the expansion of the Earned Income Tax Credit and welfare reform, with the improving economy playing a more modest role.[11] Although employment levels of single mothers (those most likely to have benefited from welfare) rose sharply during the period, the main concern was that the loss of welfare might not have been fully compensated for by increased earnings for a number of very low-income families. Meyer and Sullivan (2004) show, however, that the material well-being of single mothers at the bottom of the distribution actually improved slightly during the 1990s.

The duration of benefits can be seen conceptually as another dimension of the generosity of benefits that also creates a classical equity–efficiency trade-off. Limiting the duration of benefits improves incentives to find work and leave welfare or unemployment before benefits end. However, imposing a limit on benefits harms those who cannot find work quickly enough and thus are the most in need of support. The optimal duration of benefits should be set so as to weigh these two considerations. A number of studies have tried to calibrate such theoretical dynamic models, using estimates from the empirical literature in order to assess how long benefits should be set in practice. Most of studies have adopted the standard dynamic model in which individuals can self-insure against future unemployment spells with savings. In that context, it has been found that the size of government-provided unemployment insurance should be rather modest because self-insurance with savings is a powerful tool to insure against short-term income loss due to unemployment (see, e.g., Werning 2002). However, empirical studies such as Gruber (1997) have shown that, in contrast to the prediction of the standard dynamic optimization model, consumption falls sharply during unemployment spells. This suggests that the standard model fails to capture important aspects of the problem, and that if individuals fail to save enough against unemployment risk, government-provided unemployment insurance might be a valuable and desirable program.

THE CASE OF MYOPIC INDIVIDUALS

Our analysis has so far considered only situations where individuals are rational and able to make the best choices for themselves. There are important

reasons for believing that individuals may not always be able to make the best choices, especially when the time dimension is introduced. For example, it is very difficult to make an accurate assessment of the future benefits of investing in human capital now. Therefore, there is a concern that individuals may not realize how beneficial education is, and hence invest too little in education for them or for their children. In that situation, the optimal program should logically provide more incentives for work and invest in education than in the situation where individuals are fully rational. This element would be an additional reason to favor EITC-type programs that promote work over traditional welfare programs that discourage labor force participation.

The notion that individuals may not make the best choices for themselves raises difficult conceptual issues. Individuals may be fully rational and just happen to have a high preference for the present, which causes them to prefer not to invest in human skills today instead of investing and getting more earnings later in life. In that case, a government intervention would be a clear case of paternalism—the government wants to impose its own views over and above the preferences of citizens. That view of the government has been fiercely denounced by libertarian economists from the Chicago school.

However, the growing field of behavioral economics has shown that, in important situations, individual decisions involving the time dimension, such as investment and savings, cannot be accounted for with purely rational preferences.[12] For example, many studies have documented that individuals tend to have inconsistent time preferences that may explain behaviors such as underinvestment in education or procrastination. In those situations, government interventions may be desired by those individuals because it may help them overcome some of the shortcomings of their own behavior. (See, for example, Diamond and Koszegi 2003 for such an analysis in the case of retirement programs for the old when individuals have hyperbolic discount rates). In those situations, individuals do not have standard preferences, and thus it is not clear how the government should evaluate their utility; thus the question of defining a social welfare criterion becomes complicated. Substantial research effort is currently, and will continue to be, devoted to these new research questions in the future. (See Bernheim and Rangel 2004 for a recent exposition of some of the issues of defining a welfare concept in nonstandard situations).

There is a large literature in labor economics evaluating the costs and benefits of training programs for low-skilled and low-income earners (see, e.g., Heckman et al. 1999 for a recent comprehensive survey). Such programs are designed to improve future earnings of the trainees, and thus can be justified only if trainees are myopic (and do not realize the full benefits of being trained) or are credit constrained (and hence not able to borrow to pay for their training). The empirical literature has shown that in general training has a positive impact on future earnings, but is of modest size in most cases. Thus the gains in earnings rarely cover the training costs. How-

ever, the benefits of those programs are very heterogeneous across groups and particular programs. For example, Ifcher (2004) shows that the job placement assistance program set in place by New York City in 1999 for its welfare recipients had a very strong effect on the probability of finding a job and leaving welfare permanently, and thus was cost-effective. Therefore, this suggests that subtle variations in the way those training programs work or small differences in the environment in which they take place can have a dramatic impact on their success. Indeed, the empirical behavioral literature has shown that in a number of contexts, differences in the environment which should be irrelevant for a standard rational economic agent can have very large impacts on actual economic decisions.[13] Therefore, it is plausible to think that such framing effects might also be very important in the context of transfer programs. An important challenge for future research is to understand precisely under which circumstances such programs can be successful in helping low-income earners.

CONCLUSION

This essay has presented a critical overview of the findings of economic research on the problem of redistribution toward low incomes. There are a number of important conclusions that emerge. First, the behavioral labor supply responses to transfers, even though significant, are not so large that the costs of redistribution necessarily outweigh the benefits. Given the size of empirical behavioral responses, substantial transfers could be carried out that would greatly improve the welfare of the poorest families and individuals in American society at a reasonable cost for middle- and higher-income earners. Therefore, the difference in the size of redistribution between the United States and Northern Europe might be due to differences in the redistributive tastes and political processes of those societies.

Second, since empirical studies have found that the labor supply response is concentrated along the participation margin, it would be desirable to structure welfare programs so as to encourage work participation of beneficiaries.[14] A valuable way to do this is to lower the welfare payments for those who do not have any earnings but allow welfare recipients to keep their full earnings in addition of welfare payments up to a limit. The level of welfare payments should of course depend on the number of dependents (in particular children) but should also cover single individuals with no dependents in order not to leave out of support a significant number of persons in need of assistance as in the United States. Furthermore, in the case of two-parent low-income families, it is very important to structure the program so as to preserve incentives to work for both earners.

Third, tagging special groups, such as the disabled or the old, for assistance, as well as providing in-kind benefits instead of cash or extending workfare (i.e., provide benefits conditional on work requirements), could help to improve redistribution on the margin, but it is very doubtful that

such schemes could replace cash transfers to a large extent. Recent studies suggest that many individuals may not take full account of the future benefits of current actions, such as work and investment in human capital. Therefore, this reinforces the idea we developed that redistributive programs should be structured so as to encourage work.

ACKNOWLEDGMENTS I thank Roland Benabou for useful comments. Financial support from NSF grant SES-0134946 and the Sloan Foundation is gratefully acknowledged.

NOTES

1. The AFDC was renamed TANF (Temporary Assistance for Needy Families) following the welfare reform of 1996.

2. For example, welfare payments to single parents are now limited to five years within a lifetime, and unemployment insurance benefits are in general limited to only six months.

3. For example, the guaranteed minimum income (*revenu minimum d'insertion*) in France is a monthly payment for all families or individuals above age twenty-five with no time limits and subject only to a means test.

4. In actual policy debates, these two elements, which are conceptually distinct, are often confused. Conservative policy makers rarely state explicitly that they have little taste for redistribution per se; rather, they justify their lack of taste for redistribution because they believe negative behavioral responses to redistributive programs are large. Conversely, liberals emphasize the redistributive virtues of transfer programs and often ignore negative incentive effects.

5. This general principle is known as the Second Welfare Theorem in economics. It states that any feasible and desirable outcome, no matter how redistributive, can be obtained using appropriate transfers based on immutable characteristics.

6. In 2003, the matching rate of the American EITC was 40% for families with two or more children for the first $10,500 of annual family earnings (corresponding roughly to a single full-time, full-year, minimum-wage salary). The credit is equal to $4,200 for earnings between $10,500 and $14,700. The EITC is then phased out for earnings between $14,700 and $34,700 at a rate of 21%.

7. For example, the EITC in the United States is based on family earnings.

8. This result in favor of old-age support is weakened when one recognizes that poverty in old age might be due to savings decisions taken earlier in life which might be distorted by a generous old-age program.

9. See the final section for a discussion of the case where this rationality assumption does not hold.

10. See, for example, Nickell (1997) for a cross-country empirical analysis. Empirical work in the United States has shown convincingly that the duration of unemployment benefits significantly affects unemployment spells (see Meyer 1990).

11. This finding is confirmed by the fact that welfare rolls hardly increased during the recession of 2001–2002.

12. See, for example, the survey by Frederick et al. (2002).

13. Perhaps the most striking example is the study of Madrian and Shea (2001), showing that a change in the default option for enrollment in an employer-provided pension plan had an enormous impact on enrollment and pension contribution decisions.

14. The United States made such a move with the expansion of the EITC in 1993.

BIBLIOGRAPHY

Akerlof, George. "The Economics of Tagging as Applied to the Optimal Income Tax." *American Economic Review* 81 (1978): 8–19.

Alesina, Alberto, and Eliana La Ferrara. "Preferences for Redistribution in the Land of Opportunities." NBER Working Paper no. 8267, 2001.

Bane, Mary Jo, and David Ellwood. *Welfare Realities.* Cambridge, Mass.: Harvard University Press, 1994.

Bernheim, Douglas, and Antonio Rangel. "Behavioral Public Economics: Welfare Economics When Individuals Can Make Mistakes." Mimeo, Stanford University, 2004.

Besley, Timothy, and Stephen Coate. "Workfare versus Welfare: Incentive Arguments for Work Requirements in Poverty Alleviation Programs." *American Economic Review* 82 (1992): 249–261.

Blundell, Richard, and Thomas MaCurdy. "Labor Supply: A Review of Alternative Approaches." In *Handbook of Labor Economics.* Vol. 3A. Edited by O. Ashenfelter and D. Card. Amsterdam: North-Holland, 1999.

Bound, John. "The Health and Earnings of Rejected Disability Insurance Applicants." *American Economic Review* 79 (1989): 482–503.

Diamond, Peter, and Botond Koszegi. "Quasi-Hyperbolic Discounting and Retirement." *Journal of Public Economics* 87 (2003): 1839–1872.

Eissa, Nada, and Hilary Hoynes. "The Earned Income Tax Credit and the Labor Supply of Married Couples." *Journal of Public Economics* 88 (9–10) (2004): 1931–1958.

Frederick, Shane, George Loewenstein, and Ted O'Donogue. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature* 40 (2) (2002): 351–401.

Gradus, R. H. J. M., and J. M. Julsing. "Comparing Different European Income Tax Policies Making Work Pay." Mimeo, OCFEB, Erasmus University, Rotterdam, 2001.

Grogger, Jeffrey, and Lynn A. Karoly. *Welfare Reform: Effects of a Decade of Change.* Cambridge, Mass.: Harvard University Press, 2005.

Gruber, Jonathan. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87 (1) (1997): 192–205.

Heckman, James. "What Has Been Learned about Labor Supply in the Past Twenty Years?" *American Economic Review* 83 (1993): 116–121.

Heckman, James, Robert J. LaLonde, and Jeffrey Smith. "The Economics and Econometrics of Active Labor Market Programs." In *Handbook of Labor Economics,* Vol. 3A. Edited by O. Ashenfelter and D. Card. Amsterdam: North-Holland, 1999.

Ifcher, John. "Leaving Welfare and Joining the Labor Force: Does Job Training Help? Evidence from an Innovative Intervention in New York City." Ph.D. dissertation, University of California, 2004.

Killingsworth, Mark, and James Heckman. "Female Labor Supply: A Survey." In *Handbook of Labor Economics,* Vol. 1. Edited by O. Ashenfelter and R. Layard. Amsterdam: North-Holland, 1986.

Kremer, Michael. "Should Taxes be Independent of Age?" Mimeo, Harvard University, 1997.

Krueger, Alan, and Bruce Meyer. "Labor Supply Effects of Social Insurance." In *Handbook of Public Economics.* Vol. 4. Edited by A. Auerbach and M. Felstein. Amsterdam: North-Holland, 2002.

Madrian, Brigitte, and Denis Shea. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *Quarterly Journal of Economics* 116 (2001): 1149–1187.

Meyer, Bruce. "Unemployment Insurance and Unemployment Spells." *Econometrica* 58 (1990): 757–782.

Meyer, Bruce, and Daniel Rosenbaum. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." *Quarterly Journal of Economics* 116 (2001): 1063–1114.

Meyer, Bruce, and James X. Sullivan. "The Effects of Welfare and Tax Reform: The Material Well-being of Single Mothers in the 1980s and 1990s." *Journal of Public Economics* 88 (7–8) (2004): 1387–1420.

Mirrlees, James A. "An Exploration in the Theory of Optimal Income Taxation." *Review of Economic Studies* 38 (1971): 175–208.

Murray, Charles. *Losing Ground: American Social Policy, 1950–1980.* New York: Basic Books, 1984.

Nichols, A., and Richard Zeckhauser. "Targeting Transfers through Restrictions on Recipients." *American Economic Review* 72 (1982): 372–377.

Nickell, Stephen. "Unemployment and Labor Market Rigidities: Europe versus North America." *Journal of Economic Perspectives* 11 (3) (1997): 55–74.

Parsons, Donald. "The Health and Earnings of Rejected Disability Insurance Applicants: Comment." *American Economic Review* 81 (1991): 1419–1426.

Saez, Emmanuel. "Using Elasticities to Derive Optimal Income Tax Rates." *Review of Economic Studies* 68 (2001): 205–229.

Saez, Emmanuel. "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses." *Quarterly Journal of Economics* 117 (3) (2002): 1039–1073.

Werning, Ivan. "Optimal Unemployment Insurance with Unobservable Savings." Working Paper, MIT, 2002.