THE COST OF REDUCING ECONOMIC INEQUALITY

Edgar K. Browning and William R. Johnson

The publication of Charles Murray's (1984) book *Losing Ground* has rekindled interest in the role of government in promoting economic equality using redistributive tax and expenditure programs. Murray concludes that spending on poverty programs may well have increased the number of poor people and that eliminating many of these programs might be a better antipoverty strategy. These conclusions reflect the view that the adverse incentive effects of social programs, such as the discouragement of work effort and saving by the poor, more than offset the positive contributions the programs make by providing income transfers. Many scholars, however, have criticized Murray, arguing that much evidence suggests the disincentive effects of social programs are not large enough to offset the income-enhancing effects of the policies (see, for example, Danziger and Gottschalk 1985).

We tend to agree that Murray has overstated his case, but we also believe that the debate has obscured an important aspect of the issue, namely, the importance of distinguishing between the overall effects of social programs and the effects of small changes in the scale of these programs. Much of the recent debate has addressed the question: if we eliminate poverty programs, will poverty rise or fall? An equally, if not more, important question is: if we reduce spending on poverty programs, will poverty rise or fall? These questions may not have the same answers.

The essential distinction is between the overall, or average, effects and the effects of a small change, or marginal effects. As economists have long been aware, marginal and average effects can differ greatly,

Cato Journal, Vol. 6, No. 1 (Spring/Summer 1986).

Edgar K. Browning is Professor of Economics at Texas A&M University, and William R. Johnson is Associate Professor of Economics at the University of Virginia. The authors thank William A. Niskanen for his helpful comments.

not only in magnitude but also in direction. Moreover, for most practical purposes, the marginal effect of doing a little more or a little less is what is relevant. It is unlikely that social programs will be eliminated, but changes in their scales are frequently made, and the marginal effects are relevant in evaluating these changes.

In this paper we investigate the effects of a change in the scale of redistributive tax and transfer programs. More specifically, we focus on the marginal cost to the rest of society of changing the income of the poor. It should be emphasized that finding that this cost is substantial, as we do, does not necessarily imply that the amount of redistribution should be reduced. Marginal cost is only one side of the equation; the other side is the marginal benefits perceived when the poor are helped or a more equal distribution of income is produced. People can be expected to differ over how much they believe greater equality is worth. Nonetheless, the cost of achieving more equality is clearly an important consideration: a person may favor more redistribution when it costs $1.25 to raise the income of the poor by $1 but favor less redistribution if it costs $10 to achieve that result.

A Trade-off between Equality and Efficiency

Popular discussions of egalitarian policies often implicitly assume that income transfers are costless from a social point of view. In other words, one person's income can be increased by $1 by reducing another person's income by $1; total income remains unchanged. This conceptualization of the process would be valid (neglecting administrative and compliance costs) if incomes were unaffected by the policies used to redistribute income. In this setting, the total income pie is fixed and redistribution is a zero-sum game whereby one person's gain exactly equals another's loss. When redistribution is viewed in this way, it is understandable that ethical and moral judgments become the primary considerations.

Economists have long been aware, however, that a redistribution of income, by severing the link between effort and reward, can dull the incentive to earn income. Thus, they speak of a trade-off between equality and efficiency; that is, greater equality may come at the expense of a lower total income for the community. Put in a different but more revealing way, if an expansion in redistributive policies leads some people to decide to earn less income, then the cost to the nonpoor of increasing the income of the poor by a dollar will be greater than one dollar.

A key empirical question, then, is how much tax and transfer programs affect the supply of productive resources and hence the size of the total income pie. Twenty years ago, available evidence suggested that the effect of public policies on labor supply and saving was negligible, small if not actually zero. In recent years the use of more sophisticated statistical techniques and better data has increasingly led economists to conclude that adverse supply responses are significant. For example, studies of the data generated by field tests of the negative income tax form of welfare program in the 1970s have estimated that work effort of recipients would fall by perhaps 10 to 20 percent. This finding is relevant to a number of actual welfare programs that are similar in structure to the negative income tax. At a more aggregate level, a major survey of studies of all types of government transfer programs hazarded the conclusion that all these programs together had reduced national labor supply by about 4.8 percent (Danziger, Haveman, and Plotnick 1981).

Suppose we tentatively accept the conclusion that transfer programs have reduced labor supply by roughly 5 percent; does this imply that the efficiency cost of policies designed to reduce income inequality is small? Many people would probably view a 5 percent reduction as a small price to pay for the welfare and income support system we have. In fact, most economists seem to concur. As Danziger (1982, p. 18) states: "[I]t should be stressed that while income transfer programs create disincentives to work and save, the magnitude of these disincentives is relatively small. They pose no threat to the overall efficiency of the economy."

The major purpose of this paper is to suggest that this sanguine conclusion misinterprets the significance of disincentive effects for evaluating the cost of reducing economic inequality. We do not dispute the magnitudes of the disincentive effects most researchers have found; in fact, our analysis is based squarely on their findings. Instead, we argue that disincentive effects of the size generally found imply that the cost to the rest of society of raising the income of the poor is quite high.

The fact that the overall effect of taxes and transfers on labor supply and saving is small, if true, does not tell us much about the relevant cost of reducing economic inequality. There are two reasons for this. First, as pointed out earlier, the cost relevant for decision making is the marginal cost. A 5 percent reduction in labor supply due to current transfer programs is a measure of total cost or, if expressed relative to some base, of average cost. Average cost, however, is generally not a good guide to marginal cost; in the present case, marginal cost is certain to be greater than average cost.

An example will clarify this point. Imagine increasing the tax rate on labor earnings from zero to 100 percent. At 100 percent, a worker gets to keep none of his earnings, so labor supply will typically fall to zero at or before this point. This does not, however, imply that labor supply will fall by 40 percent when the tax rate is 40 percent (as it is for many households today). Common sense as well as empirical evidence tells us that the reduction is much smaller. An 8 percent reduction, for example, represents an average reduction of 2 percent for each 10 percentage points in the tax rate. This, however, does not mean that increasing the tax rate by another 10 percentage points will reduce labor supply by only 2 percent: Since labor supply will fall to zero when the rate reaches 100 percent, the reduction in labor supply beyond some point must become much larger than 2 percent. It seems reasonable to suppose that each successive 10-percentage-point increase in the tax rate will depress labor supply by more than the previous increment. This would then describe a situation where the marginal reduction in labor supply is greater than the average reduction. Consequently, focusing on the total or average effects of tax and transfer programs will normally understate the responses produced by a small change in these policies.

The second reason why a small total reduction in labor supply does not necessarily imply that the marginal cost of reducing inequality is low is that the disincentive effect (the reduced labor supply) is not properly related to the income gains and losses of transfer recipients and taxpayers. The aggregate reduction in labor supply does not tell us what the marginal cost of increasing the incomes of low income households will be. Clearly, if we know that labor supply will fall by X percent following a certain policy change, that is related in some way to the marginal cost, but the exact relationship is not obvious. It turns out, as we show below, that a seemingly moderate change in labor supply can imply a surprisingly high marginal cost.

Marginal Tax Rates, Productive Incentives, and Redistribution

It is important to understand how tax and transfer programs tend to undermine incentives to earn income. One important way such programs affect productive incentives is through the application of high marginal tax rates to earnings. Seeing why incentives are affected requires distinguishing carefully between marginal and average tax rates.

Consider a worker with $30,000 in earnings, and suppose the government wishes to collect $7500 in taxes from him. As a first approx-

imation, it could do this by exempting $20,000 (so taxable earnings are $10,000) and applying a constant tax rate of 75 percent to taxable earnings. If the worker continues to earn $30,000, his average tax rate is 25 percent ($7500/$30,000), but his marginal tax rate is 75 percent, since any change in earnings from the current level is subject to the 75 percent rate. Note that it is the marginal tax rate—not the average amount of earnings taken in taxes—that gives the worker less incentive to work. It does this by reducing the net wage rate received by the worker for changes in the amount he works. In this case, should the worker earn $1000 less, his disposable income falls by only $250. In effect, a 75 percent marginal tax rate reduces the net wage rate at the margin to one-fourth of its before-tax level. It pays less to work, and the higher the marginal tax rate the lower is the net rate of pay.

In addition to taxes, government transfers to low income households also often impose marginal tax rates on the earnings of recipients, by relating benefits to the level of earnings. If the transfer a person receives is reduced when earnings are increased, then this reduction in benefits has exactly the same effect as a marginal tax rate applied to earnings. Suppose, as in the food stamp program, that the transfer is reduced by $0.30 for each dollar earned. When earnings rise by $10, disposable income only goes up by $7, the same as if a 30 percent marginal tax rate is applied to earnings.

Therefore, to evaluate how taxes and transfers affect incentives, it is important to consider how high marginal tax rates are under present programs. This is difficult to establish because the relevant marginal tax rate is the combined rate that results from all the separate taxes paid and/or transfers received, so it will differ by income level and even among households at the same income level. In recent research, however, we have developed estimates of the average effective marginal tax rates for different income classes of households using 1976 data (see Browning and Johnson 1984). Although there have been several changes in tax and transfer policies since 1976, we do not believe current marginal rates would be significantly different on balance. For the lowest income class, containing the 20 percent of households with the lowest incomes, the average marginal tax rate was 54.4 percent. Most of this rate is attributable to the benefit reduction rates for transfers received, and some households receive transfers from more than one program. For the next four income classes from lowest to highest, each containing 20 percent of households, our estimated marginal tax rates are: 47.1 percent, 40.8 percent, 38.8 percent, and 44.6 percent.

Marginal tax rates at these levels can be expected to adversely affect incentives to work and save, but, as mentioned earlier, it is

possible that the total effect would be relatively small. The more relevant question, however, is how a small change in the amount of redistribution will affect incentives to earn for both taxpayers and transfer recipients so as to determine the marginal cost of reducing economic inequality. A small increase in transfers and taxes from their present level will obviously increase marginal tax rates and thereby aggrevate disincentive effects. But what is probably not obvious is how much marginal tax rates will rise in comparison with the share of national income that is transferred from upper to lower income groups.

Suppose taxes and transfers were increased in such a way that the increment in marginal tax rates is the same for all persons. Under these conditions, to redistribute 1 percent of national income from upper to lower income households increases everyone's marginal tax rate by approximately 5 percentage points. To see why this happens, assume that income from all households with above average income is transferred to all households with below average income, and that average income is $30,000. To collect 1 percent of national income in taxes from households earning over $30,000, we must take more than 1 percent of their combined incomes, since the combined incomes are less than national income. But there is a more important reason why the marginal tax rate on upper income households will be greater than 1 percent: if only above average incomes are taxed, then only that portion of income in excess of the average can be taxed.[1] Thus, taxable income for a household earning $40,000 would be $10,000, or the earnings in excess of $30,000, and to collect just 1 percent of this household's total income requires a marginal tax rate of 4 percent on its taxable income.

For a similar reason, to redistribute 1 percent of national income to households with below average incomes means that the transfer program must embody a marginal tax rate of about 5 percent. To restrict transfers to those with low incomes, the transfer must decline as income rises until until the transfer is zero at an income of $30,000.[2]

[1]It is often suggested that the total incomes of high income households could be taxed, while not taxing low income households at all. This procedure, however, encounters a serious problem. Suppose that those earning above $30,000 were required to pay a tax of 10 percent of their total incomes while those earning less pay no tax. A person earning $33,000 would actually have a higher disposable income by reducing his earnings to just below $30,000 and avoiding the tax. In other words, such a policy creates a strong incentive for those just above $30,000 to reduce their incomes to below that level.

[2]If all households with incomes below $30,000 were given the same transfer, say $1000, this would mean a person earning $30,500 could increase his disposable income by reducing his income below $30,000 to quality for the transfer.

The required benefit reduction rate, or marginal tax rate, would be 5 percent. This necessary benefit reduction rate is the main reason why low income households now face such high marginal tax rates. Even though the lowest two quintiles receive a net redistribution that is probably no more than 5 or 6 percent of national income, their marginal tax rates are typically close to 50 percent, over half of which is due to the benefit reduction rates embodied in transfer programs.

Of course, the exact way an increased redistribution raises marginal tax rates depends on the particular changes made in tax and transfer policies, and it is not necessary for everyone's tax rates to rise by the same amount. But the major point—that marginal tax rates rise more than the percentage of national income redistributed—remains valid. We might, for example, choose to redistribute income to those with incomes below $15,000 and tax those with higher incomes. In this case, a redistribution of 1 percent of national income might increase marginal tax rates of recipients by perhaps 10 percentage points, higher than the previous case because the same volume of transfers must be tapered off more rapidly as income rises (a higher benefit reduction rate) so that the transfer falls to zero at $15,000. For taxpayers, the increment in marginal tax rates might be only 3 percent because now income in excess of $15,000 can be taxed. In general, when 1 percent of national income is redistributed, if the increment in marginal tax rates for some people is kept below 5 percentage points, this will necessitate increases for other people greater than 5 percentage points.

A Hypothetical Example

In this section we use a simple numerical example to illustrate how the various factors discussed above interact to determine the marginal cost of increasing the income of low income households. In Table 1, we assume that society is composed of five households with labor earnings as shown in column 2. Initially, each household is assumed to confront a marginal tax rate of 40 percent under current tax and transfer policies. (Exactly how much each household initially pays in taxes or receives in transfers has no bearing on our problem except insofar as they combine to produce the 40 percent marginal tax rate.) What we propose is to increase slightly the marginal tax rate in a way that redistributes income in favor of low income households. Specifically, we assume that the marginal tax rate on total earnings is increased by 1 percentage point, and the additional tax revenue produced is returned as equal per household transfers. On

TABLE 1

EFFECTS OF A CHANGE IN REDISTRIBUTIVE POLICY

Household (1)	Initial Earnings (2)	Change in Earnings (3)	Net Additional Tax Revenue[a] (4)	Transfer (5)	Net Transfer (6)	Change in Disposable Income (7)
A	10,000	−50	80	240	160	110
B	20,000	−100	160	240	80	−20
C	30,000	−150	240	240	0	−150
D	40,000	−200	320	240	−80	−280
E	50,000	−250	400	240	−160	−410
Total	150,000	−750	1,200	1,200	0	−750

[a]These figures are rounded to the nearest $10.

balance, this will redistribute income in favor of households with below average incomes.

We also need to specify how the households' labor supply, and hence earnings, will respond to the higher marginal tax rate. We therefore assume the elasticity of labor supply with respect to the net wage rate is 0.3 for all workers. The labor supply elasticity simply indicates the percentage change in labor supply divided by the percentage change in the net wage rate. An elasticity of 0.3 thus means that a 10 percent reduction in the net wage rate will result in a 3 percent reduction in labor supply.

Any labor supply elasticity should be interpreted as the average response of a large number of people, some of whom may respond either more or less than the amount predicted by that elasticity. Furthermore, workers can respond to changed incentives in several ways: changing hours worked per week, vacation days, time of entry into the labor force, time of retirement, moonlighting, or intensity of effort on the job. Obviously, many of these changes will be apparent only in the long run as workers change jobs, new workers enter the labor force, or unions alter their bargaining demands.

The actual effects of our incremental redistributive policy depend significantly on the size of labor supply responses, as indicated here by the elasticity. Our choice of 0.3 reflects the available empirical evidence in the economics literature. While there is certainly no consensus that the exact value is 0.3 (interpreted as an average of the different elasticities of different groups of workers), most economists would view it as a reasonable figure in light of available evidence.[3] In any case, the implied reduction in labor supply is not implausibly large. Indeed, our intention is to show how a moderate reduction in labor supply translates into a surprisingly high marginal cost of raising the incomes of low income households.

When the marginal tax rate rises from 40 to 41 percent, the net wage rate falls from 60 to 59 percent of its before-tax value, a reduction of 1.67 percent. The percentage reduction in labor supply is thus 1.67 percent times the elasticity of 0.3, implying that labor supply will decline by 0.5 percent. Assuming that earnings decline in proportion to labor supply, the reductions in gross earnings are shown in column 3 of Table 1.

[3]For a recent survey of empirical studies, see Hansson and Stuart (forthcoming). For simplicity, the numerical example developed here assumes that uncompensated and compensated labor supply elasticities are equal: labor supply is taken to depend only on the marginal net wage rate, effectively ignoring income effects. The empirical results presented in the next section, however, are based on an analysis that accounts for income effects.

The additional tax revenue collected from each household is shown in column 4. Tax revenues do not rise by 1 percent of total earnings when the tax rate is increased by 1 percentage point. When earnings fall by $750, the government loses $300 in revenue under its initial 40 percent marginal tax rate, so the 1 percent increase in the marginal tax rate only adds $1200 to revenue rather than $1500. This illustrates one reason why the initial level of marginal tax rates is relevant: the higher the rates, the greater the revenue loss from any reduction in earnings caused by a change in tax or transfer policy.

The additional tax revenue is returned to households as transfers of $240, shown in column 5. The figures in column 6 are the difference between the transfer received and the additional tax paid. On balance, our policy transfers $240 from households D and E to households A and B. (The amount redistributed, $240, is about one-sixth of 1 percent of total income, but marginal tax rates had to rise by a full percentage point to accomplish that redistribution.) This transfer, however, does not indicate how the disposable incomes of households are affected, since each household's own earnings have fallen in response to the policy. The change in disposable income is given by the sum of the net transfer and the change in earnings, with the results shown in column 7. Overall, disposable income falls by $750, which equals the reduction in total earnings caused by the labor supply reduction.

A small expansion of taxes and transfers therefore raises the disposable income of household A by $110 and reduces the disposable incomes of the other households by $860. In other words, the marginal cost of raising the income of the poorest household by $1 is about $8. We believe that this is the appropriate way to look at the effects of a change in redistributive policy since the gain—$1 to the poor—is weighed against its cost—$8 to the nonpoor. As noted above, the magnitude of the marginal cost does not tell us whether it is desirable to expand or curtail redistribution, but it does present the consequences in a way that makes a more informed judgment possible. This hypothetical example illustrates the point that a moderate disincentive effect can imply a surprisingly high marginal cost associated with reducing income inequality. In the next section we show that this conclusion remains valid when actual data on U.S. households are used rather than hypothetical figures.

Before proceeding, however, it is necessary to consider further the specific type of policy change that is used in Table 1 and in our simulations discussed below. This analysis has been criticized as biased toward finding a high marginal cost because we assume that

transfers are made not only to the poor but also to the nonpoor (see column 5). In response, we would make two points.

First, the assumption that everyone pays taxes and receives transfers is used to simplify the calculations in our example; the policy need not work in exactly this way. It is important to recognize that our tax-transfer policy is equivalent to a negative income tax that restricts transfers to those with below-average incomes and finances these transfers with taxes that fall exclusively on those with above-average incomes. For example, the government could simply make the transactions recorded in column 6. In this case, there would be a benefit reduction rate (marginal tax rate) for transfer recipients of 1 percent up to an income of $30,000, and a marginal tax rate of 1 percent on earnings above that level; each person's marginal tax rate would rise by 1 percentage point as before, so labor supply would be affected in exactly the same way as in our example. This avoids having everyone pay taxes and receive transfers, but the results are the same as in our example.

In our example the nonpoor receive transfers, but that feature is not essential to the redistributive policy we are analyzing; it is merely an accounting procedure that simplifies the exposition and arithmetic. In practice, we would not expect most actual changes in redistributive policies to involve transfers to everyone. Nevertheless, our hypothetical redistributive policy is still capable of evaluating such policies if their net effects are similar to those shown in column 6.

Second, we have studied the actual pattern of taxes and transfers in the United States and compared the net transfers at each income level to those implied by our hypothetical policy. We found that the average effect of all taxes and transfers on the distribution of income is similar to the distributional effects produced by our hypothetical policy. In other words, a proportionate expansion in all taxes and transfers would produce distributional results much like our policy. This result is perhaps the most important reason for believing that our hypothetical policy can serve as a reasonable proxy for an actual expansion or contraction in the scale of redistributive policies in the United States.

Thus, we think examining the effects of changing the distribution of income using our hypothetical policy provides a good indication of the actual cost of redistributing income. Of course, a different type of policy might perform better, and in the Appendix we explore such a possibility by examining, in the context of the same five-person society as in our example, a policy that restricts transfers to those earning below $20,000 and collects taxes from those with incomes above that level. We found that the marginal cost is about 10 percent

less than in the example discussed here. This, along with other policies we have examined, suggests that the hypothetical policy emphasized here is not likely to misrepresent the effects of actual redistributive policies to any significant degree.

Marginal Cost and Labor Supply Responses

Using data from a representative sample of about 50,000 U.S. households in 1976, we employed computer simulation techniques to calculate the marginal cost of reducing economic inequality. A 1 percentage point increase in marginal tax rates on labor earnings was assumed to be added to the actual system of taxes and transfers, and the revenues returned as equal per capita transfers, as in our previous numerical example. The ultimate effects on disposable incomes were worked out under a variety of assumptions regarding labor supply responses.

Some of the results are presented in Tables 2 and 3. Households were ranked according to income and grouped into quintiles, each containing 20 percent of all households. Thus, quintile one contains the poorest 20 percent of households and quintile five contains the wealthiest. Table 2 gives the results for four different assumptions about the size of labor supply responses. We regard the benchmark case, with its elasticity of 0.31, as best supported by available evidence. Labor supply elasticities, however, are difficult to estimate with any degree of accuracy, as illustrated by the range of estimates found in the literature. Thus, we also performed the simulations for higher and lower assumed values for the labor supply elasticity.[4] Finally, for purposes of comparison, we calculated the effects when there was no change in labor supply.

The results in our benchmark case are similar to the hypothetical example discussed earlier. A small expansion in redistribution that raises everyone's marginal tax rate by 1 percentage point leads to an increase in the average disposable incomes of households in the two lowest quintiles (but primarily in the lowest quintile) and to a reduction for the upper three quintiles. Per dollar of additional income for the lowest two quintiles combined, the marginal cost to the upper

[4]As a technical note regarding our modeling of labor supply responses, two points should be made. First, the elasticities are economy-wide average figures; in our formulation, each household's elasticity depends on its marginal and average tax rates, and so the elasticities vary among households. Second, the figures in Tables 2 and 3 are the compensated labor supply elasticities. The corresponding uncompensated elasticities are 0.2, 0.31, and 0.05 for the benchmark, high elasticity, and low elasticity cases, respectively. For more detail about the technical aspects of our work, see Browning and Johnson (1984).

TABLE 2

CHANGES IN DISPOSABLE INCOME PER HOUSEHOLD

Case	Elasticity	Quintile 1	2	3	4	5	Marginal Cost
Benchmark	0.31	36.49	7.54	−43.15	−108.17	−267.49	9.51
High Elasticity	0.47	25.23	−15.13	−76.46	−150.65	−324.95	22.48
Low Elasticity	0.21	45.73	25.60	−16.41	−71.63	−211.82	4.20
No Response	0	64.48	62.69	32.54	−22.82	−136.89	1.0

SOURCE: Browning and Johnson (1984, Table 6).

three quintiles is estimated at $9.51. In our benchmark case, we thus find that it costs the nonpoor about ten dollars to raise the income of the poor by one dollar when labor supply responses are taken into account.

Predictably, the results depend strongly on how large the labor supply response is. With an elasticity of 0.47 (implying that a 10 percent reduction in net wage rates results in a 4.7 percent decline in labor supply), marginal cost is more than double the benchmark case. For an elasticity of 0.21, it is less than half the benchmark case. Since we do not know with any precision how large labor supply responses will be, we also cannot determine exactly how large the marginal cost really is. In our view, however, Table 2 suggests that marginal cost can be surprisingly high even when labor supply responses are quite modest.

The results of the redistribution when everyone continues earning the same income are presented in the last row of Table 2. In this case, marginal cost is, of course, equal to $1. Of greater interest is how the gains and losses for quintiles vary from the zero elasticity case to the others. The gain to the two lowest quintiles together is twice as large in the zero elasticity case than in the low elasticity case, and three times as large as in the benchmark case. Therefore, studies that simply calculate gains and losses without taking labor supply effects into account greatly overstate the extent to which an increase in redistribution raises the incomes of the poor.

These figures all pertain to changes in disposable money income and would not accurately reflect how the well-being of households is affected. This is because people are working less and consequently have more leisure time. For example, in the benchmark case, the average household in the lowest quintile gains $36.49 in money income and also has more time available for nonwork pursuits, so the gain in money income alone understates the total benefit of the policy. Table 3 provides estimates of the changes in real income per household; these figures add to the change in money income an estimate of the value of the additional leisure time gained when labor supply falls, using the net wage rate to value time. The marginal cost of increasing the real income of the poor is lower when this adjustment is made.

Of the two measures of marginal cost—one based on money incomes and the other incorporating the value of leisure—which is the relevant one to consider? Most economists would emphasize marginal costs reckoned to include the value of time, because they believe all gains and losses—not just changes in spendable income—should be counted. For that reason, the estimates in Table 3 give a better

TABLE 3

CHANGES IN REAL INCOME PER HOUSEHOLD

Case	Elasticity	Quintile					Marginal Cost
		1	2	3	4	5	
Benchmark	0.31	47.30	32.60	−11.49	−71.59	−196.21	3.49
High Elasticity	0.47	38.55	18.07	−30.32	−93.45	−222.00	6.11
Low Elasticity	0.21	56.12	45.85	3.46	−56.80	−185.071	2.29

SOURCE: Browning and Johnson (1984, Table 8).

indication of the size of the trade-off between equality and efficiency, since our concept of efficiency recognizes time as a scarce resource.

But insofar as our goal in redistributing income is to increase the consumption of goods and services by the poor, the estimates in Table 2 are more relevant. Poverty is usually identified with a lack of such material things as food and housing, and it is the level of money income that largely determines consumption. Since the government's poverty lines are based on money income and take no account of how much or whether people work, a concern over the level of money incomes alone is indicated. Therefore, we conclude that both measures of marginal cost provide useful information.

Marginal Cost Including Savings Responses

The estimates discussed above incorporate only the effects of reduced labor supply, but there are other ways that tax and transfer policies may contribute to the high cost of helping the poor. For example, there are collection, administrative, and enforcement costs associated with raising tax revenues and dispensing transfers. These costs also enter as a wedge between the benefit to recipients and the burden on taxpayers, and incorporating these costs would further increase the marginal cost (but not by a great deal, we suspect). The focus in this section is on yet another potentially important cost: the cost of a change in redistribution on the level of private saving.

An increase in the amount of redistribution is likely to reduce saving for two reasons. The first is well known and is based on the belief, which has some empirical support, that the well-off save more (at the margin) than do the poor. Taking a dollar from a wealthy person might reduce his saving by 20 cents, while the recipient of that dollar might spend it all, implying that total saving falls by 20 cents. The second and more important reason is based on our earlier analysis. Since redistribution lowers the total income of the community (because of the labor supply reduction), saving can be expected to fall even if the marginal propensity to save out of income is the same for the poor and the nonpoor. These two points together strongly suggest that increased redistribution will depress the level of saving.

But why is a reduction in saving a cause for concern? Under certain circumstances, it would not be. Specifically, if the tax on labor income is the only distortion in an otherwise competitive economy, whether people respond to a reduction in income by cutting saving or consumption is irrelevant to the determination of the cost of helping the poor. In this case, a person who saves receives the entire gain associated with providing funds that finance capital formation. If the

interest rate is 5 percent, then an additional dollar in saving can finance capital that (net of depreciation) will augment future production at an annual rate of return of 5 percent. A reduction in saving means that the saver sacrifices future consumption, but the sacrifice is equal to the future production the saving would have made possible. Thus, if a $1 reduction in income led to a $1 reduction in saving, only the saver bears a cost, and that cost will be correctly measured as $1.

The situation is different, however, if the government intervenes in the operation of capital markets by levying taxes on capital income. In the United States, several taxes fall on the return to capital, notably corporate income taxes, property taxes, and personal income taxes. The combined effect of these taxes is to make the rate of return received by the saver less than the rate of return actually generated when that saving finances capital investments. For example, a recent study (Feldstein, Poterba, and Dicks-Mireaux 1981) has estimated that the before-tax real return to capital in the corporate sector averaged about 10 percent in the 1970s, but after taxes savers received a real rate of return that averaged about 3 percent. In other words, the return to corporate capital was taxed at a rate of about 70 percent. Although tax rates on capital income are somewhat lower today, we will use these figures to illustrate how a change in saving is relevant to the marginal cost of redistribution.

With a before-tax return of 10 percent and an after-tax return of 3 percent, consider what happens when a person saves less, and assume that the principal would have been held intact in the future and only the net return consumed. If saving declines by $1, the person sacrifices a stream of interest payments of $0.03; that sacrifice is borne by the saver himself. But there is also a sacrifice of a stream of future tax payments of $0.07 per year, and that cost is borne by people other than the saver, presumably by whomever would have benefited from the expenditure of those tax revenues. Consequently, if redistribution causes saving to decline, a tax on the return to saving imposes costs in the future on people other than those whose saving falls. This cost takes the form of reduced capital income tax revenue, and it is not taken into account in our earlier estimates.

Exactly how this cost should be measured is a complex theoretical issue, but we can explain why this cost is likely to be large. The net (after-tax) rate of return of 3 percent can be taken as a measure of how much future benefits are discounted: a benefit of $1.03 one year in the future has a present value of $1.00. When saving falls by $1, as in the example above, society loses a stream of annual benefits of $0.10 ($0.07 of which would have gone to the government and $0.03

to the saver). The present value of an infinite stream of annual benefits of $0.10 when discounted at the 3 percent rate is $3.33 ($0.10/.03). Thus, when saving falls by $1, the present value of the future costs is $3.33; one dollar of this is the cost borne by the saver himself and the other $2.33 represents the present value of lost future tax revenues. In this example, a reduction of $1 in saving actually imposes costs on people other than the saver of $2.33.

The above calculation is a simplified example of what economists call the shadow price of capital. It is intended to measure the present value of future sacrifices produced when saving (and hence real investment) falls by $1. The shadow price, therefore, is the real social cost of a reduction in saving. The actual calculation would take into account how long the principal is held intact and the extent to which some of the net return to the saver may additionally be saved in the future. In a recent survey of issues related to the calculation of the shadow price of capital, Lind (1982) concludes that 3.80 is a reasonable estimate. This estimate is based on the assumption that the tax rate on capital income is 54 percent. We use Lind's estimate of the shadow price of capital in the following discussion.

Now consider how the cost of reduced saving, as measured by the shadow price of capital, influences the marginal cost of helping the poor, using the benchmark case from Table 2. The marginal cost of raising the incomes of the bottom two quintiles by $1 ($0.83 for the lowest and $0.17 for the second) was a loss of $9.51 to the top three quintiles. Suppose that only the top three quintiles save and that their marginal propensity to save out of disposable income is 20 percent. Then a decline of $9.51 in income implies a reduction of $1.90 in saving. The cost of this reduction in saving is $3.80 per dollar of saving, or $7.22. Of this, we have already counted $1.90 as the loss to the savers, but that leaves an additional cost of $5.32 to be allocated.

It is not clear who will actually bear this cost, but since it takes the form of reduced tax revenues, presumably the losses result from lower government spending in the future. This raises the further question of how the benefits of government spending are distributed among income classes. The distribution for total government spending is unknown, but we do have estimates of how government expenditures on transfers are distributed among income classes. The lowest quintile receives about 21.7 percent of all transfers.[5] If we assume

[5]This figure may seem low, but it should be pointed out that the lowest quintile of households contains only 11.5 percent of all persons, so 21.7 percent represents a per person transfer of nearly twice the average. Moreover, such programs as social security, medicare, and unemployment insurance, which involve large transfers to the nonpoor, also work to keep the percentage low.

that the lowest quintile will bear 21.7 percent of the unallocated cost of reduced saving, the cost on the lowest quintile is $1.15. When this is combined with their gain in current income of $0.87, the overall effect is a reduction of $0.28. The reduction for the second quintile would be larger.

Note carefully what this implies: an expansion in redistributive policies of the sort described reduces the incomes of not only the nonpoor but also the poor, when the present value of future losses is included. Conversely, a reduction in redistribution would actually benefit the poor under these conditions.

We do contend that this is necessarily the outcome. Obviously, the actual effects depend on labor supply responses, savings responses, tax rates on labor and capital, and other factors, many of which are not known with accuracy. What we have shown, however, is that plausible assumptions about the relevant magnitudes imply that redistribution at the margin may actually harm the poor. Even if this extreme result is not correct, it should be clear that, when the effects of reduced saving are incorporated, the marginal cost of raising the incomes of the poor is likely to be much higher than the estimates of the last section that evaluated only the labor supply responses.

The Time Pattern of Gains and Losses

Our finding that redistribution at the margin may harm both the nonpoor and the poor may seem reminiscent of conclusions drawn from supply-side economics using the Laffer curve. In particular, some supply-siders have argued that a reduction in tax rates would increase tax revenues (and therefore benefit everyone) because the lower rates would so stimulate the supply of productive services. Our analysis is related to this view to the extent that we also stress the importance of effects on the supplies of labor and capital. But there is a difference: for the extreme supply-side outcome to occur, resource supplies must increase sharply when tax rates are reduced, but our analysis assumes much smaller responses. For example, for a reduced tax rate on labor income to yield more revenue, as some supply-siders have suggested it would, it has been calculated that the labor supply elasticity would have to be in the 1.5 to 2.5 range, well beyond the value most economists find plausible. In our benchmark case, we assume the labor supply elasticity is only 0.31. In fact, the most important implication of our analysis is the finding that seemingly moderate responses imply that marginal cost is quite high.

Exactly how the poor may be harmed by an expansion of redistributive taxes and transfers can be better understood by considering

how the effects would unfold over time. In the discussion above, we discounted the future costs and expressed them in present value terms. We also assumed that labor supply adjusted immediately when marginal tax rates rose. The timing of the responses, however, suggests that the short-run, or temporary, effects will be quite different from the long-run, or ultimate, effects.

Suppose, for example, there is a small expansion in redistributive programs. Since it takes time for people to adjust their work habits to a change in circumstances, the immediate effect would be a small change in labor supply. Since earnings fall very little, savings would also decline very little at first. The short-run effect, therefore, will be to raise the incomes of the poor, and the marginal cost to the nonpoor would initially be only slightly greater than one dollar, ignoring subsequent future effects. But as people adjust their labor supply, the benefits to the poor fall and the costs to the nonpoor rise. In addition, as saving is lower year after year, the cumulative effect on the stock of capital grows and tax revenues from capital income taxes fall progressively further below what they would otherwise have been. Government spending on transfers will probably grow more slowly as a consequence, further harming the poor and the nonpoor in the future.

By distinguishing between short-run and long-run effects, we can see why it is virtually certain that an expansion in redistribution will at least temporarily raise the incomes of the poor. But the benefits gradually dissipate over time and the costs rise, until in some future year the poor may have incomes that are no higher than if the initial expansion had never taken place, and in later years have even lower incomes. It is also important to recognize that the adjustment process can take quite a long time. Studies have found that it may take from 2 to 4 years for half the eventual change in labor supply to occur following a change in wage rates. It will take even longer—perhaps 20 to 25 years—before the full effects of reduced saving on capital income tax revenues are realized. Thus, the immediate effects of a change in redistributive policy can be expected to be quite different, and appear much more favorable, than the ultimate effects.

This difference between the short-run and long-run effects of changes in the amount of redistribution suggests how society might inadvertently embark on a course that ultimately harms the poor who are the intended beneficiaries. If the political process tends to emphasize the short-run effects of policies, welfare programs could be extended to the point where the marginal effects are actually harmful to the interests of the poor themselves. This could be true even if the actual consequences are known; when they are not known and are as dif-

ficult to evaluate accurately as the analysis here may suggest, the possibility of an overexpansion in welfare programs cannot be dismissed. Again, we do not argue that this has actually occurred, but there is no doubt that a tendency to stress the immediate effects of welfare policies inclines public policy in that direction.

Conclusion

The relevant cost of helping the poor through redistributive tax and transfer programs of the sort now commonly used is likely to be quite high. Moreover, there is the possibility that further efforts to raise the incomes of the poor will be counterproductive, especially if a long-run point of view is taken. If one's conception of a good society includes an adequate standard of living for the poorest households, these are distressing implications. But two qualifications to these conclusions should be kept in mind.

First, our analysis has been concerned exclusively with the cost of a particular type of redistributive tax and transfer policy; a different type of redistributive policy could result in substantially lower costs. For example, there may be identifiable subgroups among the poor (for example, disabled persons) whose labor supply responses may be nonexistent or considerably lower than our assumed average response. The cost of targeting welfare aid on such groups would be less than the cost of aiding low-income people generally.

Second, redistribution is not the only way the incomes of the poor can be increased; another way is through economic growth. If real per capita income grows at a rate of 2 percent per year, in nine years income would be 20 percent higher. This is a larger gain than can probably be accomplished through increased redistribution, and it can occur without increasing the tax rates on the poor or the nonpoor. Thus, one of the implications of our analysis is to strengthen the case for economic growth as a way to reduce poverty. This is not to suggest that increasing the rate of economic growth is an easy matter, but there are steps the government could take in this direction.

Economic growth as an antipoverty policy has fallen into disfavor in recent years, in part because it is believed that the poor do not reap significant benefits from it. The basis for this belief is that many of the poor have low or zero earnings; a person with no earnings is not directly benefited when economic growth raises real wage rates. What this argument ignores, however, is that economic growth makes it possible to finance larger transfers without raising marginal tax rates. Insofar as expenditures on transfers are kept a constant proportion of national income, growth in national income benefits recip-

ients of transfers proportionately as much as it benefits other people. Economic growth therefore may be a better antipoverty policy than attempts to redistribute a larger share of national income.

Appendix

In this appendix, we modify the numerical example in the text to evaluate the marginal cost of redistributing income using a different policy than the one developed earlier. Figure 1 can be used to contrast the two policies graphically. Earnings of households are measured horizontally and disposable incomes are measured vertically. The 45-degree line illustrates the situation before the incremental redistributive policy is adopted; for simplicity, it is assumed that there are no preexisting tax or transfer policies, so earnings equal disposable incomes. The hypothetical policy examined in the paper then produces the relationship ABC. Disposable incomes are raised above earnings for those with earnings below $30,000 (they receive net transfers), while disposable incomes are reduced for those with higher incomes (they pay net taxes). Schedule ABC shows only the combined effects of the tax and transfer programs, and, as the diagram illustrates, the net effect is the same as a negative income tax applied to those with incomes below $30,000, with taxes collected from those

FIGURE 1

ALTERNATIVE REDISTRIBUTIVE POLICIES

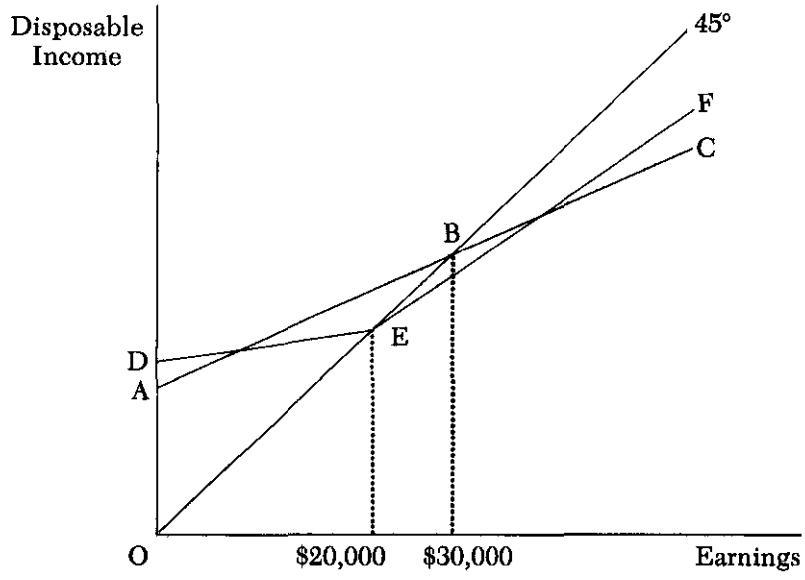

with incomes above. The marginal tax rate is indicated by (one minus) the slope of ABC. Since ABC is a straight line, the marginal tax rate is the same for transfer recipients and taxpayers. In our example, the marginal tax rate was 1 percent.

As an alternative redistributive program, let us consider a policy that restricts transfers to those with earnings below $20,000 and redistributes approximately the same amount of income. For transfer recipients, the relationship becomes DE, with larger transfers going to those with the lowest incomes; note that the marginal tax rate is higher now for low income households because the transfers must be tapered off more quickly to reach zero at $20,000 rather than $30,000. For taxpayers, the relationship becomes EF, showing that a lower tax rate can be used for taxpayers in this case when income in excess of $20,000 is taxed rather than income in excess of $30,000. Basically, this policy is simply another negative income tax, but one where the break-even level of income (the earnings level where the transfer is zero) is $20,000 rather than $30,000.

To evaluate whether this alternative policy involves a lower marginal cost, we use the same assumptions employed in the numerical example in the text (see Table 4). The one important point not shown explicitly in Table 4 is the increment in marginal tax rates. It was assumed that a 1 percent tax was applied to earnings above $20,000. By requiring that the government's budget be balanced (additional tax revenues sufficient to finance the additional transfers), we determined by trial and error that the marginal tax rate for those with incomes below $20,000 would be 2.2 percent.

Faced with an increase of 2.2 percentage points in their marginal tax rates, households A and B reduce their earnings by more than in the example in the text, as shown in column 3. The reductions for C, D, and E are the same because their marginal tax rates rise by 1 percentage point in both cases. The net effect on tax revenue for each household is shown in column 4: note that tax revenues fall for A and B because their initial marginal tax rates were 40 percent and their earnings fall. Total additional net revenue is $228, and this is distributed as transfers to A and B, as shown in column 5. The net transfer, the sum of columns 4 and 5, is shown in column 6. Finally, the change in disposable income, the sum of columns 6 and 3, is shown in column 7.

Households B, C, D, and E have their disposable incomes reduced by a total of $1086 with this policy, while household A has its disposable income increased by $156. This translates into a marginal cost of $6.96. Thus, the marginal cost to the upper four households of increasing the disposable income of the poorest household is just

TABLE 4

AN ALTERNATIVE TAX-TRANSFER POLICY

Household (1)	Initial Earnings (2)	Change in Earnings (3)	Net Additional Tax Revenue (4)	Transfer (5)	Net Transfer (6)	Change in Disposable Income (7)
A	10,000	−110	−44	222	266	156
B	20,000	−220	−88	6	94	−126
C	30,000	−150	40	0	−40	−190
D	40,000	−200	120	0	−120	−320
E	50,000	−250	200	0	−200	−450
Total	150,000	−930	228	228	0	−930

under $7, compared with the $7.82 marginal cost for the policy evaluated in the text. In this example, which concentrates transfers more completely on the lowest income households, the marginal cost is about 10 percent.

The important point suggested by this example is that the marginal cost is not likely to be significantly different for policies that may appear to be quite different from the one on which we based the estimates in the text. Although one example does not prove this general point, we have experimented with several other policies with the same results. Thus, we believe that the estimates of marginal costs presented in the paper are likely to be reliable indications of the general order of magnitude involved.

References

Browning, Edgar K., and Johnson, William R. "The Trade-off between Equality and Efficiency." *Journal of Political Economy* 92 (April 1984): 175–203.

Danziger, Sheldon; Haveman, Robert; and Plotnick, Robert. "How Income Transfers Affect Work, Savings, and the Income Distribution." *Journal of Economic Literature* 19 (September 1981): 975–1028.

Danziger, Sheldon. "Poverty and Inequality Under Reagonomics." *Journal of Contemporary Studies* 5 (Summer 1982).

Danziger, Sheldon, and Gottschalk, Peter. "The Poverty of *Losing Ground*." *Challenge* 28 (May/June 1985): 32–38.

Feldstein, Martin; Poterba, James; and Dicks-Mireaux, Louis. "The Effective Tax Rate and the Pretax Rate of Return." Working Paper no. 740, National Bureau of Economic Research, 1981.

Hansson, Ingemar, and Stuart, Charles. "Tax Revenue and the Marginal Cost of Public Funds in Sweden." *Journal of Public Economics* (forthcoming).

Lind, Robert C. "A Primer on the Major Issues Relating to the Discount Rate for Evaluating National Energy Options." In *Discounting for Time and Risk in Energy Policy*. Edited by R. Lind et al. Washington, D.C.: Resources for the Future, 1982.

Murray, Charles. *Losing Ground*. New York: Basic Books, 1984.