

NBER WORKING PAPER SERIES

OPTIMAL INCOME TRANSFERS

Louis Kaplow

Working Paper 12284

<http://www.nber.org/papers/w12284>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

May 2006

Harvard University and National Bureau of Economic Research. I am grateful to workshop participants and the referees for comments, Mary Bear, Ivan Chen, Eun Young Choi, Vicki Chou, Zachary Gubler, Jeffrey Harris, Jonathan Lin, Ned Locke, Kathleen Saunders, and Kevin Terrazas for research assistance, and the John M. Olin Center for Law, Economics, and Business at Harvard University for financial support. The current paper traces its origins to a book project, Kaplow (forthcoming). The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2006 by Louis Kaplow. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Income Transfers
Louis Kaplow
NBER Working Paper No. 12284
May 2006
JEL No. H21, H53, I38

ABSTRACT

A substantial literature addresses the design of transfer programs and policies, including the negative income tax, other means-tested transfers, the earned income tax credit, categorical assistance, and work inducements. This work is largely independent of that on the optimal nonlinear income tax, yet formulations of such a tax necessarily address how low-income individuals should be treated. This paper draws on the optimal income taxation literature to illuminate the analysis of transfer programs, including the level and shape of marginal tax rates (including phase-outs), the structure of categorical assistance, and the role of work inducements in an optimal income transfer scheme.

Louis Kaplow
Harvard University
Hauser Hall 322
Cambridge, MA 02138
and NBER
moverholt@law.harvard.edu

1. Introduction

Income transfers to low-income individuals have been the subject of a great deal of scholarly attention, and for good reasons. Many programs target the poor and near-poor, and these systems have been a continuing subject of controversy and reform. Low-income individuals are often believed to have certain traits – disabilities, age, or dependents – that render work difficult or inappropriate, thereby requiring income support. Moreover, under most standard social welfare functions and plausible individual utility functions, the marginal social welfare weight per dollar is substantially greater at the bottom end of the income distribution, so the design of income transfer programs is of great social consequence.


Most existing literature focuses on specific transfer schemes or particular characteristics thereof. Over the years, subjects receiving the most attention include the negative income tax and related systems, means-testing of transfer programs, earnings subsidies such as the earned income tax credit (EITC) in the United States tax system, categorical assistance, and other work inducements.¹ Analyses often isolate one or more features, such as program cost, work incentives, or effectiveness in targeting the needy. This work does not, however, typically provide a comprehensive assessment or attempt to determine what overall plan would be optimal.

There also exists an important body of work on optimal nonlinear income taxation.² This literature offers both analysis – which, although technically demanding, does yield first-order conditions that may be interpreted intuitively – and simulations. Of particular relevance for present purposes, the “income tax” examined in this literature does not correspond simply to a typical personal income tax schedule but rather encompasses all taxes *and transfers*. The schemes examined feature a grant, implicitly received by everyone, combined with a tax that depends on the level of income.

¹See, for example, Atkinson (1995), Garfinkel (1982), Green (1967), Meyer and Holtz-Eakin (2001), and Moffitt (2002, 2003).

²See, for example, Mirrlees (1971), Atkinson and Stiglitz (1980), Stiglitz (1987), and Tuomala (1990).

Figure 1. Income Tax and Transfer Schedule



In Figure 1, the schedule $T(y)$, showing taxes as a function of income, is a single, integrated tax and transfer scheme. Individuals who earn no income receive the grant $g = -T(0)$, those earning income y° are at the breakeven point at which the taxes they owe just equal the grant, and those earning more than y° pay positive net taxes.³ However, one can just as easily interpret Figure 1 as depicting an income tax that exempts income below y° combined with a separate transfer scheme that provides g , which is fully phased out when income reaches y° . (There are additional possibilities: $T(y)$ could be viewed as a standard income tax with an exemption less than y° , including one with no exemption, combined with a transfer scheme that provides g , which is not fully phased out until income reaches a level above y° .) Furthermore, in examining actual systems of transfer programs, one can interpret g as the sum of all forms of assistance available to those earning no income, and the marginal tax rate $T'(y)$ can be taken as the sum of the explicit marginal tax rates of the income tax and other pertinent taxes and the phase-out rates of various transfer programs. In other words, each tax and each transfer program can be represented by its own schedule $T^i(y)$, and we can let $T(y) = \sum T^i(y)$.⁴

Suppose that one has determined an optimal income tax schedule for a given society; perhaps it is that depicted in Figure 1. Then, a fortiori, one has determined the optimal system of income transfers, for that is indicated by the lower end of the very same income tax schedule. Likewise, one has determined how best to reform an existing system: If one subtracts the aggregate existing schedule from the optimal schedule, the difference indicates what reform of taxes and transfers would move the system to the optimum.

³As will be explored in section 3, one can readily extend the analysis to allow different schedules for different types of individuals or family units.

⁴A complication is that some transfer programs have a so-called cliff or notch effect, such that when income reaches a particular point, certain benefits are lost altogether; that is, some $T^i(y)$ may be discontinuous, so $T'(y)$ may not be defined at such points. (There also may be kinks in the tax schedule, but these are less troublesome for most analysis.) Such complications will generally be ignored here.

This approach to optimal income transfers differs substantially from that in most of the pertinent literature. Existing work, as noted, tends to focus on particular transfer programs or transfer programs as a whole. Changes in such programs, however, are not generally revenue-neutral and are almost never distribution-neutral, so unless one also considers effects on the rest of the population and aggregates them with a social welfare function, it is difficult to know what changes would be optimal. Even clear indicators may be misleading because they are partial: A program may reduce outlays or labor supply distortion, but at the expense of redistribution; a reform may improve work incentives, but since the optimum involves some degree of work disincentive, one cannot know whether the move is in the right direction. Moreover, it is familiar from research on optimal income taxation that an important effect of changing the marginal income tax rate at any given level of income – and particularly at low levels of income – is the inframarginal impact on those with higher incomes. Accordingly, it is difficult to illuminate questions about optimal income transfers by confining attention to explicit transfer programs at the lower end of the income distribution.

An additional limitation is that much work focuses on very specific programs, whether existing, proposed, or hypothetical. As just one component $T^i(y)$ of the overall schedule $T(y)$, there is no sense in which any particular program can be assessed or optimized without regard to the rest of the system. To be sure, if all other taxes and transfers are accounted for and taken as given, one can speak meaningfully of optimal reform of a single component, but if the component is restricted in various ways – for example, if the reform cannot change the grant g or if any adjustment to marginal tax rates is limited to lower-income individuals and must have a prescribed shape (say, that of the EITC) – the results may obscure the best paths for reform. (If aggregate marginal tax rates, including phase-outs, were too high in both the phase-in and phase-out ranges of the EITC, making the EITC more or less generous might be desirable if no other change is possible, but such a reform would be largely orthogonal to the true problem.)

Furthermore, it is necessary to be consistent (and plausible) regarding what policy instruments are assumed to be feasible. For example, if a form of work requirement necessitates that hours and thus implicitly wage rates be observable, such information could be used to implement a tax-transfer scheme that was less distortionary than (and qualitatively different from) typical work inducement proposals. But if such a scheme could readily be circumvented due to the manipulability of reported hours, which is commonly supposed, then the more conventional work inducement may be impractical as well.

The purpose of this article is to analyze optimal income transfers as part of the broader optimal income taxation problem. This approach is initially advanced in Diamond's (1968) review of Green's (1967) book on the negative income tax. Diamond focused on how to choose transfer schedules optimally, foreshadowing important elements of the literature on optimal income taxation that appeared shortly thereafter. Furthermore, characterizing optimal income transfers was one of the motivations offered by Mirrlees (1971) for his exploration of the optimal nonlinear income tax, and he remarked on the inability to address transfers without regard to the rest of the income tax schedule. In a sense, the present article aims to pick up where Diamond

(1968) and Mirrlees (1971) left off, taking advantage of insights in the subsequent literature.⁵

Section 2 provides the most direct application of results on optimal nonlinear income taxation. Initially, the standard first-order condition for that problem is presented in an intuitive form (which has appeared in some of the writings since Mirrlees (1971)) and interpreted with special emphasis on the lower end of the income distribution. Additionally, results from simulations are noted. Both the analysis and the simulation results support – very roughly – fairly high marginal rates at the bottom, although lower than existing aggregate marginal rates that can approach or even exceed one hundred percent. Furthermore, in the low to moderate income range, the level of optimal marginal rates is fairly flat or gradually declining, which stands in sharp contrast to the very large drops in marginal rates (notably, at the end of phase-out ranges) in existing systems and under many proposals. It is suggested that existing thought and practice reflects misplaced emphasis on the idea that transfers need to be phased out quickly and, once they are, that marginal tax rates should be fairly low on moderate-income individuals. These conclusions are applied to existing discourse about extremely high aggregate marginal tax rates, the design of the EITC, and other matters.

Section 3 examines categorical assistance, such as special programs targeted at the disabled, elderly, or families with young children. Analytically, the approach involves a modest modification to the optimal income tax model already set forth, which can then be interpreted to illuminate the optimal design of categorical assistance. Most discussion focuses on the realistic setting in which categorization is imperfect; for example, some classified as able may be disabled and some deemed to be disabled may nevertheless have high earning ability. It turns out that the optimal tax and transfer schemes for such two-category systems are qualitatively different from each other. It is not simply the case that the more able group receives less generous assistance. It also seems plausible that the schedule of marginal tax rates is distinctive, and in a manner that deviates from common thinking and practice. Specifically, it may be optimal to apply higher marginal tax rates (phase-outs) to the group receiving lower assistance, even though there are only modest benefits to “phase out,” while applying lower rates to the group receiving more generous assistance. Once again, understandings based on the supposed need to phase out program benefits can be misleading.

Section 4 analyzes schemes that embody work inducements of various sorts. As will be explained, many work incentive plans do not literally require work but instead adjust benefits in light of earnings or work effort. In other words, individuals who are classified as able but do not meet target levels of work effort have their benefits reduced or eliminated. It turns out that the optimal design of earnings-based work inducements is implicit in the analysis of section 3.

⁵A subset of the literature on income transfers seeks to examine optimal schemes. See, for example, Liebman (2001), the writing on categorical assistance surveyed in subsection 3.3, and some of the literature on work inducements discussed in subsection 4.4. However, little research takes advantage of the general, flexible, and encompassing framework offered by the work on optimal nonlinear income taxation, and accordingly it is not able to generate many of the insights produced here.

Having already determined the optimal scheme for individuals classified as able, the analysis is essentially complete. Moreover, the optimal scheme does not have the characteristics of standard work incentive programs. A different case arises if not only earnings but also the amount of work effort can be observed. If there is perfect observability of effort, schemes better than and qualitatively different from standard work requirements become feasible. In the more realistic case of imperfect observability, features of the preceding analysis become applicable, with the implication that important features of various work requirement plans are not optimal.⁶ Finally, this section explores the important possibility – which may underlie welfare reforms in recent decades – that there are various sorts of externalities to work of low-income individuals and then considers how such externalities would affect the analysis.

2. Optimal Income Taxation and Optimal Income Transfers

2.1. Framework

The standard optimal nonlinear income tax model has the following basic elements. An individual's utility is given by $u(c,l)$, where c is consumption and l denotes labor effort. An individual's consumption (equivalent to disposable income) is given by

$$(1) \quad c = wl - T(wl),$$

where w is the individual's exogenously given wage rate. Individuals' pre-tax earnings are the product of their wage and effort level, that is, $y = wl$. The motivation for redistributive taxation is that individuals differ in their wages, that is, their earning abilities. The distribution of abilities is given by $F(w)$, with density $f(w)$, the population being normalized to have a total mass of one. Furthermore, it is supposed that the government cannot directly observe individuals' abilities – if it could, any distributive objective to be achieved with nondistortionary individualized lump-sum taxes – and thus must rely on distortionary income taxation.

Individuals choose the level of labor effort l that maximizes $u(c,l)$ subject to their budget constraint (1).⁷ The government's problem is taken to be the choice of a tax-transfer schedule

⁶A range of additional topics – including whether assistance should be provided in cash or in kind, concerns involving two-earner families, labor market interventions such as a minimum wage or subsidies to the development of human capital, effects of transfers on wage levels, and insurance such as for temporary unemployment – are beyond the scope of this article, although the framework advanced here should prove useful for exploring them as well.

⁷Although the formulation in the text is standard in the literature, one might question the extent to which individuals, especially low-income individuals who tend to be less sophisticated and face multiple marginal rates from a variety of complex programs, choose labor supply as the theory of the perfectly informed rational maximizer would predict. Low-income individuals may be more aware of average than marginal rates, but may approximately learn marginal rates over time as they vary their work effort or obtain information from others similarly situated. In the absence of empirical evidence on the structure of individuals' misperceptions, it is not apparent

$T(wl)$ to maximize social welfare,

$$(2) \int W(u(c(w), l(w))) f(w) dw,$$

where c and l are each expressed as functions of w to refer to the level of consumption achieved and labor effort chosen by an individual of type (ability) w . This social maximization is subject to a revenue constraint and to a set of constraints regarding individuals' behavior. The former is

$$(3) \int T(wl(w)) f(w) dw = R,$$

where R is an exogenously given revenue requirement. Here, revenue is to be interpreted as expenditures on public goods that should be understood as implicit in individuals' utility functions; because these expenditures are taken to be fixed, they need not be modeled explicitly. Regarding the latter set of constraints, individuals are assumed to respond to the given tax schedule optimally, that is, by choosing labor effort to maximize their own utility, taking the tax schedule as given.

2.2. Analysis

Much of the analysis can be summarized in a first-order condition for the optimal marginal tax rate at any income level y^* , where w^* and l^* correspond to the ability level and degree of labor effort supplied by the type of individual who would just earn y^* . Following the presentation in Diamond (1998) – which makes the simplifying assumptions that utility is separable between consumption and labor effort and that marginal utility u_c is constant – the condition can be expressed as⁸

how this potentially important complication is best analyzed, but absent reasons to believe that there are large, systematic errors, one suspects that the rough guidance provided by standard analysis remains illuminating.

⁸In addition to Diamond (1998 (expression (10) on page 86), see Atkinson and Stiglitz (1980) (expression (13-54) on page 417), Stiglitz (1987) (expression (25) on page 1007 and the expression in note 17 on page 1008), Dahan and Strawczynski (2000) (expression (2) on page 682), and Auerbach and Hines (2002) (expressions (4.12) and (4.15) on pages 1381-82). Expression (4) is also similar to the two formulations in Saez (2001, p. 215); an important difference involves the attention he devotes to translating the results from ones in terms of the distribution of abilities (which is unobservable) to ones in terms of the distribution of income (taking into account that the tax rules and other parameters will determine the relationship between the ability distribution and the resulting income distribution).

$$(4) \frac{T'(w^*l^*)}{1 - T'(w^*l^*)} = \frac{1 - F(w^*)}{\xi^* w^* f(w^*)} \frac{\int_{w^*}^{\infty} \left(1 - \frac{W'(u(w))u_c}{\lambda}\right) f(w)dw}{1 - F(w^*)},$$

where primes indicate derivatives with respect to a function's only argument, λ is the shadow price of government revenue, and $\xi^* = 1/(1+l^*u_l/u_c)$ – which, when utility is quasi-linear as assumed here, equals $\varepsilon/(1+\varepsilon)$, where ε is the elasticity of labor supply. Note that this formulation (like those in recent literature) includes $1 - F(w^*)$ in both the numerator and the denominator on the right side. The motivation is that, in the first term, $(1 - F(w^*))/f(w^*)$ is purely a property of the distribution of w , and, in the second term, because the numerator is an integral from w^* to ∞ , the term as a whole gives an average value for the expression in brackets in the integrand. Both aspects facilitate interpretation, as will be seen in the discussion to follow.

To aid in understanding expression (4), it is helpful to have in mind the simple perturbation of the income tax schedule that underlies this first-order condition.⁹ If one begins with some tax schedule $T(wl)$, assumed to be optimal, it must be that no slight adjustment to the schedule will change the level of social welfare. Consider an adjustment that slightly raises the marginal tax rate at some income level, y^* (say, in a small interval from y^* to $y^*+\delta$), leaving all other marginal tax rates unchanged. There are two effects of such a change. First, individuals at that income level face a higher marginal rate, which will distort their labor effort, a cost. Second, all individuals above income level y^* will pay more tax, but these individuals face no marginal distortion. That is, the higher marginal rate at y^* is inframarginal for them. Since those thus giving up income are an above-average slice of the population (it is the part of the population with income above y^*), there tends to be a redistributive gain.

Expression (4) can readily be interpreted in terms of this perturbation.¹⁰ Begin with the first term. Revenue is collected from all individuals with incomes above y^* , which is to say all ability types above w^* ; hence the $1 - F(w^*)$ in the numerator. One distorts only the behavior of the marginal type, which explains the $f(w^*)$ in the denominator. The larger is the fraction of the

⁹This sort of perturbation is discussed in Diamond (1968), and he offers the conjecture that optimal rates at the bottom are high, for reasons that overlap with those offered in the text to follow. See also Saez (2001).

¹⁰There are a variety of reservations to any such interpretation. One set concerns the fact that all values are endogenous, so differing presumed values of one component affect the values of others at the optimum. Another is that the problem is not necessarily well behaved. In particular, in regions of falling marginal tax rates or fixed costs of labor market participation, small changes in the tax schedule may lead individuals to “jump” from one level of income to another, in which case a different condition governs their behavior and there would be gaps in the resulting income distribution. In addition, simplifications such as the assumption that u_c is constant, which rules out income effects, have been employed. Such complexities are largely ignored for present purposes, but they may prove important, especially in designing simulations. See also the discussion of the participation decision in subsection 4.4.

population paying more tax and the smaller is the group being distorted – or put more succinctly, the greater is the ratio of the former to the latter – the higher is the optimal tax rate. The denominator also contains weights of ξ^* , indicating the extent of the distortion, and w^* , indicating how much production (and thus tax revenue) is lost per unit of reduction in labor effort.

The second term applies a relative social weighting to the revenue that is collected from inframarginal individuals; it is an average weight for the portion $1 - F(w^*)$ that pays more tax. Regarding the numerator, the term in parentheses is the difference between the marginal dollar that is raised and the dollar equivalent of the loss in welfare that occurs on account of individuals above w^* paying more tax: u_c is the marginal utility of income to such individuals, W' indicates the impact of this change in utility on social welfare, and division by λ , the shadow price on the revenue constraint, converts this welfare measure into dollars.¹¹

We now consider what expression (4) reveals about optimal marginal rates at the lower end of the income distribution. It appears that fairly high rates may be optimal. Focusing on the first term on the right, we can see that three factors contribute to this result. First, the numerator, $1 - F(w^*)$, is large: Raising marginal rates on very low incomes raises substantial revenue (allowing for a higher g) because most of the population has incomes higher than this level. Moreover, high rates at the bottom are inframarginal for this large group of individuals.¹² Second, in the denominator, $f(w^*)$ is not very high, indicating that only a moderate portion of the population has their labor supply distorted by high marginal rates in this income range. Third, w^* is low, so there is little lost productivity and thus revenue when such individuals reduce their labor supply. Fourth, although the elasticity component, ξ^* , is often taken to be constant, there is some evidence indicating that, if anything, the pertinent elasticity may be lower for low-income individuals, which would also contribute to a higher optimal marginal tax rate.¹³ There is,

¹¹Another natural way to think of the experiment of raising the marginal rate $T'(w^*l^*)$ is to suppose further that the additional revenue will be used to increase the uniform grant. The marginal social value of increasing the grant will, at the optimum, necessarily equal the shadow price λ of government revenue. Hence, the formulation in the text in which the rebate is not carried through but is given a shadow value is formally equivalent to what would obtain if one literally adhered to the revenue constraint by increasing the grant until the budget balanced. (This relationship of course is standard for constrained optimization that employs shadow prices.)

¹²Additionally, if one relaxed the assumption that u_c is constant, the resulting income effect would induce greater labor supply, which further increases revenue.

¹³Empirical evidence on the elasticity of taxable income tends to find higher elasticities among higher-income individuals, apparently due to differences in tax avoidance opportunities. See Alm and Wallace (2000), Auten and Carroll (1999), Gruber and Saez (2002), and Moffitt and Wilhelm (2000). By contrast, Juhn, Murphy, and Topel (1991) find that uncompensated labor supply elasticities for men fall with income. It is suggested, however, that the broader concept of elasticity of taxable income is more relevant to welfare analysis because all margins of behavioral response to changes in tax rates involve the same marginal distortion. See Feldstein

however, a countervailing factor: The second term, reflecting the average welfare weight for individuals with abilities above w^* , provides some offset to the argument for high marginal rates at the bottom.¹⁴ In this range, those inframarginal include some fairly low-income individuals, and their welfare weight is relatively high.

2.3. Simulations

The conjecture that optimal marginal tax rates will be high at the bottom of the income distribution (and possibly higher than at middle or upper levels of income) is confirmed by simulations on the optimal nonlinear income tax. Mirrlees (1971) original simulations largely have marginal rates in the 20's and 30's at the low end of the income distribution although one case has rates in the high 50's and as high as 60%. These low-end rates are the highest of the marginal rates in his simulations. Tuomala (1990) offers a wide range of simulations; with a utilitarian (rather than more concave) welfare function and his preferred substitution elasticity (table 6.1, cases 4-6), marginal rates at the tenth percentile of the income distribution are in the 40's. Kanbur and Tuomala (1994) report results for differing degrees of inequality in the distribution of abilities and find rates at the bottom in the range 49% to 60%. Saez's (2001) utilitarian base case has marginal rates at the bottom in the 70%-80% range. Additionally, Slemrod et al.'s (1994) simulations for the optimal two-bracket system find that the optimal lower bracket rate is approximately 60%, which is slightly higher than the optimal upper-bracket rate.

Whatever the absolute level of marginal rates in a particular simulation, it is typical for the optimal marginal rates just above the low end – running, say, from levels of income where welfare phase-outs ordinarily are complete to the median level of income – to be almost as high. In most instances, the falloff is only a few percentage points, with somewhat higher drops when the initial (bottom-end) rates are higher. Exceptions are Kanbur and Tuomala (1994), who report a rise in marginal rates in two of their three simulations and Saez (2001) who finds drops (from an initial level of 70-80%) of approximately 30% (with the lowest rate being reached at an income of about \$75,000).¹⁵ The finding that marginal rates do not rapidly fall is to be expected

(1999).

¹⁴At y^* near zero, this offset would be complete if even the lowest-ability types worked. This is the familiar “zero at the bottom” result of the optimal nonlinear income taxation literature. See Brito and Oakland (1977), Seade (1977), and Ebert (1992). The usual argument for high marginal rates near the bottom of the income distribution is inapplicable when all earn above some minimum level because it is more efficient to reduce the grant g than to apply a positive marginal tax at the bottom. Put another way, raising the marginal rate from zero for the bottom person accomplishes no redistribution (everyone pays more tax and receives a correspondingly higher grant) but does distort that individual's labor supply. As the simulations to follow reveal, however, this force toward a zero rate at the bottom is not nearly complete if a significant number of low-ability individuals do not work at all, as is actually the case.

¹⁵The drop in Saez (2001), although atypical, is not that surprising in light of the foregoing discussion of the first term of expression (4), specifically the ratio $(1 - F(w^*)) / w^* f(w^*)$.

in light of the first-order condition. Although some factors favor falling rates, not all do and none is likely to favor precipitously declining rates for modest differences in income.

Finally, the optimum is typically characterized by a nontrivial fraction of individuals, those with the lowest abilities, choosing not to work at all. Mirrlees (1971) and Tuomala (1990) have approximately 5%-10% not working under the optimal scheme in most simulations, although this percentage goes as high as 20%. This finding also is not surprising. Given that high marginal tax rates are optimal near the bottom and that a generous grant is provided, individuals of very low ability are unlikely to find work worthwhile. (It is of course possible that some would have sufficiently low ability to render it optimal to abstain entirely from work even if the marginal rate at the very bottom were zero.) Furthermore, the productivity and therefore the revenue loss when such individuals do not work is small.

2.4. Implications

In the United States, there are multiple taxes and transfer programs relevant to many low-income individuals. Some transfer programs have very high phase-out rates, others more gradual, and the tax system includes the EITC, which provides roughly a 40% marginal subsidy to low-income workers with two or more children but adds approximately a 20% tax in its phase-out range. Taken together, Giannarelli and Steuerle's (1995) microsimulation suggests that the aggregate marginal tax rate (ignoring work expenses) is 75% or more for many low-income workers and above 100% for some. Examining programs as of 1998 (after welfare reform), Sammartino et al. (2002) find that a single parent with two dependents faces average marginal tax rates of about 60-70% when increasing earnings from half the poverty line to 125% thereof, of over 100% in moving to 150% of the poverty line, and of just over 30% beyond that point. Other studies reach similar conclusions.¹⁶

Such high aggregate marginal rates (inclusive of phase-outs) are generally criticized. As suggested above, optimal marginal rates at the low end are quite high, but simulations do not suggest that they should be as high as actually exists. Furthermore, once past the phase-outs of transfers (including the EITC), existing aggregate marginal rates seem to be below what is optimal. Likewise, many proposals over the years, including the negative income tax and variants, contemplate a system with very high marginal rates on the poor, dropping (often instantaneously, at a breakeven point) to far lower rates on the near-poor. This combination of excessive marginal rates initially, followed by suboptimal marginal rates, seems to be a product of thinking oriented toward phase-outs. If the only question is what structure will minimize program costs – taking initial benefit levels as given – very rapid phase-outs may be cheapest (although they may not, due to work disincentives), and there is no apparent need to continue to tax individuals heavily once benefits are fully phased out. By contrast, if one views the system as

This explanation is suggested noting the similarity between Saez's Figure 4, which graphs a similar relationship as a function of income, and his Figure 5, which graphs simulated optimal income tax schedules.

¹⁶See, for example, Dickert, Houser, and Scholz (1994), Keane and Moffitt (1991), and Wilson and Cline (1994), and after welfare reform, Acs et al. (1998), Gokhale, Kotlikoff, and Sluchynsky (2002), Hepner and Reed (2004), and Shaviro (1999).

a whole, recognizes that phasing out any particular benefit that is just one component of the system is not a particularly meaningful objective, and cares about overall welfare maximization (noting, for example, that higher marginal rates beyond normal phase-out ranges may finance basic grants), the foregoing analysis and simulations suggest a different understanding of transfer programs that leads to different prescriptions.

Finally, the high marginal rates at the bottom are seen as an important contributor to low work effort (often, no market labor effort) by welfare recipients. It should be kept in mind, however, that even if marginal rates were set optimally, they would be high, and a significant number of low-ability individuals would not work. Work effort, to be sure, is important – both overall productivity and revenue are relevant – but it is not the ultimate objective; as one aspect of the optimization, it must be traded off against others. Furthermore, even confining attention to work effort, it must be recalled that high marginal rates at the bottom are inframarginal to individuals above the bottom, allowing more revenue to be raised without distortion of their behavior. (For further discussion, see section 4 on work inducements.)

A number of these points can be illustrated by reference to the EITC. The 40% marginal subsidy is sometimes viewed as a pure earnings subsidy, but the justification for subsidizing earnings per se is not clear. Sometimes it is rationalized as an offset to payroll and other taxes for low-income earners. In addition, it can be seen as a reduction in welfare phase-outs: Indeed, for qualifying individuals, reducing the welfare phase-out rate (say, in TANF and/or food stamps) from 100% to 60% is identical to keeping the rate at 100% and providing a 40% subsidy through the EITC. This correspondence conflicts with the common characterization of the EITC as the opposite of a welfare expansion, the latter being seen as involving greater welfare expenditures and the extension of welfare benefits to less needy individuals. The difference in rhetoric masks an underlying similarity in reality.

The EITC phase-out can also be described in various ways. Most commonly, it is understood as a precise phase-out of the EITC's subsidy. However, the system would be identical if there were no phase-out whatsoever, and ordinary tax rates were correspondingly higher on the pertinent group of individuals. Indeed, many in the phase-out range also pay a positive marginal tax rate on the same income, reported on the same income tax form.

These alternative interpretations of the EITC's subsidy (phase-in) and tax (phase-out) ranges are related to the verbally varied yet substantively identical descriptions of Figure 1 and the cross-over point y° given in the Introduction. That is, one can view the diagram as depicting a pure welfare grant g that is phased out by y° , followed by an income tax that exempts the first y° of income, as a welfare system that is not phased out until incomes above y° combined with a positive income tax that begins below y° , and so forth.

A further observation is that it is most unlikely that the optimal tax-transfer system looks anything like the EITC, which has marginal rates rising by 60 percentage points over a short range of income, given that the optimal system has fairly flat (gently falling) marginal rates in this range. However, if the system without the EITC has marginal rates that are too high at the bottom and too low on the near-poor, which may well be the case, then the EITC's otherwise odd

combination of subsidy and tax may make sense. This justification, however, is expressly conditional on the EITC offsetting other suboptimal features of the system, and it does not favor the EITC in particular, by contrast, for example, to a reform that directly makes phase-outs of welfare more gradual.¹⁷ Moreover, this view does not rationalize analyses of the EITC that are conducted as if the poorest individuals really face a net earnings subsidy of 40% and the near poor really face a net tax of only 20%, by contrast to the reality revealed by the previously noted microsimulations, which do include effects of the EITC.

Brief consideration of the EITC makes clear that most of what can be said about any single component of the tax-transfer system is necessarily incomplete and often misleading. What matters, to low-income individuals themselves and to society as a whole, is the aggregate, not the form of any particular piece or how one may choose to describe that component. Furthermore, if one seeks to design an optimal system, or to identify directions of reform, thinking that focuses on particular programs, phase-outs, notions like means testing or earnings subsidies, and the like provides little illumination. By contrast, existing analysis of optimal nonlinear income taxation directly addresses the pertinent questions.

3. Optimal Categorical Assistance

Transfer programs and, to a lesser extent, income tax codes often divide individuals into categories. Eligibility or the degree of generosity or stringency of various provisions may depend on disabilities, age, household characteristics (the presence of children, whether there is a single parent), and other factors. An important motivation for classification relates to the fact that the labor-leisure distortion caused by the tax-transfer system is due to the inability to observe earning ability directly. Various traits, however, may be correlated with ability and thus may be useful in designing a more efficient redistributive system.¹⁸

3.1. Framework and Analysis

To analyze this possibility, it is helpful to think of the government as being able to set a

¹⁷When different systems are used, the whole often lacks coherence due to different income definitions, eligibility standards, and so forth – so one ends up with income bands having even higher than the already excessive marginal rates for some groups and having far too low rates for others, as documented by many of the microsimulation studies reported above. Administrative considerations and somewhat different program purposes may explain some of the irregularities, but Giannarelli and Steuerle (1995), Wilson and Cline (1994), and others suggest that many proposals are formulated without a clear picture of the overall situation of lower-income individuals and the true impact of the reforms.

¹⁸There are other possible differences among groups. For example, one could allow the utility function in expression (5), below, to differ among groups, recognizing, for example, that both utility levels and the marginal utility of consumption could be affected by disabilities or family composition. See Kaplow (forthcoming, chapter 12). Additionally, externalities to work effort (noted in subsection 4.5) may vary.

separate income tax and transfer schedule for each identifiable group. Letting θ denote the observed parameter (which can be interpreted as an index of discrete classifications or as a continuous variable), we can modify the first-order condition for the optimal nonlinear income tax (4) as follows:

$$(5) \frac{T'(w^* l^*, \theta)}{1 - T'(w^* l^*, \theta)} = \frac{1 - F(w^*, \theta)}{\xi^* w^* f(w^*, \theta)} \frac{\int_{w^*}^{\infty} \left(1 - \frac{W'(u(w)) u_c}{\lambda}\right) f(w, \theta) dw}{1 - F(w^*, \theta)}.$$

The only difference between expressions (5) and (4) is that the tax function, T , in (5) is allowed to depend on θ and the density and distribution functions, f and F , depend on θ . The optimizations for each value of θ are linked by the common shadow price on revenue, λ . One can think of the marginal dollar being distributed pro rata across the entire population or being concentrated on certain types; which assumption is made does not matter because, at the optimum, the marginal social value of increasing the transfer will be the same for each type.

Initially, consider a simple case in which it is possible to observe perfectly which individuals have abilities below some low level, w° . Then that group can be given a high transfer g , which would not be that costly to finance because g could be fairly low for everyone else without fear that such individuals would be destitute because, by assumption, they all can earn at least a minimal income. Relatedly, it would be optimal not to tax low levels of earnings in the group for whom $w \geq w^\circ$ because this would avoid any labor supply distortion at the bottom of the group. (This is an instance of the optimal nonlinear income tax result that a zero marginal rate at the bottom is optimal in the case in which everyone works.¹⁹)

More realistically, signals about ability will be noisy. First, many traits, notably various forms of disability, can be observed only imperfectly. Second, even if perfect observation is possible, which may be nearly so with age and certain disabilities, there usually will be differences in ability associated with these features. For example, those with no physical disability obviously vary greatly in earning ability, and some individuals with significant physical impairments nevertheless have high earning ability.²⁰

For concreteness, suppose that a very low-cost signal (sufficiently cheap that its cost will be ignored) makes it possible to divide the population into two groups, group L consisting mostly of individuals with very low ability and group H containing relatively few such individuals. That is, by reference to the population density function $f(w)$, the density $f^L(w)$ is heavily concentrated at low levels of w and the density $f^H(w)$ is thin at the bottom. It may be useful to think about the

¹⁹See the discussion in note 14. Ordinarily – that is, without perfect categorization of those least able – it is optimal for some individuals not to work, so this result is inapplicable.

²⁰These points are emphasized by Diamond and Sheshinski (1995), among others. In addition, Benitez-Silva, Buchinsky, and Rust (2004) offer evidence that both types of classification errors are common in the social security disability system in the United States.

case in which group L consists of individuals with physical disabilities and group H consists of everyone else. Group L would then have the stated features and, once group L is removed from the rest of the population, those who remain, group H , would necessarily have a thinner density at the bottom.

To begin, it seems plausible that $g^L > g > g^H$. Because transfers to the low-ability group are limited to a subset of the population and, moreover, these individuals are disproportionately the neediest, their optimal grant will be high, both by comparison to the optimal grant under a single schedule (g) and, even more so, by comparison to the optimal grant for the group consisting mostly of more able individuals.

The validity and strength of this conjecture, however, depends on the shape and level of the tax schedules T^L and T^H above $y = 0$. To examine these tax schedules and gain insights into how they may differ from each other, the first-order condition (5) proves helpful. It should be kept in mind, however, that, as noted previously (see note 10), interpretations based solely on the first-order condition are inevitably speculative. Ultimately, the problem may best be illuminated through simulations of actual programs, wherein the greatest challenge will be ascertaining with reasonably accuracy each group's underlying density function, which as will be seen drives the analysis in important respects.

With this caveat in mind, begin with the more able group and focus initially on low levels of income. The first term in expression (5) will be notably higher than in the single-group version of the problem. The $1-F$ component in the numerator will be somewhat greater because almost everyone in the group will have higher incomes. More significantly, the f component in the denominator will be smaller, indeed, very small if the categorization is reasonably accurate. This suggests that the optimal marginal tax rate at low levels of income should be substantially higher than in the standard problem. (The resulting work disincentive at low levels of income will not be very costly because so few individuals fall in that range, and also because, as in the standard problem, productivity and therefore lost tax revenue is low in any event.)

Some offset to the optimality of very high marginal rates at the bottom is provided through the second term of expression (5). If $g^H < g$ and higher marginal tax rates are employed at low income levels, then individuals at higher income levels (associated with abilities $w > w^*$) will have somewhat higher welfare weights. This offset is likely to be most significant at the very bottom of the income scale, but only if g^H is significantly less than g , which may not be optimal.²¹ Note that, with very high marginal rates at the low end of the income distribution – rates that optimally are even higher if g^H is higher – a larger grant for the able group is less costly

²¹The problem of misidentifying very-low-ability individuals as having higher ability might be mitigated, even with a fairly low g^H , by making available public service employment, which might be designed as a screening device that would tend to be attractive only to individuals truly of low ability, who were not eligible for g^L because of misclassification. See the discussion of Brett (1998) in subsection 4.4.

than in the standard problem.²²

At higher levels of income, there may be less deviation between the optimal tax schedule for high-ability individuals and the optimal common tax schedule for the standard problem. At any given y , $1-F$ will be somewhat larger, favoring higher rates; f will also be somewhat larger, favoring lower rates; and the second term will continue to favor somewhat lower marginal rates. The direction of the net impact of these factors is a priori ambiguous. In any event, the magnitude of each of these factors will diminish as y increases, so it seems that for higher-income individuals in the high-ability group, the optimal schedule would be close to that in the case without categorization.

For the less able group, the results at the low end of the income distribution reverse. The $1-F$ component will, after extremely low levels of income, be substantially smaller than in the combined problem, and f will be much larger; these factors favor quite low marginal rates in this income range. Some offset will be provided by the second term: If the grant is more generous and initial marginal rates are low, then the welfare cost of higher payments by those with greater income will be less than otherwise. At higher levels of income, both $1-F$ and f will be much lower than in the standard problem, which has no clear effect on the optimal level of marginal tax rates, and the reduced second term will differ less from that in the common problem as y increases.

3.2. Implications

As a very crude first approximation, the system in the United States can be described as having the following elements: greater levels of g to individuals or families deemed unable, marginal tax rates at very low levels of income consisting substantially of phase-outs of grants and hence much higher rates for those receiving significantly more generous grants, and a common income tax schedule on incomes beyond the phase-out ranges.²³ That individuals in

²²The extent to which this is true depends on how able are those in the able group; for example, if the able (ignoring those misclassified) are individuals only capable of at least a minimum wage job, there may be many individuals who would not work at all if g^H was fairly generous and low-end marginal tax rates were very high. Note that the situation described in the text reverses for the low-ability group: As the text to follow explains, optimal marginal tax rates are likely to be fairly low at the bottom, which makes a higher grant more costly. Therefore, it is not inconceivable that the optimal grant would be higher for the more able group. (Consider, for example, a density $f^H(w)$ having a mass of individuals with abilities near zero, virtually no one else of low ability, and everyone else of abilities above a moderate level.) However, for some categorizations such a scheme may not be incentive compatible with regard to the classification system, as low-ability individuals would attempt to mimic high-ability individuals, such as by denying that they were really disabled.

²³For this purpose, the EITC can, as noted in subsection 2.4, usefully be viewed as a reduction in the phase-out rate for a certain group, namely, families, especially single heads of households with children, who also tend to be eligible for more generous welfare assistance.

lower-ability groups receive more assistance when they have no earnings is likely to be optimal, although the tendency to give very meager assistance to able individuals without earnings is optimal only if there is negligible misclassification.²⁴ The foregoing analysis had no strong implications for how the tax schedules for low- or high-ability groups should deviate from the optimal common schedule once income passes the lower end of the distribution, so the failure of the existing income tax to make significant distinctions is not obviously problematic.²⁵

Existing phase-outs, however, do not seem to reflect the basic features of optimality. It was suggested that optimal marginal tax rates on fairly low incomes may be rather high for the high-ability group and low for the low-ability group. But welfare phase-outs tend to have the opposite character: When benefits are high, as they are for low-ability groups, aggregate phase-out rates are correspondingly high because more benefits are being phased out. But it was just shown that optimal aggregate (phase-out inclusive) marginal tax rates for such individuals may be low. Moreover, this result holds even if it implies that a substantial grant is not fully phased out until income reaches fairly high levels.²⁶ For high-ability groups, existing benefits are low so there is little to phase out, and phase-out rates are correspondingly low. However, the foregoing analysis explains that high marginal tax rates may nevertheless be optimal. To be sure, work would be discouraged, but by assumption there are few whose abilities would put their incomes in this range and the high marginal rates are inframarginal for everyone else.

This apparent deviation from optimality seems to be a product of the beliefs that, when transfers are granted, they must be phased out and that the phase-out must be complete at reasonably modest levels of income, lest welfare become too expensive and available to non-needy individuals.²⁷ And when there is little welfare to be phased out, there is thought correspondingly to be no need for high marginal tax rates. Such reasoning views components of the tax and transfer system in a vacuum and fails to engage explicitly in optimization of the system as a whole by reference a well-specified social welfare function. Thus, the point y° in Figure 1, the income level at which taxes and transfers net to zero, should be determined as a byproduct of optimization, not chosen as a policy target, somehow arising exogenously. Furthermore, in a categorical system, there will be different $T(y)$ schedules for different groups, each with their own associated breakeven point y° . These points need not be the same, and as the

²⁴For evidence that significant numbers of disabled individuals are misclassified as able, see Benitez-Silva, Buchinsky, and Rust (2004). Note, however, that individuals thus misclassified are likely to be less severely disabled.

²⁵Income tax schedules above the low end of the income distribution do adjust for marital status and whether a taxpayer is a head of household, but these adjustments may be motivated by differences in utility functions rather than by differences in earning abilities.

²⁶Parsons (1996) argues that a sharp departure from optimality seems to arise in some disability programs, wherein any significant work may be taken as evidence of mistaken classification and hence terminate eligibility. Given that nontrivial classification errors are made, it may be better if misclassified individuals are not subject to prohibitive marginal tax rates.

²⁷This assumption about transfers and phase-outs often characterizes formal analyses, not just political debate, as reflected, for example, in Moffitt's (2002) survey.

foregoing analysis suggests they might differ significantly in an optimal transfer system.

3.3. *Prior Literature*

Others have examined how classification may enhance the ability to redistribute income. In Akerlof (1978), some individuals can be identified (“tagged”) as low-ability types. In his analysis, all those so identified are in fact of low ability (though some low ability individuals are not thus identified), there are only two levels of ability and two types of jobs, and he allows only for lump-sum transfers across groups.²⁸ Parsons (1996) extends Akerlof to allow both types of classification errors; however, he also considers only two types of individuals and two levels of effort (work and no work), so many features of the optimal tax schedule cannot be explored.

Stern (1982) compares an undifferentiated nonlinear income tax with a purely proportional income tax supplemented by differential lump-sum transfers that depend on a noisy signal of ability. A significant finding is that, the greater the concern for equality, the less desirable is the classified scheme because errors involving low-ability individuals misclassified as high-ability types are particularly costly. He does not allow for a nonlinear income tax in his classified scheme and his model has only two types, so the nonlinear income tax in the unclassified scheme can perform reasonably well without supplementation. In the analysis here, the misclassification problem is partly mitigated by allowing for a higher grant coupled with high marginal tax rates at the low end of the income distribution, taking advantage of the ability to use a nonlinear income tax with classifications and taking into account that, in practice, any category is likely to contain a range of different types.

Of greatest relevance are those who consider a nonlinear income tax with categorization. Diamond’s (1968) early essay discusses this case only in passing.²⁹ Immonen et al. (1998) focus on how the pattern of marginal rates over the entire income range differs across ability groups – with particular attention to whether marginal rates are rising or falling. They address this question using simulations and find that marginal rates are rising for the low-ability group and falling for the high-ability group. They attribute this result to differences in implicit revenue requirements on account of the fact that the high-ability group pays higher taxes overall to fund a subsidy to the low-ability group; this observation is supplemented by further simulations

²⁸Akerlof (1978), after examining the special model described in the text, states a more general version of the problem but does not analyze it. Bennett (1987) also explores lump-sum transfers between different types of individuals. Viard (2001) considers a linear income tax constrained to have the same rate for all groups but allowing for a group-specific transfer.

²⁹Diamond (1968, p. 296) offers the conjecture that the low-ability group might best be given a higher grant and higher marginal tax rates, the latter because distortion is not that great. However, as explained in subsection 3.1, when one considers the lower range of income, lower rates are likely to be optimal compared to the rates optimal in the standard problem and for the higher-ability group – although, as discussed in subsection 2.2, in the noncategorical case higher rates do tend to be optimal at the low end of the income distribution compared to higher levels of income for the reasons Diamond offers.

suggesting that lower (especially highly negative) revenue requirements tend to be associated with rising rather than falling rates. The discussion in subsection 3.1, however, suggests alternatively that differences in the shape of the ability distributions may be an important source of the different patterns in rates.³⁰

3.4. *Endogenous Categorization*

Offering more generous tax and transfer schedules to some groups than to others creates incentives to change one's category. If transfers are more generous when children (or a greater number of children) are present or if there is a single head of household, incentives to procreate and to marry will be affected. Likewise, preferential treatment of individuals who are disabled will produce moral hazard regarding the choice of occupation, levels of care, and efforts at rehabilitation. This consideration may favor reducing the degree of differentiation, to an extent that depends on the elasticity of the pertinent behavior. This implication, however, may not follow in all cases; notably, the proper social assessment of a change in the level of procreation is hardly uncontroversial.

If categorization is employed, it is also necessary to design the classification system itself. There are issues of burden of proof (the optimal tradeoff of type one and type two errors), optimal investments in accuracy, and developing methods (such as application fees or waiting times) to induce individuals to self-select at the application stage.³¹ Optimal classification and optimal treatment of those receiving a given classification are obviously interrelated problems.³²

³⁰This interpretation is reinforced by the informal, diagrammatic analysis in Dilnot, Kay, and Morris (1984, pp. 74-77), which is presented in Immonen et al. (1998). Immonen et al. state, however, that they do not consider such explanations because they use distributions of the same form for both ability groups (which, one may note, seems unlikely to characterize actual classifications based on criteria like disabilities that seek to segregate by type). Nevertheless, although their distributions have the same functional form, they have a different mean and variance. Most obviously, a distribution with a lower (higher) mean will tend to have a lower (higher) first term in expression (5) at the low end of the income distribution because $1 - F$ in the numerator is lower (higher) and f in the denominator is higher (lower). Also notable is their assumption that the low-ability group is the one with the higher variance; previous simulations performed by some of the same authors, Kanbur and Tuomala (1994), showed that greater dispersion in the distribution of abilities may result in optimal rates that rise rather than fall over much of the income distribution.

³¹Various of these and related issues have been explored, largely in somewhat different contexts, by Diamond and Sheshinski (1995), Kaplow (1998), and Parsons (1991, 1996).

³²Mirrlees (1990) raises the possibility that there are errors of income measurement, and suggests, by analogy to Stern's (1982) analysis of errors in classification in an ability-tax scheme, that larger errors may favor lower marginal income tax rates and that the extent of this adjustment may be greater the stronger the social preference for redistribution (because errors in which low-income individuals are misclassified as high-income are particularly socially costly).

4. Work Inducements

4.1. Preliminary Analysis

Over the history of welfare programs, a variety of work incentives have often been proposed and sometimes been implemented. Work inducements may be weak, such as when transfers are modestly reduced for individuals who fall short of targets, or they may be strong, such as when all transfers are forfeited for failing to work the required amount. The main motivations for rewarding work relate to the disincentive effects of welfare and the sense that those who are able to support themselves should do so. It might be supposed that work incentives can partially counter the distortionary effect of the tax and transfer system, at least at the lower end of the income distribution.

From the perspective of designing an optimal income transfer scheme, however, work inducements are puzzling in an important respect: If work incentives are tantamount to adjustments to the tax-transfer schedule for an identified group of workers – typically, those classified as able in various respects – why is there anything left to analyze? After all, the assessment in section 3 already determined, in principle, the optimal form of categorical assistance for able individuals as a function of the distribution of abilities in the relevant population. That reductions in work effort entail lost productivity and thus reduced tax revenue was taken fully into account. (The possibility of externalities to work is considered below.)

Furthermore, the optimal scheme outlined in section 3 does not appear to resemble a system with significant positive inducements to work. For the more able group, which is likely to be the sort of group that might be rewarded for higher work effort, two features were identified: First, a positive grant is included – at a level below, but not obviously greatly below, the level offered to those less able. Second, the optimal marginal tax rates on very low incomes are high, assuming that there is nontrivial classification error. To be sure, with no error a low grant, if any, is offered, followed by marginal rates of zero at the bottom, which does more closely resemble some work incentive schemes. Nevertheless, since nontrivial error is likely, the analysis to follow will continue to focus on that case.

Truly rewarding work (beyond market rewards) through the transfer system entails negative net marginal tax rates in the pertinent range. However, it is hard to see from the first-order condition (5) how such an approach could be optimal. Indeed, an important result on the optimal nonlinear income tax (which should extend to the case of classification, given the relationship between conditions (4) and (5)) is that optimal marginal tax rates are positive.³³ The intuition behind this result can be drawn from the perturbation used in subsection 2.2 to interpret condition (4): If the marginal tax rate in some range were negative, raising it would, through the marginal effect, reduce rather than increase distortion and, through the inframarginal effect, produce a redistributive benefit, so keeping the marginal rate below (or equal to) zero cannot be

³³This result holds except possibly at the endpoints of the distribution. See Mirrlees (1971), the generalization in Seade (1982), and the subsequent work of Ebert (1992).

optimal. Optimal marginal tax rates may be falling somewhat when incomes are, say, at the level one would earn working full time at the minimum wage or at the poverty line, but they are still likely to be high and the decline is likely to be modest and certainly will not entail a plunge below zero. Even though the character of the optimal tax and transfer schedule seems on its face to conflict with the optimality of standard work incentive schemes, it is useful to examine some specific forms to see more concretely when and why they are suboptimal.

4.2. Illustrations

Modest work inducements. – Consider a moderate work incentive program under which benefits are reduced by an earnings deficit, such as the extent to which actual earnings fall short of what one would earn working thirty hours per week at the minimum wage. Translated into the framework employed here, such a program is equivalent to a large reduction in g coupled with a modestly negative marginal tax rate (a net work subsidy) in the range from zero to the target level. When earnings are low, the marginal return to work has three components: one's wage, the penalty reduction (which equals, for example, 100% of one's wage rate if it equals the target wage rate), and the aggregate of taxes and phase-outs (which may be high, but are assumed for the moment to be less than 100% of one's earnings). Because the second component is taken to exceed the third, a net marginal earnings subsidy results.

If everyone subject to such a regime has the requisite ability such that in an optimal scheme they all would work at least at the target level, then the system is unproblematic. This conclusion is consistent with the previous observation that, with perfect classification, there should be no marginal tax on the bottom (able) individuals; this regime produces close to that result.

However, if there are classification errors – notably, if some subject to this work incentive scheme have a lower ability – then the previous analysis suggests that this scheme is not optimal. The very low implicit grant level may well be too low. Also, as explained in subsection 3.1, a higher grant would not be that costly if (contrary to the hypothesized scheme) marginal tax rates near the bottom were high rather than low. Furthermore, as stated in the preceding subsection, marginal rates should be high, possibly quite high, rather than negative near the target. Higher marginal tax rates near the target may, to be sure, reduce the work effort of lower-ability individuals among the able group, but the productivity and thus tax revenue loss is small whereas the inframarginal redistributive gain with respect to individuals having higher abilities is substantial. Additionally, it is quite unlikely to be optimal for the marginal tax rate to jump on the order of 100 percentage points at the target income level.

Strong rewards for work. – Next, consider a much stronger work inducement that eliminates all welfare benefits (or a significant portion thereof) if an individual's earnings fall even slightly below the target earnings level.³⁴ Once again, if classification were perfect and in

³⁴In considering this scheme, it is useful to keep in mind that due to positive and possibly high aggregate marginal tax (including phase-out) rates, the remaining benefit at the target level

the optimal scheme everyone would meet the target in any event, this sort of approach would be benign. However, with classification errors, this extreme approach would be worse than the aforementioned gradual penalty, a result implied by the analysis in subsection 3.1.

To further explore why such strong rewards are not optimal, consider a scheme that makes a work-reward grant (or a significant increment to a meager base-level grant) available only if one reaches a target income and applies a zero marginal tax rate on earnings up to that level. Could such a regime be optimal? To see why not, suppose that one reduced this work-reward grant slightly and increased the grant available at zero income in a revenue-neutral manner. (Note that one could not raise the base grant dollar for dollar because additional individuals, those with such low ability that they do not reach the target, would receive the base grant in addition to those who earned enough to receive the work-reward grant.) Three sets of effects from such a reduction in the work-reward grant may be identified, and all are favorable.

First, hold labor supply fixed and consider the purely redistributive effect on social welfare. Those receiving a greater grant are all poorer than those receiving lower grants, so this effect is positive. Second, regarding labor supply, very low-ability individuals who did not previously reach the target would work less due to the income effect of the higher base grant; this adjustment raises the actors' utility and has no effect on revenue. Additionally, some moderate-income individuals who were just at the target may reduce their labor supply due to the change in levels of the two grants; if they do, their utility will increase, and revenue will also increase since they forgo the remaining work-reward grant. Third, the income effect on those who previously earned above the target (they receive both grants, but the increase in the base grant is less than the reduction in the work-reward grant) results in greater labor supply, increasing their utility and also raising additional revenue.

To be sure, the resulting scheme is still not optimal. Nevertheless, the analysis to this point suffices to illustrate how a modification that reduces a strong work inducement is welfare improving. Further reductions would raise welfare even more, so the optimal scheme would not involve any work-reward grant at all.

4.3. Observability of Work Effort

Perfect observability. – The sorts of work inducements just examined tend not to be optimal despite the problem of labor supply distortion. It is, however, possible in principle to fashion effective work requirements if work effort itself – rather than just levels of earnings – can be observed. In the present formulation, if hours are observable, it is possible to infer the wage (ability) level, which is simply y/l , earnings divided by hours. This inference, however, is only possible for those who work a positive amount.

To analyze this case, begin with the stronger, simpler setting in which the wage rate can be observed for everyone (including those who do not work). Then one could fashion a work

of earnings may be modest.

requirement for each wage level that corresponds to the first-best level of labor supply. That is, one could withhold all assistance unless one worked at the first-best level and provide the first-best level of assistance at that level. For those for whom no work is optimal in the first best, an appropriate transfer could be paid.

Nevertheless, the observability of wage (ability) levels does not imply the desirability of work requirements per se. In fact, a work requirement is entirely unnecessary when the wage is fully observable because it then would be possible to implement an individualized lump-sum tax based purely on ability, which wholly avoids the labor-leisure distortion. Specifically, an individual with a very low wage would receive a significant grant and face a zero marginal tax rate on earnings: The grant would be set at the level that would be optimal assuming that the individual works the optimal amount, which he or she would do given the zero marginal tax rate. A somewhat higher-wage individual might pay a modest (lump-sum) tax and similarly face a zero marginal tax rate on earnings. Even higher-wage individuals would pay a higher tax, but still be subject to a zero marginal rate. In this setting, there is no reason (and it would be inefficient) to condition any grant or tax reduction on an individual's willingness to work more than he or she would otherwise find optimal.

Now, return to the case in which there is only conditional perfect observability, that is, ability is inferred from the hours and earnings of those who work a positive amount, but ability cannot be observed for those who do not work at all. In this setting, the first-best cannot be achieved because some who ideally would work would mimic lower types who do not work, qualifying for their level of transfer (but avoiding the need to exert effort). To combat this incentive problem, the rate of reduction in transfers (or increase in net taxes paid) as ability rises needs to be reduced. Net transfers to types that do work would still fall by type – and, after a point, net taxes paid would raise – and, as in the first best, zero marginal tax rates on the earnings of those who work would remain optimal. See Dasgupta and Hammond (1980).³⁵

This overall design is qualitatively different from work requirement schemes that indeed assume hours to be observable. Consider, for example, a regime that reduces the grant by an hours shortfall multiplied by a target wage rate, such as the minimum wage.³⁶ This is similar to the previously examined modest work inducement, except that the formula is based on an hours shortfall rather than an earnings shortfall. Specifically, let w° and l° denote the target wage and required labor supply and t the (assumed constant, for simplicity) preexisting aggregate (inclusive of phase-outs) marginal tax rate below the target income level. Then, for $l < l^\circ$, disposable income available for consumption, c , is

³⁵If individuals differed only by earning ability, as in Dasgupta and Hammond (1980), the optimal average tax rate is confiscatory as ability increases; that is, beyond the marginal type who just chooses to work, higher-ability individuals would pay a lump-sum tax that left them indifferent; all individuals have the same utility at the constrained optimum. If individuals also (unobservably) vary in their disutility of effort, an element of some of the writing on the extensive margin noted in subsection 4.4, this sharp result would be softened.

³⁶See, for example, Mead (2001), describing such a program in Wisconsin.

$$(6) \quad c = g + wl(1 - t) - w^\circ (l^\circ - l), \text{ or} \\ = [g - w^\circ l^\circ] + wl(1 - \tau),$$

where $\tau = t - w^\circ/w$. Note that, when $w = w^\circ$ and $t < 1$, $\tau < 0$. (Also observe that wage subsidies are somewhat similar.³⁷) For work above l° , ordinary income taxation is applicable, perhaps including phase-outs of remaining transfers, producing nontrivial and possibly high marginal tax rates. Likewise, although the specific version of the strong work incentive scheme examined in subsection 4.2 made the work-reward grant contingent on reaching a target level of income, there exist programs under which the target is in terms of hours – that is, eligibility is contingent on whether $l \geq l^\circ$.³⁸ However, as explained, if reliable information on hours is obtainable, better schemes could be devised – and, in particular, ones that do not reward work beyond market compensation.

Imperfect observability. – The primary inhibition to the use of the foregoing schemes is that ability cannot readily be inferred because hours are difficult to measure reliably and also are subject to manipulation. Regarding the latter, if rewards were contingent on higher reported hours of work, employers could raise hours and reduce the wage rate (such as by counting breaks, permitting individuals to start work early but not really do anything for awhile, allowing a more leisurely pace, failing to carefully monitor presence and effort), circumventing a purely hours-based scheme. Outside the formal sector’s large employers, potential circumvention could be far worse. Consider the self-employed, individuals working (even for employers) in their own homes (piecework, for example), those performing services in others’ homes or away from an employer’s premises, and many others.³⁹ If work restrictions were imposed on only certain job types where requirements were more readily enforceable, significant distortion in employment patterns would result, whereas if welfare eligibility (along with work requirements and rewards) were confined to this sector of the economy, opposite distortions would occur and many individuals may be left without needed support.

With a binding minimum wage – or a program rule that requires work to be rewarded by

³⁷Suppose that, conditional on $w < w^\circ$, $c = g + wl(1-t) + s(w^\circ - w)l$, where s is the subsidy rate. This can be rewritten as $c = g + wl(1-\tau)$, where $\tau = t - s((w^\circ - w)/w)$. Further, a wage subsidy is likely to be associated with a reduced grant, making this formulation even closer to that in (6). Differences include that w rather than l is assumed to be observable, the subsidy rate is chosen (though one could reduce the grant by a fraction of the earnings shortfall under the work requirement as well), and the ceiling applies to the wage rate rather than to an amount of hours worked.

³⁸See, for example, Blundell and Walker’s (2001) examination of the Working Families Tax Credit in the United Kingdom and Michalopoulos, Robins, and Card’s (2005) of an experimental program in Canada.

³⁹In the short run, such as in an income maintenance experiment, manipulation may not be that great, but in a permanent, comprehensive regime, one would eventually expect more substantial responses by workers and employers.

at least some base wage level w° to count – this problem is reduced, at least for individuals of sufficiently low ability. (Individuals for whom $w > w^\circ$ could still game the system, but only to the point that their nominal wage fell to w° .⁴⁰) However, moving toward this type of alternative makes the scheme depend more on earned income and less on hours per se. Moffitt (2002) suggests that the manipulability of wages and hours explains why earnings rather than wage subsidies are more common.

At that point, the preceding analysis of categorization based on a noisy signal, where within each category only income is observed, becomes applicable. Put another way, it may be that the best that can be implemented at reasonable administrative cost is a crude system that yields noisy signals of hours and ability. Then individuals can be categorized, but only imperfectly. The analysis in section 3, in turn, indicates how to design optimal transfer systems in such a setting.⁴¹

To be more precise, suppose that there are not merely two categories – such as very-low-ability individuals and everyone else – but a number of categories at the bottom, each corresponding to a narrow band of wage rates.⁴² Assume further that there is some noise (perhaps due in part to manipulations) in these observations so that, in each category, true wage rates are concentrated in a narrow range, but there are some misclassified individuals near the very bottom, some scattered in between the bottom and the target range, and some scattered above. In each such category, the optimal scheme will probably involve a nontrivial grant (because of those at or near the bottom), a steep marginal rate (still below 100%) in the next segment (because there are few individuals there and many above, making the marginal effect quite small relative to the inframarginal effect), a very low marginal rate where most individuals are concentrated (because the marginal effect is very high), and a hard-to-specify marginal rate beyond this range (both the marginal and inframarginal effects are quite small, so the optimal marginal rate will depend on the precise shape of the distribution, among other factors).

The weak work requirement scheme outlined in expression (6) differs. Instead of a high grant and a steep phase-out near the bottom, it has a low (or zero) grant and a very low, even negative marginal rate. As noted, this would only be optimal if classification were nearly perfect. That the marginal rate is very low near the target is optimal, because of the heavy concentration of individuals at that point in the distribution. However, at the target point there is a very large jump in the marginal rate, whereas it is optimal to continue a very low marginal rate (as long as the target is not above nearly everyone in the pertinent distribution). Thus, in most respects, this

⁴⁰The higher the target wage level, the less evasion, but the greater the number of individuals who cannot reach the target (or who suffer great disutility from the attempt).

⁴¹It is useful to keep in mind that the analysis is not only pertinent to discrete categorization but also to the case in which θ (which may be multi-dimensional) is continuous.

⁴²One might think of the (noisy) signal as the wage rate itself or hours (they are interchangeable if earnings are also observed, although one may also think of the case, especially with manipulation, in which hours are observed but are a poor signal of the wage rate because of unobserved variation in effort).

sort of work incentive regime is not optimal given the information assumed to be available.

The strong scheme, which provides a significant work-reward grant when reaching the target that is subsequently subject to a steep phase-out, also is not optimal under present assumptions about available information. The bottom end (supposed to be a reasonably generous grant and a steep phase-out) is roughly optimal. However the jump – involving an infinite negative marginal rate – is certainly not optimal. Likewise, the steep phase-out of the work-reward grant starting just at the target is inefficient: Assuming that the target is in the middle range of the cluster of individuals, that rate should be low, not high.⁴³ (Indeed, in optimal schemes, the marginal rate is likely to be falling just past the middle of the cluster since $1-F$ is falling rapidly. To see this point, compare the optimal marginal tax rate just to the left and right of the mode of the distribution, where f has the same value but $1-F$ is lower at the right. Typical work requirement schemes, by contrast, have the marginal rate rise steeply in this region.)

In sum, under a range of assumptions about available information – no observability of hours and wage rates, perfect observability, and imperfect observability – various work inducement regimes appear to be suboptimal. To know what is optimal, it is necessary to identify what information is actually available and then to be consistent in employing the generic analysis of the optimal tax and transfer scheme to a categorical system reflecting the available information.

4.4. Prior Literature

Previous investigations of work inducements take different approaches from that adopted here. Fortin, Truchon, and Beauséjour (1993) perform simulations and find that introduction of a strong work requirement may be welfare improving, but this is by reference to their benchmark of a negative income tax with a 100% grant phase-out rate, which is known not to be optimal. Their simulations suggest that other, less extreme negative income tax schemes with no work requirement are even better. Besley and Coate (1995) reach a contrary result in a model in which the objective is not welfare maximization but rather minimization of the cost of bringing all individuals up to a target level of consumption without regard to the utility they thereby achieve. They find that a form of workfare, under which the poorest individuals perform unproductive public service jobs with high disutility, is optimal. In essence, individuals are subject to a nonmonetary penalty if they wish to receive cash assistance rather than earning in the private sector; because this penalty eliminates the utility benefit of welfare, more able individuals are not induced to participate. When Besley and Coate instead consider a goal of providing minimum utility rather than consumption, such workfare is no longer optimal. Brett (1998) considers productive public service employment with a more standard objective function and finds that, if

⁴³Note that the large jump in assistance at the target point followed by a steep phase-out creates an incentive for individuals with incomes well above the target point to reduce work effort. This effect does not appear, for example, in Michalopoulos, Robins, and Card's (2005) examination of a Canadian experiment because, as a one-time unanticipated program, eligibility could be limited to individuals who in the preceding period had very low earnings.

public sector work is sufficiently productive, it may be optimal for low-ability types to be subject to public workfare requirements. The rationale is that tying benefits to public employment is selectively advantageous to low-ability individuals because they forgo less market income than do high-ability individuals when they switch from market to public employment.⁴⁴

An additional line of work, although not focused on work inducement schemes per se, is highly relevant to understanding the problem. The seminal paper is Diamond (1980), who analyzes a model in which labor supply decisions are exclusively at the extensive (participation) margin: Each individual decides whether or not to work full time, and those who work effectively reveal their type. (This revelation of type by workers was similarly present in Dasgupta and Hammond (1980), mentioned above.) More able types are assumed unable to mimic the less able by working part time or full time at lower effort. This framework has recently received increased attention and has been extended in Boone and Bovenberg (2004), Choné and Laroque (2005), Laroque (2005), and Saez (2002).

The core ideas can best be understood in the basic model. Raising the grant received by those who do not work reduces participation by high as well as low types because there is assumed to be preference heterogeneity, and therefore it has a relatively high revenue cost. By contrast, raising the work reward for low types is assumed not to cause higher types to reduce labor effort because of the assumption that more able individuals cannot earn less income by exerting less labor effort. Accordingly, it may be optimal for the net received by lower types who participate in the labor force to exceed the base grant by a sufficient amount that marginal tax rates are negative.⁴⁵ In some simulations by Saez (2002), which also allow for limited substitution on the intensive margin, the optimal scheme does indeed involve negative rates at the very bottom (the first few thousand dollars of income), with a quick rise to high positive rates thereafter. The policy relevance of these models depends on the degree to which workers are unable to adjust effort or hours in the long run and the presence of fixed costs to employment that render part-time work of all sorts insufficiently remunerative.⁴⁶

⁴⁴See also Drèze and Sen (1989), who discuss the self-selection benefits of public employment in disbursing famine relief in developing countries. Additional work that empirically examines existing or experimental programs is not primarily concerned with analyzing what sort of scheme would be optimal.

⁴⁵Choné and Laroque (2005) and Laroque (2005) find negative marginal rates possible if adjacent ability types' disutilities of work are sufficiently concentrated (and the government can observe this), a possibility that they show is ruled out by a regularity assumption on the distribution of disutilities of work. Boone and Bovenberg (2004) show that negative marginal rates may be optimal if the base grant is constrained to be too low because work subsidies to low-income individuals are then the only way to redistribute income.

⁴⁶Blundell and Walker (2001) present evidence for the United Kingdom that a substantial portion of childcare obtained by single parents involves the use of friends and relatives and that the marginal cost of childcare rises with hours of work as cheaper forms of childcare are exhausted first, suggesting that marginal costs of working may exceed rather than fall short of average costs for these individuals. Blundell et al. (1999) present UK data indicating that part-

4.5. Externalities to Work

Types of externalities – The analysis throughout this article assumes that labor effort has no special social significance beyond that already incorporated in the optimal income tax framework. Yet there may be other grounds for being concerned about the employment of lower-income individuals. Most notably, externalities to work may be present. First, there may be a psychological externality to other citizens. Taxpayers may derive utility (or avoid negative utility) from lower-income transfer recipients working or working harder. In some countries, including the United States, this seems to be a political fact. Whether it reflects an actual source of utility or merely conjectures about optimal policy that might be based on misunderstandings is a more complicated question. It may be that most individuals do not appreciate the subtle reasoning underlying the optimality of high marginal tax rates at the bottom of the income distribution. On the other hand, the analysis in section 2 suggested that, although optimal marginal tax rates at the bottom are high, they are probably not as high as actual aggregate (phase-out inclusive) marginal tax rates have in fact been in the past and even after welfare reforms including the expansion of the EITC. To that extent, popular perceptions may align with the teachings of optimal transfer analysis.

Additionally, this issue must be considered in the context of section 3's analysis of categorical assistance. It was suggested that (crudely consistent with actual practice) it may be optimal to divide low-income individuals into two (or more) groups, in the basic case corresponding to individuals with identifiably low ability and everyone else. A possibility is that certain individuals – notably, parents with young children – are now viewed as in the more able group, contrary to understandings in the past. In any case, it was seen that optimal within-group marginal rates for the more able group were very high at the bottom of the income distribution. However, under an optimal scheme this may not substantially discourage work because the optimal grant for that group is lower than for others.

These considerations are also pertinent to a second type of externality, those of the more conventional sort. It is often suggested, for example, that parents serve as role models for their children, so the failure to work perpetuates poverty, harming both the children (and future generations) and others, such as through increased crime. Additionally, individuals who work, having less free time, may themselves be less likely to engage in criminal activity. On the other hand, if parents, especially single parents, leave home for a greater period of time, care and supervision of children may suffer. The absolute and relative significance of these concerns and how they vary by family configuration, age of children, and income level are empirical questions about which little is known with confidence.⁴⁷ It might be supposed that, like with returns to

time employment at all levels is very common for both single mothers and women in couples.

⁴⁷Grogger and Karoly (2005) find that welfare reforms tended to raise children's well-being only when the reforms raised family incomes (through generous work incentives), while reforms that were more likely to reduce welfare dependency were more likely to be associated with negative effects on children. Levine and Zimmerman (2005) explore whether welfare spells adversely affect children and find that, with proper controls, there is no significant detrimental

investment in human capital, these effects are largely internal to the family (though effects on crime and the like obviously are not). However, this is not the case: Parents consider only the benefit to themselves (even if it is a highly altruistic benefit) and not the additional benefit to the children. Accordingly, if there are significant effects on children, one way or the other, they may justify departures from what the basic analysis would otherwise indicate to be optimal.⁴⁸

Work and related activities by the poor also might augment their human capital, leading to increases in effective earning ability in the future. If individuals are not myopic, however, they should already take this benefit into account. (To be sure, the benefits are discounted in light of positive marginal tax rates, but this distortion is no different in kind from the basic distortion in labor supply already considered.) Nevertheless, myopia might be a real concern; indeed, more myopic individuals who accordingly make insufficient investments in their own human capital may for that very reason be disproportionately concentrated at the lower end of the income distribution. It is familiar that myopia may be analyzed similarly to externalities, for one can view the problem as one of the present self imposing an externality on the future self.

Analysis of externalities – Supposing that there are net external effects of work of one sort or another, it is possible to modify the optimal nonlinear income tax model to take this into account. For concreteness, consider the case of net positive externalities. First, ceteris paribus, lower grants than otherwise would be optimal, on account of the income effect. Second, lower marginal tax rates at the bottom of the income distribution would be optimal.

To elaborate this latter point, reexamine the first-order conditions (4) and (5). The denominator of the first term in these expressions contains w , indicating that higher productivity of work is associated with lower marginal tax rates being optimal. In the standard analysis, productivity as measured by the wage rate reflects the marginal social product of additional work. If one postulates that there is an additional, external social benefit to lower-income individuals' work effort, it would be as if this productivity factor were higher. A higher denominator, in turn, indicates that the optimal marginal tax rate is lower.⁴⁹

Another way to view the point is to suppose that the government were to correct the

effect. Antel (1992), however, finds that a mother's welfare participation increases the likelihood of her daughter's subsequent participation. Duncan, Hill, and Hoffman (1988) survey earlier literature that provides mixed evidence, some of which suggests a degree of transmission of dependency, although determining causation is difficult.

⁴⁸See, for example, Kaplow (1995). This is one of the reasons for offering some welfare benefits in-kind, in ways that target children.

⁴⁹There are additional considerations. For example, another consequence of admitting income effects is that higher marginal rates on very-low-income individuals may be optimal because, like lower grants, they lead higher-income individuals, including low-ability individuals just above the marginal type, to work more. Additionally, higher true productivity implies higher revenue, which helps meet the revenue constraint, whereas a higher social return in terms of others' utilities does not.

externality with a Pigouvian tax. In the case of a positive externality, this would entail a subsidy, which could be implemented as a wage subsidy for the pertinent individuals. Such a subsidy would (ignoring complications such as multiple employers in the case of nonlinear schemes) be functionally equivalent to an adjustment to the income tax and transfer schedule.⁵⁰

The foregoing discussion indicates that externalities may influence the generosity of transfers and the level of marginal rates. However, it seems unlikely that the presence of externalities would change the qualitative results and, in particular, justify the forms of work inducement schemes considered previously in this section. Specifically, if there were, say, a notable positive externality to work, which could be reflected indirectly by substituting a higher social-product wage rate for the standard private-product wage rate in the first-order condition, the shape of the optimal rate schedule would not change radically. Large upward or downward jumps in marginal tax rates for modestly different earnings levels would remain difficult to justify. (One could suppose that there was a large positive externality to, say, the thirtieth hour of work per week but none to the thirty-first, but this supposition seems implausible.)

It is also interesting to reflect on the magnitude of externality necessary to warrant a substantial change in the level of optimal marginal tax rates at the bottom of the income distribution. Suppose, for example, that one wished to reduce marginal tax rates by ten percentage points on income between 0 and \$10,000 in the United States. For all individuals who earn in excess of \$10,000, the forgone revenue would be \$1000 (plus any further loss due to reduction in labor effort on account of the income effect). There are more than one hundred million such individuals, so the revenue cost would exceed one hundred billion dollars. Analyzing the welfare consequences of this shift is an involved task since minimizing the welfare cost requires re-optimizing the rest of the tax and transfer system, yet it seems clear that rather substantial positive externalities to work would be necessary to justify such a shift. Furthermore, it should be kept in mind that this revenue cost is associated with reducing rates only by ten percentage points and only on incomes below \$10,000.⁵¹

⁵⁰These observations also relate to the explanation in subsection 4.1 of why, in the standard optimal income taxation model, the optimal marginal rate must be positive (except possibly at the endpoints of the distribution). Part of that argument was that raising a negative marginal tax rate would, through the marginal effect, reduce distortion. If there is an externality to work, however, the nondistorting marginal tax rate is not zero, but negative. Therefore, once the marginal rate, even if negative, exceeds this level, raising the negative marginal tax rate from that point would increase rather than reduce distortion, just as does raising a positive marginal tax rate in the standard model.

⁵¹One way to fund the rate reduction would be to concentrate the necessary tax increase on individuals just above this range, say those in the \$10,000 to \$20,000 income band. This is similar to the method of phasing out EITC benefits. It should be kept in mind, however, that the phase-out only reduces the benefits to zero at the end of the phase-out range; it does not attempt in addition to raise the over \$30 billion necessary to fund the credits themselves, which if financed by individuals in the phase-out range would require substantially higher marginal tax rates. One effect of such adjustments, of course, would be to discourage work by the near-poor.

5. Conclusion

This article uses results from the literature on optimal nonlinear income taxation to illuminate the optimal design of income transfers. Initially, the standard model and first-order condition were interpreted with emphasis on the lower end of the income distribution. The resulting insights suggest modifications of the manner in which transfer programs are currently analyzed and avenues for reform. These conclusions were illustrated by reference to the EITC. Next, the model was adjusted to allow for categorical assistance, more precisely, different tax and transfer schedules for different groups in the population. Again, some of the results differ from conventional analysis and existing practice. Finally, the approach was used to illuminate common forms of work inducements, which were found not to be optimal.

The overall theme of this article is that questions involving the design of transfer programs are best analyzed by reference to what is already understood about optimal income taxation and what can be learned from extensions of that framework. There is a substantial gap between the existing literature on optimal nonlinear income taxation and that on the design of transfer programs. The former literature is abstract, highly technical, and at a level of generality that excludes many features of transfer programs and many characteristics of low-income individuals that may be important. The latter writing tends to be more concrete, less technical, and highly connected to the peculiarities of existing systems and reform proposals. Bridging the gap between these two literatures seems a priori to be a valuable enterprise. This article is meant to provide some specific insights but also to motivate further work along these lines.

There is a natural reluctance to tether the study of transfer programs to work on optimal nonlinear income taxation. Explicit analysis is difficult to understand, much less perform, there are few general analytical results, and the outcomes of simulations depend on elasticities, evidence on the distribution of skills, and the choice of a social welfare function, all of which are subject to controversy. These difficulties, however, can only be hidden, not overcome, by other approaches. The optimal design of transfer programs does depend on the pertinent empirical evidence and on how society trades off gains to some against losses to others. And the problem is analytically difficult; for example, setting the marginal tax rate at a particular, low level of income does affect everyone with higher incomes, and the effects on their utility and on the amount of revenue raised are important considerations, not details that can be set aside. Accordingly, there really does not seem to be a viable alternative to the suggested path, no matter how challenging that path proves to be.

References

- Acs, G., N. Coe, K. Watson, and R.I. Lerman. 1998. Does Work Pay? An Analysis of the Work Incentives Under TANF. Occasional Paper No. 9. The Urban Institute, Washington, D.C.
- Akerlof, G.A. 1978. The Economics of "Tagging" as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning. *American Economic Review* 68, no. 1, 8-19.
- Alm, J., and S. Wallace. 2000. Are the Rich Different? In Slemrod, J.B. (Ed.), *Does Atlas Shrug?: The Economic Consequences of Taxing the Rich*. Russell Sage Foundation, New York, pp. 165-187.
- Antel, J.J. 1992. The Intergenerational Transfer of Welfare Dependency: Some Statistical Evidence. *Review of Economics and Statistics* 74, 467-473.
- Atkinson, A.B. 1995. *Public Economics in Action: The Basic Income/Flat Tax Proposal*. Clarendon Press, Oxford.
- Atkinson, A.B., and J.E. Stiglitz. 1980. *Lectures on Public Economics*. McGraw-Hill Book Company, New York.
- Auerbach, A.J. and J.R. Hines, Jr. 2002. Taxation and Economic Efficiency. In Auerbach, A.J., Feldstein, M. (Eds.), *Handbook of Public Economics*, vol. 3. North-Holland, Amsterdam, pp. 1347-1421.
- Auten, G., and R. Carroll. 1994. Behavior of the Affluent and the 1986 Tax Reform Act. *National Tax Association Proceedings, Eighty-Seventh Annual Conference*, pp. 70-76.
- Benitez-Silva, H., M. Buchinsky, and J.P. Rust. 2004. How Large are the Classification Errors in the Social Security Disability Award Process? Working Paper No. W10219, National Bureau of Economic Research Working Paper Series. National Bureau of Economic Research, Cambridge, Mass.
- Bennett, J. 1987. The Second-Best Lump-Sum Taxation of Observable Characteristics. *Public Finance* 42, no. 2, 227-235.
- Besley, T., and S. Coate. 1995. The Design of Income Maintenance Programmes. *Review of Economic Studies* 62, no. 2, 187-221.
- Blundell, R., A. Duncan, J. McCrae, and C. Meghir. 1999. Evaluating In-Work Benefit Reform: The Working Families Tax Credit in the UK. (Preliminary Draft.)
- Blundell, R. And I. Walker. 2001. Working Families' Tax Credit: A Review of the Evidence, Issues and Prospects for Further Research. Inland Revenue Research Report No. 1.
- Boone, J., and L. Bovenberg. 2004. The Optimal Taxation of Unskilled Labor with Job Search and Social Assistance. *Journal of Public Economics* 88, no. 11, 2227-2258.
- Brett, C. 1998. Who Should Be on Workfare? The Use of Work Requirements as Part of an Optimal Tax Mix. *Oxford Economic Papers* 50, 607-622.
- Brito, D.L., and W.H. Oakland. 1977. Some Properties of the Optimal Income Tax. *International Economic Review* 18, no. 2, 407-423.
- Choné, P., and G. Laroque. 2005. Optimal Incentives for Labor Force Participation. *Journal of Public Economics* 89 nos. 2-3, 395-425.
- Dahan, M., and M. Strawczynski. 2000. Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates: Comment. *American Economic Review* 90, no. 3, 681-686.
- Dasgupta, P., and P. Hammond. 1980. Fully Progressive Taxation. *Journal of Public Economics*

- 13: 141-154.
- Diamond, P.A. 1968. Negative Taxes and the Poverty Problem—A Review Article. *National Tax Journal* 21, no. 3, 288-303.
- Diamond, P.A. 1980. Income Taxation with Fixed Hours of Work. *Journal of Public Economics* 13, 101-110.
- Diamond, P.A. 1998. Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates. *American Economic Review* 88, no. 1, 83-95.
- Diamond, P.A., and E. Sheshinski. 1995. Economic Aspects of Optimal Disability Benefits. *Journal of Public Economics* 57, 1-23.
- Dickert, S., S. Houser, and J.K. Scholz. 1994. Taxes and the Poor: A Microsimulation Study of Implicit and Explicit Taxes. *National Tax Journal* 47, 621-638.
- Dilnot, A., Kay, J. and Morris, N. 1984. *The Reform of Social Security*. Clarendon Press, Oxford.
- Drèze, J., and A. Sen. 1989. *Hunger and Public Action*. Clarendon Press, Oxford.
- Duncan, G., M. Hill and S. Hoffman. 1988. Welfare Dependence Within and Across Generations. *Science* 239, 467-471.
- Ebert, U. 1992. A Reexamination of the Optimal Nonlinear Income Tax. *Journal of Public Economics* 49, 47-73.
- Feldstein, M.F. 1999. Tax Avoidance and the Deadweight Loss of the Income Tax. *Review of Economics and Statistics* 81, 674-680.
- Fortin, B., M. Truchon, and L. Beauséjour. 1993. On Reforming the Welfare System: Workfare Meets the Negative Income Tax. *Journal of Public Economics* 51, 119-151.
- Garfinkel, I. (Ed.). 1982. *Income-Tested Transfer Programs: The Case for and Against*. New York: Academic Press.
- Giannarelli, L., and E. Steuerle. 1995. The Twice-Poverty Trap: Tax Rates Faced by AFDC Recipients. Working Paper. The Urban Institute, Washington, D.C.
- Gokhale, J., L.J. Kotlikoff, and A. Sluchynsky. 2002. Does It Pay to Work? Working Paper No. 9096, National Bureau of Economic Research Working Paper Series. National Bureau of Economic Research, Cambridge, Mass.
- Green, C. 1967. *Negative Taxes and the Poverty Problem*. Brookings Institution, Washington, D.C.
- Grogger, J., and L.A. Karoly. 2005. *Welfare Reform: Effects of a Decade of Change*. Cambridge, Mass.: Harvard University Press.
- Gruber, J., and E. Saez. 2002. The Elasticity of Taxable Income: Evidence and Implications. *Journal of Public Economics* 84, 1-32.
- Hepner, M., and Reed, W.R. 2004. The Effect of Welfare on Work and Marriage: A View from the States. *Cato Journal* 24, 349-370.
- Immonen, R., Kanbur, R., Keen, M., and Tuomala, M. 1998. Tagging and Taxing: The Optimal Use of Categorical and Income Information in Designing Tax/Transfer Schemes. *Economica* 65, no. 258, 179-192.
- Juhn, C., K.M. Murphy, and R.H. Topel. 1991. Unemployment, Nonemployment, and Wages: Why Has the Natural Rate Increased Through Time? *Brookings Papers on Economic Activity* 1991, no. 2, 75-142.
- Kanbur, R., and M. Tuomala. 1994. Inherent Inequality and the Optimal Graduation of Marginal Tax Rates. *Scandinavian Journal of Economics* 96, no. 2, 275-282.
- Kaplow, L. 1995. A Note on Subsidizing Gifts. *Journal of Public Economics* 58, 469-477.

- Kaplow, L. 1998. Accuracy, Complexity, and the Income Tax. *Journal of Law, Economics, and Organization* 14, 61-83.
- Kaplow, Louis. (Forthcoming.) *The Theory of Taxation and Public Economics*.
- Keane, M., and R. Moffitt. 1998. A Structural Model of Labor Supply and Multiple Welfare Program Participation. *International Economic Review* 39, no. 3, 553-589.
- Laroque, G. 2005. Income Maintenance and Labor Force Participation. *Econometrica* 73, 341-376.
- Levine, P.B., and Zimmerman, D.J. 2005. Children's Welfare Exposure and Subsequent Development. *Journal of Public Economics* 89: 31-56.
- Liebman, J.B. 2001. The Optimal Design of the Earned Income Tax Credit. In Meyer, B.D., Holtz-Eakin, D. (Eds.), *Making Work Pay: The Earned Income Tax Credit and Its Impact on America's Families*. Russell Sage Foundation, New York, pp. 196-233.
- Mead, L.M. 2001. Implementing Work Requirements in Wisconsin. Discussion Paper No. 1231-01, Institute for Research on Poverty.
- Meyer, B.D., and D. Holtz-Eakin (Eds.). *Making Work Pay: The Earned Income Tax Credit and Its Impact on America's Families*. Russell Sage Foundation, New York.
- Michalopoulos, C., Robins, P.K., and Card, D. 2005. When Financial Work Incentives Pay for Themselves: Evidence from a Randomized Social Experiment for Welfare Recipients. *Journal of Public Economics* 89: 5-29.
- Mirrlees, J.A. 1971. An Exploration in the Theory of Optimum Income Taxation. *Review of Economic Studies* 38, 175-208.
- Mirrlees, J.A. 1976. Optimal Tax Theory: A Synthesis. *Journal of Public Economics* 6, 327-358.
- Mirrlees, J.A. 1990. Taxing Uncertain Incomes. *Oxford Economic Papers* 42, 34-45.
- Moffitt, R.A. 2002. Welfare Programs and Labor Supply. In Auerbach, A.J., Feldstein, M. (Eds.), *Handbook of Public Economics, vol. 4*. North-Holland, Amsterdam, pp. 2393-2430.
- Moffitt, R.A. (Ed.). 2003. *Means-Tested Transfer Programs in the United States*. University of Chicago Press, Chicago.
- Moffitt, R.A., and M.O. Wilhelm. 2000. Taxation and the Labor Supply Decisions of the Affluent. In Slemrod, J.B. (Ed.), *Does Atlas Shrug?: The Economic Consequences of Taxing the Rich*. Russell Sage Foundation, New York, pp. 193-234.
- Parsons, D.O. 1991. Self-Screening in Targeted Public Transfer Programs. *Journal of Political Economy* 99, 859-876.
- Parsons, D.O. 1996. Imperfect "Tagging" in Social Insurance Programs. *Journal of Public Economics* 62, 183-207.
- Saez, E. 2001. Using Elasticities to Derive Optimal Income Tax Rates. *Review of Economic Studies* 68, 205-229.
- Saez, E. 2002. Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses. *Quarterly Journal of Economics* 117, 1039-1073.
- Sammartino, F., E. Toder, and E. Maag. 2002. Providing Federal Assistance for Low-Income Families Through the Tax System: A Primer. Discussion Paper No. 4. Urban-Brookings Tax Policy Center, Washington, D.C.
- Seade, J.K. 1977. On the Shape of Optimal Tax Schedules. *Journal of Public Economics* 7, 203-236.
- Seade, J. 1982. On the Sign of the Optimum Marginal Income Tax. *Review of Economic Studies* 49: 637-643.

- Shaviro, D.N. 1999. *Effective Marginal Tax Rates on Low-Income Households*. Employment Policies Institute, Washington, D.C.
- Slemrod, J., S. Yitzhaki, J. Mayshar, and M. Lundholm. 1994. The Optimal Two-Bracket Linear Income Tax. *Journal of Public Economics* 53, 269-290.
- Stern, N. 1982. Optimum Taxation with Errors in Administration. *Journal of Public Economics* 17, 181-211.
- Stiglitz, J.E. 1987. Pareto efficient and optimal taxation and the new new welfare economics. In Auerbach, A.J., Feldstein, M. (Eds.), *Handbook of Public Economics*, vol. 2. North-Holland, Amsterdam, pp. 991-1042.
- Tuomala, M. 1990. *Optimal Income Tax and Redistribution*. Clarendon Press, Oxford.
- Viard, A.D. 2001. Some Results on the Comparative Statics of Optimal Categorical Transfer Payments. *Public Finance Review* 90, 148-180.
- Wilson, P., and R. Cline. 1994. State Welfare Reform: Integrating Tax Credits and Income Transfers. *National Tax Journal* 47, 655-676.