

The Efficiency Case for Universality

Nicholas Rowe and Frances Woolley*

Revised 26 August 1998

Abstract

The paper compares universal social programs with targeted social benefits. We define "universality" as separability of the tax/transfer system in income and other non-monetary attributes, and introduce the concept of "parallel equity" - a requirement that like differences in needs should be treated alike. We develop a model based on optimal tax theory which shows that under specified assumptions universality is efficient as well as equitable; departures from universality result in an inefficient structure of net marginal tax rates. Our key assumptions are that individual-specific needs necessitate a fixed exogenous monetary expenditure, and that needs are exogenous, observable, and uncorrelated with other relevant variables. We then relax these assumptions and describe the limited circumstances under which departures from universality may be desirable.

The Efficiency Case for Universality

1. Introduction

This paper compares “universal” social programs, where benefits to individuals with special needs are paid to rich and poor alike, with targeted social programs, where benefits are targeted to the poor only. We develop a model consistent with optimal tax theory which shows that under specified assumptions universality is efficient as well as equitable. “Clawback”-- the reduction of benefits as income increases-- results in an inefficient structure of net marginal tax rates. If poor people get free eyeglasses and prescription drugs, then so should the rich.¹

The universality debate has been surprisingly separate from the more general tax design literature.² Arguments for universality generally focus around higher benefit take-up, the stigma of targeted benefits, administrative issues, political coalition-building, or collective responsibility for raising children. The arguments against universality are that clawing back benefits from the rich increases the aggregate progressivity of the tax system, and increases “target efficiency” by targeting benefits on the poor.³ Our paper does not address other arguments in favour of universality, but counters both arguments against it.

Our paper is related to Akerlof’s (1978) paper on “tagging”. Akerlof argues that a simple tax/transfer system like the negative income tax, where everybody gets the same basic transfer and then pays tax as a function of income only, essentially throws away information on attributes which signal earning capacity. We agree with Akerlof that information on the non-income attributes of individuals should be incorporated into the tax/transfer system. But the attributes Akerlof looks at are indicators of an individual’s wage. The attributes we consider are indicators of special needs,

that is, the extra expenditures that some individuals need to give them the same standard of living and utility as others.

Our theoretical argument is simple; we merely apply an argument that is standard in other contexts within public finance and even macroeconomics. The total deadweight costs of taxation are minimized when the marginal deadweight cost of tax revenues is equalized across revenue sources. This in turn implies that marginal tax rates should be equalized across revenue sources, (unless the marginal deadweight cost functions differ because of, say, different supply or demand elasticities.) Different groups of people, such as those who do and who do not need eyeglasses, should face the same marginal tax rate, unless the groups differ in their behavioural responses to taxation. But if benefits are not universal, there will be higher effective marginal tax rates for the shortsighted than for the normally-sighted, because when a shortsighted person earns an extra dollar of income, she not only pays higher income tax, but also receives lower benefits.

One of the contributions of our paper is to offer a formal definition of “universality” and to introduce the concept of “parallel equity”. We do this in section 2 below. In section 3, we develop a formal model to show that, under certain circumstances, the most efficient and equitable means of addressing differing needs across the population is through a universal system of benefits. In sections 4 through 7 we show that modifying the assumptions in our model could justify departures from universality, but that universality remains a useful benchmark model, for these departures are as likely *a priori* to require paying greater, instead of lesser, benefits to the “needy” rich.

2. Defining Universality

Consider an individual i with income Y_i and having a vector of other observable⁴ (to the government) non-monetary attributes A_i indicating that individual’s special needs. For example, an

attribute could be the number of children dependent on the individual, presence of diabetes and hence the need for insulin, or shortsightedness and hence the need for glasses.

A tax/transfer system is a mapping from an individual's market income and attributes into a disposable income, $D_i = F(Y_i, A_i)$, for that individual. A tax/transfer system with "universality" takes the form $F(Y_i, A_i) = Y_i + B(A_i) - T(Y_i)$, where $T(Y_i)$ are taxes paid, and $B(A_i)$ are benefits received. *Universality means the tax/transfer function is additively separable in income and other non-monetary attributes.*


The crucial feature of universality, as we have defined it, is that individuals with the same market income face the same marginal tax rate, regardless of their other attributes (though their average tax rates will differ).⁵ We believe that our definition captures one important aspect of what is commonly meant by "universality", namely that the same benefits are paid to rich and poor alike.⁶

Universality can be visualized by drawing curves showing disposable income as a function of market income for people with different levels of need. Under universality the curve representing each level of need would simply be a vertical displacement of any other curve. All curves would have the same slope, and hence the same marginal tax rate, at any given level of income. Figure 1 shows a universal tax/transfer system with a progressive income tax (the solid lines), and a tax/transfer system where benefits are "clawed back" or reduced as income increases (the dashed line).

Universality is satisfied under a "demogrant," a transfer of a given amount to all members of a particular group in the population. Canadian old age security and family allowance programs were once universal.⁷ Provincial health insurance plans and primary and secondary education fit our definition of universality, in that they completely offset the cost of specific needs (the need for

health care and the need for education), regardless of income. Non-refundable tax credits, such as the disability amount, are not universal because they provide no benefits to people with very low taxable incomes. Deductions for specific needs, such as the deduction for dependants in the US income tax system or the Canadian northern residents deduction, are not universal either. Given a progressive rate structure, they are worth more to individuals who pay a higher marginal tax rate. In terms of Figure 1, deductions shift the disposable income curve horizontally rather than vertically. Refundable tax credits such as the child tax benefit are not universal because benefits are reduced or “clawed-back” as income increases.

The demand for universality is not a demand for vertical equity (the rich should pay more tax), nor horizontal equity (similar people should pay similar tax), but is instead a demand for



"parallel equity"-- a demand that like *differences* be treated alike. Suppose a diabetic person who needs insulin needs \$100 more disposable income to achieve the same standard of living as if he were not diabetic. According to parallel equity, if a poor diabetic deserves \$100 in compensation for his insulin needs, then a rich person deserves the same level of compensation.

Our argument is that universality is efficient as well as equitable. If a diabetic with insulin behaves in exactly the same way as a person without diabetes, so that for example he has the same labour supply elasticity, then both should face the same marginal tax rate if deadweight costs of taxation are to be minimized. Needing to spend \$100 on insulin or glasses is just like a capricious lump-sum tax imposed on some individuals by nature, which the government should correct by paying those individuals an extra \$100 lump-sum transfer. One way of understanding why these transfers are efficient is to note the parallel with an insurance market. Transfers from the normally-sighted to those who need glasses provide perfect insurance against defective vision through a universal social program administered by the government.⁸

3.0 Formal Models

In this section we demonstrate that universality is optimal under specified assumptions about people's preferences and the distribution of needs. The first part of the section takes a four-person example, with a utilitarian social welfare function. These restrictive assumptions are made for simplicity. The second part shows how individual-specific needs can be introduced into the canonical optimal tax problem. We allow both needs and income to vary continuously, and use a more general social welfare function. In later sections we relax the assumptions of the formal models, in order to establish the (limited) circumstances under which deviations from universality may be efficient.

3.1 A Simple Four-Person Example

Consider first a simple example with a utilitarian social welfare function and with only four types of people. People differ by need, N_i , and are either high-need, N_H , or low-need, N_L . Need is a perfectly observable exogenous attribute. People also differ by wage, W_j , $j=P,R$. The low-wage, W_P , are “poor”, the high-wage, W_R , are “rich”. Wage is exogenous, but not observable by the government. There are thus four types of people: low-need poor (LP), high-need poor (HP), low-need rich (LR), and high-need rich (HR). There is an equal number of individuals of each type in the economy, so wage and need are independently distributed across the population. The government knows the population distribution of wages. It can observe income, Y , as well as needs, but cannot observe the underlying determinants of income -- wage and labour supply.

All individuals have identical utility functions of the form:

$$U^{ij} = U(C_{ij} - N_i, L_{ij}), \quad i=L,H; j=P,R \quad (1)$$

where C_{ij} is consumption expenditure, N_i is individual-specific need, and L_{ij} is labour. A ‘need’ is thus modelled as a fixed additional expenditure that some individuals must make to give them the same standard of living and utility as other individuals without that special need. N_i could thus represent the cost of insulin for a diabetic, or the cost of glasses for a short-sighted person. We assume the utility function is concave, to ensure a unique equilibrium to the individual’s choice problem, and we assume diminishing marginal utility of consumption, to motivate governmental transfers from rich to poor and from low-need to high-need.

There is no saving, and so the individual budget constraint is given by:

$$C_{ij} = W_j L_{ij} - T_{ij} \quad (2)$$

Consumption is labour income minus taxes. Each individual pays taxes (net of transfers) T_{ij} , contingent upon his type ij . If $T_{ij} < 0$, the individual receives a transfer.

An individual's income affects the amount of tax he pays only insofar as his income (in equilibrium) reveals whether he is a rich type or a poor type. In this context, universality can be defined as a tax function in which the difference in taxes between rich and poor is independent of need, that is, one which satisfies $T_{HP} - T_{LP} = T_{HR} - T_{LR}$.

The government chooses taxes to maximize a utilitarian social welfare function:

$$S = \sum_{i=L,H; j=P,R} U^{ij} \quad (3)$$

subject to the government budget constraint

$$\sum T_{ij} \geq 0 \quad (4)$$

With perfect information, the government would transfer money from the rich to the poor, and also from the low-need to the high-need, in order to equalize the marginal utility of consumption -- or, equivalently, $C_{ij} - N_i$ -- across all individuals. With unobserved wages, however, the rich types may reduce their labour supply and masquerade as the poor types in order to get the same consumption with less labour. To masquerade as poor, a type LR would work $W_P L_{LP} / W_R$ hours (the income of a LP divided by the wage of an LR), and an HR would work $W_P L_{HP} / W_R$. This creates incentive-compatibility constraints on the government's problem, as given by:

$$U^{iR} \geq U^{iR/P} \quad i=L, H \quad (5)$$

where $U^{iR/P} = U(W_P L_{iP} - T_{iP} - N_i, L_{iP} W_P / W_R)$ represents the utility enjoyed by a rich individual masquerading as a poor individual with the same level of needs. Equation (5) requires that neither rich type is better off masquerading a poor type.

There are no incentive-compatibility constraints arising from differences in need. Low-need types cannot masquerade as high-need, because we assume (for now) that the government can observe need directly. Nor is there a moral hazard problem, since an individual's need is assumed exogenous. The government, therefore, will be able to equalize marginal utility of consumption between low-need and high-need types at a given wage. Indeed this observation drives our result that the optimal tax system must be universalistic.

We solve the problem as a two-stage game. In the first stage the government chooses an optimal tax policy (which is a tax/transfer for each of the four types), anticipating that individuals will react optimally in the second stage. In the second stage, individuals observe choose their labour supplies, L_{ij} , taking the government's choice of tax policy as given. Our conjectured solution is a subgame-perfect equilibrium, provided the government is not allowed to renege on its announced tax policy after the second stage.

The solution to the second stage of the game is defined by the individual's first order conditions. These are found by maximizing utility (1), subject to the individual budget constraint (2) :

$$U_1^{ij} W_j = U_2^{ij} \quad i=L,H \quad j=P,R \quad (6)$$

In the first stage the government chooses taxes and transfers to maximize the social welfare function (3), subject to the government budget constraint (4) and the incentive compatibility constraints (5). The Lagrangean for this problem is given by

$$\mathcal{Q} = S + \lambda \sum T_{ij} + \mu_L (U^{LR} - U^{LR/P}) + \mu_H (U^{HR} - U^{HR/P})$$

Differentiating the Lagrangean gives the first order conditions

$$\frac{\partial \mathcal{L}}{\partial T_{iP}} = -U_1^{iP} + \lambda + \mu_i U_1^{iR/P} = 0 \quad i=L, H \quad (7)$$

$$\frac{\partial \mathcal{L}}{\partial T_{iR}} = -U_i^{iR} + \lambda - \mu_i U_1^{iR} = 0 \quad i=L, H \quad (8)$$

The first term in equations 7 and 8 is the marginal utility of consumption. It is negative because higher taxes reduce utility. The second term represents the benefits of additional public funds. The third term represents the effect of changing taxes on the incentive compatibility constraints. This term is positive in equation 7, representing the fact that increasing taxes on the poor slackens the constraint, because the rich then have less incentive to masquerade as poor. It is negative in equation 8, because increasing taxes on the rich tightens the constraint, because the rich then have more incentive to masquerade as poor.

We conjecture that the solution to the government's problem requires:

$$T_{Lj} + N_L = T_{Hj} + N_H \quad j=P, R \quad (9)$$

For both poor (low wage) and rich (high wage), the difference in taxes just compensates for their difference in needs. This conjectured solution exhibits universality because the difference in taxes between low-need and high-need individuals is the same whether they are rich or poor.

We begin by showing that, with universality, labour supply and the marginal utility of consumption are independent of individual needs. With universality (9), for the individual first order conditions (6) to be satisfied, we must have $L_{LP} = L_{HP}$ and $L_{LR} = L_{HR}$. This is because the income

effect of high need on labour supply is exactly offset by the lower taxes of our conjectured solution, and so labour supply varies only according to an individual's wage, not according to need. From the individual budget constraints (2), net consumption $C_{ij} - N_i$ will also be the same for any two individuals with the same wage, regardless of whether they differ in need. It follows that $U^{Lj} = U^{Hj}$, $U^{LR/P} = U^{HR/P}$, and $U_1^{Lj} = U_1^{Hj}$ -- utility levels and the marginal utility of consumption are the same for high-need and low-need types (but are different for rich and poor).

We now construct an optimal universalistic tax system by choosing taxes for the low-need types T_{LP} and T_{LR} to satisfy (with equality) the incentive compatibility constraint for the low-need types (5 with $i=L$), and the government first order conditions (7 and 8 with $i=L$). Universality (9) implies that taxes for the high-need types must be

$$T_{HR} = T_{LR} - (N_H - N_L) \text{ and } T_{HP} = T_{LP} - (N_H - N_L).$$

However with universality, as argued above, utility and marginal utility is the same for both low- and high-need types. Therefore the incentive compatibility constraint for the high-need types (5 with $i=H$) and the government first order conditions (7 and 8 with $i=H$) will also be satisfied, with $\mu_2 = \mu_1$. Intuitively, we have broken down the system of equations which defines the optimal tax system into two symmetric sets, one for low-need types and one for high-need types, linked only through the government budget constraint (4) and its Lagrange multiplier λ . Under universality, where tax differences between high- and low-need types just compensate for differences in need, a tax differential between rich and poor which satisfies the set of equations for the low-need must also satisfy the set of equations for the high-need. Indeed the two sets of equations become identical. We have thus proven that if a (unique) solution to the optimal tax problem exists, that solution must be universalistic.

To understand the intuition behind this result, suppose that, instead of universality, the government chose to give rich types a smaller need-contingent tax difference than poor types, so that $T_{LR} - T_{HR} < T_{LP} - T_{HP}$. This implies that $T_{HR} - T_{HP} < T_{LR} - T_{LP}$, so those with high need gain more from masquerading as poor than those with low need. A tax differential which just prevented those with high need from masquerading as poor would, therefore, be more than sufficient to stop those with low need from masquerading as poor (the incentive compatibility constraint 5, $i=L$ would be slack). But this could not be an optimal tax system, for the government could take money away from the low-need rich (raise T_{LR}) and give it to the low-need poor (lower T_{LP}), who have a higher marginal utility of consumption, without violating the incentive compatibility condition.

3.2: A Reformulation in terms of the Canonical Optimal Tax Problem

In the previous simple example the social welfare function was utilitarian, and need and wage took on only two values, high and low. In the canonical optimal tax problem (Mirrlees 1971) the social welfare function is the more general Bergson-Samuelson, and individuals' wages varying according to a continuous probability distribution, but all individuals are otherwise identical -- there are no needs. The government faces a trade-off between efficiency and equity because it cannot observe individuals' wages, only their income. A higher marginal tax rate improves equity but worsens efficiency, because it distorts individuals' labour supply decisions. In this section we reformulate the canonical optimal tax problem to incorporate varying needs, and give a heuristic argument that universality is optimal.

Individuals vary by need N and wage W . Need and wage vary continuously, according to the density functions $g(N)$ and $f(W)$. We assume, as in section 3.1, that need and wage are independently distributed. The trick to converting our problem (where individuals differ in two

dimensions, need and wage) into the standard canonical optimal tax problem (where individuals differ only in one dimension, wage) is to introduce two new variables. We define “net consumption” as $C_i^* = C_i - N_i$, and “gross taxes” as $T_i^* = T_i + N_i$. Net consumption represents consumption *excluding expenditures on special needs like insulin or eyeglasses*, and gross taxes represents taxes *plus the taxes imposed by nature on the diabetic and shortsighted*.

The individual utility function (1) can then be rewritten in terms of the new variables as:

$$U_i = U(C_i^*, L_i) \quad (10)$$

The individual chooses net consumption C_i^* and labour supply L_i to maximize utility (10) subject to the budget constraint

$$C_i^* = W_i L_i - T_i^* . \quad (11)$$

(The budget constraint (11) is obtained by subtracting N from both sides of equation 2).

The government chooses the gross tax/transfer function $T_i^* = F(Y_i)$ which maximizes the Bergson-Samuelson social welfare function

$$S = \int_{\underline{W}}^{\overline{W}} \Psi(U_i) f(W) dW \quad \Psi' > 0 \quad (12)$$

subject to the individuals' optimization problem, and subject to the government budget constraint:

$$G + \int_{\underline{N}}^{\overline{N}} N_i g(N) dN = \int_{\underline{W}}^{\overline{W}} T_i^* f(W) dW \quad (13)$$

The problem defined by equations (10) through (13) is formally identical to the canonical optimal tax problem, except for the addition to government revenue needs caused by the second term in the

government budget constraint, which represents the “net neediness” of the population. And this term can be eliminated, without loss of generality, if we redefine the individual-specific needs N_i as the difference between an individual’s need and the population mean need, so that “net neediness” is zero. We have thus converted our two-dimensional problem into one which is mathematically identical to the canonical one-dimensional optimal tax problem.

The solution to equations (10) through (13) is analyzed rigorously in Mirlees (1971), and we will not reproduce it here. For our purposes, all that is required is one result: the solution to (10) through (13) defines a gross taxes function of the form

$$T_i^* = F(Y_i)$$

But $T_i^* = T_i + N_i$. Therefore taxes must be set according to $T_i = F(Y_i) - N_i$. Observe that this solution exhibits universality, because this tax/transfer function is additively separable in income and need.

Now it could be argued that by writing gross taxes as $T_i^* = F(Y_i)$ we ruled out any possibility other than universality. But suppose we allowed the government to condition gross taxes on need as well as on income, that is, $T_i^* = F(Y_i, N_i)$. Would the chosen gross tax schedule change? No. As long as wage and need are independently distributed, observing need gives the government no information on the individual’s wage, which is the only unobserved variable in the canonical problem. If it were optimal to have anything other than universality, then it would also be optimal to make purely random transfers, regardless of income, in the canonical case where all individuals are identical except for wage. An optimal tax theorist who has solved the canonical problem for the optimal gross tax function $T_i^* = F(Y_i)$, and then is informed of each individual’s N_i , sees the latter as irrelevant information, like learning the individuals’ names, and sees no reason to recalculate the optimal gross tax function $F(Y_i)$.

To make the same point another way, introducing individual-specific needs into the canonical optimal tax problem is exactly like supposing nature makes random, observable, lump-sum transfers between individuals. The optimal tax theorist would simply reverse those transfers to restore the original optimal tax system. But this is universality.⁹

3.3 The Costs of Deviations from Universality

As noted in section 2, deviations from universality take a number of forms: reduction of benefits to higher income individuals, use of non-refundable tax credits, and use of deductions instead of credits. All of these deviations have both equity and efficiency costs, which we consider verbally in this section.

We consider first the equity costs of non-universal benefits. Suppose for simplicity the government's Social Welfare Function is utilitarian, as in equation (3). Suppose a benefit program excludes certain income groups, say the rich. High-need rich will have a lower consumption net of needs, $C-N$, and a higher marginal utility of consumption than low-need rich. Total utility of the population is reduced by this violation of the equimarginal principal. Deviations from parallel equity are thus costly on distributional grounds -- like differences in needs should be treated alike.

The second cost of deviations from universality is an efficiency cost. Income taxes create a tax wedge between the marginal product of labour and the marginal rate of substitution of consumption for leisure, and so distort the individual's choice towards leisure rather than consumption. Deviating from universality by, say, reducing benefits paid to higher income individuals effectively increases the marginal tax rate on high-need people while reducing the marginal tax rate on low-need people. The tax wedge is increased for the high-need, and reduced for the low-need. So long as both have the same Hicks-compensated labour supply curve (labour

demand is the same because we assume N_i and W_i are independent), the deadweight cost triangle for the high-need has increased by more than the deadweight cost triangle for the low-need has been reduced. Optimal tax theory says that to minimize the deadweight cost of raising a given revenue, sources with the same supply and demand elasticities should face the same marginal tax rate. Deviations from universality are inefficient.

This argument that universality is efficient as well as equitable contradicts the popular argument that universality violates "target efficiency" -- scarce government dollars should be paid to the neediest, and cannot be wasted on paying benefits to the rich. But the concept of "target efficiency" rests on a fallacy; it assumes there is an exogenously given amount of tax revenue and therefore all expenditures must be evaluated at the same opportunity cost. But the deadweight costs of tax revenue are a direct result of the higher *marginal* tax rates required to collect those extra revenues. If an expenditure, such as maintaining child benefits across income classes, reduces those net marginal tax rates, the deadweight costs are negative.

What is truly at issue in the universality debate is not the overall progressivity of the tax system, or taking from the poor to give to the high-need rich, but taking from the low-need rich to give to the high-need rich. Moving towards universality means lowering net marginal tax rates for, say, the shortsighted, and raising net marginal tax rates for the normally sighted to pay for it. This promotes parallel equity by transferring from normally-sighted rich to shortsighted rich, and promotes efficiency by equalizing marginal tax rates across the two groups. There is no trade-off between parallel equity and efficiency.

4. Non-random distribution of needs

So far we have assumed that needs N_i are distributed randomly across the population. In this

section we consider non-random distribution of needs.

A positive correlation of needs and wages motivates paying for less than the full cost of needs *for rich and poor alike*. If shortsighted people tend to have a higher earning capacity, low income shortsighted people are probably just consuming more leisure than normally sighted people at the same income level. For needs which tend to be negatively correlated with earning capacity, disabling illnesses for example, the government should pay benefits that exceed the costs, for rich and poor alike. Needs then act as a way of tagging the low-waged, as in Akerlof (1978). Again it does not motivate paying a smaller benefit for those with high observed income than for those with low observed income.

What *would* motivate departures from universality is a correlation between needs and labour supply elasticities. If high-need people have a less elastic labour supply than the low-need then, according to the inverse elasticity rule, efficiency is increased when they face a higher marginal tax rate. This can be implemented by paying higher benefits to the poor than to the rich.

In summary, non-random distribution of needs can justify deviations from universality, but only under special, arguably convoluted, circumstances.

5. Imperfectly observed needs

So far we have assumed that the government can observe the needs N_i perfectly. Here we suppose instead that observable non-monetary attributes A_i are imperfectly correlated with the true underlying needs N_i . We consider two types of imperfectly observed needs. First, needs which are completely unobservable and, second, needs which are observable, but subject to measurement error.

To model unobserved needs suppose there are two needs, an observable need A_i and an unobservable need E_i , which together give total need as $N_i = A_i + E_i$. If the unobserved need E_i were

negatively correlated with income, then expected need, conditional on observing A_i , would be less for the rich than for the poor, and clawback of benefits at higher incomes would be justified. However reversing this assumption would justify “*clawforward*”, i.e. paying higher differential benefits to the rich.

It could be argued that it is not fair to ask the poor who have unsubsidized unobservable needs to pay higher taxes to provide universal benefits to the rich who have the observable need. However universality does not necessarily require higher taxes on the poor. What it requires is that high-need and low-need face the same marginal tax rates, which could also be satisfied by higher taxes for the low-need rich.

Imperfectly observed needs can also be modelled as $A_i = N_i + e_i$, where e_i is a random measurement error with zero mean and uncorrelated with N_i . Under certainty, a utilitarian government maximizes social welfare by equalizing marginal utility of consumption across individuals with different needs. Under uncertainty, the government's first order condition instead requires it to equalize *expected* marginal utilities, that is, set $EU'(C_i - N_i) = EU'(C_j - N_j)$ for all individuals i, j with the same wage. This in turn requires it to equalize expected net consumption by setting $E(C_i - N_i) = E(C_j - N_j)$, if and only if the marginal utility function is linear, that is $U'' = k$ and $U''' = 0$ where k is a constant. With a linear marginal utility function then, universality is reestablished in a certainty-equivalence form. The government should pay differential benefits equal to the *expected* needs of individuals, given its available information on attributes A_i .

Paying differential benefits equal to the *expected* needs of individuals may well entail partially discounting differences in observed A_i , because the A_i provide a noisy signal of true N_i . Using a standard signal processing result we have

$$E(N_i|A_i) = \frac{\text{var}(N)}{\text{var}(N) + \text{var}(e)} A_i + \frac{\text{var}(e)}{\text{var}(N) + \text{var}(e)} \bar{N}$$

suggesting that the larger the variance in the unobserved component, $\text{var}(e_i)$, the more observed needs should be discounted. However uncertainty *per se* gives no reason for clawback, which means discounting differences in observed A_i *more* for the rich than for the poor. This would only be justified if $\text{var}(e)$ was positively correlated with income as would be the case if, for example, medical expenditures were a less reliable signal of medical need at higher income levels. In this case paying benefits equal to expected N_i would require that differences in observed A_i be discounted more for the rich than for the poor. However in this example the departure from universality is only apparent, for differential benefits continue to reflect differential *expected* needs.

What could justify *true* clawback, defined as differential benefits for given differential expected needs being less for the rich than for the poor, would be the combination of both heteroscedastic uncertainty and a non-linear marginal utility function. Specifically, suppose the marginal utility function were locally linear for the rich but concave for the poor. By Jensen's inequality, if $U'(\cdot)$ is concave, then setting $EU'(C_i - N_i) = EU'(C_j - N_j)$ requires setting $E(C_i - N_i) > E(C_j - N_j)$ if N_i is uncertain but N_j is certain. That is, expected net consumption must be set higher, or benefits paid must exceed expected needs, when needs are uncertain. This is because the cost of underpaying benefits to an individual whose true needs are greater than expected exceeds the cost of overpaying someone whose true needs are less than expected.¹⁰ And if the uncertainty of needs is higher for those with high expected need (heteroscedasticity), we get an increased *differential* payment for the needy poor (where uncertainty is high and high uncertainty warrants high transfers), relative to the

differential payment to the needy rich (where uncertainty is equally high but does not warrant high transfers, because of a less concave marginal utility function).

Suppose for example that people with dependent children have higher needs than those with no children, but that the needs of those with children are uncertain, because some kids are more costly than others. Suppose also that the marginal utility function is non-linear for the poor, because falling below a given "poverty line" causes utility to start decreasing quite sharply. The poor should then receive child benefits greater than the expected cost of raising children, as a precaution against their falling below the poverty line if their kids happen to be unexpectedly costly. In contrast the rich, with linear marginal utility functions, should receive child benefits equal to the expected cost of raising children. This is a departure from universality, and would require partial clawback of child benefits from the rich, but a total clawback cannot be justified in this way.

What this example shows is that, while it is possible to think up a story to justify clawback, the assumptions needed are highly contrived. Universality should still be seen as the benchmark case.

6. Moral Hazard

Many needs are partially endogenous. One hundred percent compensation for needs may cause people to take fewer precautions against becoming needy, a problem known as "moral hazard."

The standard response to moral hazard is to provide less than full insurance, forcing people to bear some of the risk associated with their own actions. Generally, more insurance is provided where risk aversion is greater, and less insurance is provided for events which individuals can readily influence. Only a few elements of Canada's non-universal tax/benefit structure can be seen as a response to moral hazard. The limitation of the medical expenses tax credit to expenses over or \$1614 (in 1997) or three percent of income could be viewed as a "deductible" deterring

unnecessary medical expenditures. However several provincial health plans, including Ontario's, cover 100 percent of the cost of visiting a doctor about, say, an ear infection, but none of the cost of the antibiotics prescribed to cure that infection (except for some population groups, such as social assistance recipients). This provides the unfortunate combination of no protection for governments against moral hazard in the use of medical services, and no insurance for individuals against the cost of pharmaceuticals (for people below the threshold for the medical expenses tax credit).

Other limitations on universality are also difficult to explain using moral hazard arguments. The clawback of old age security from higher income individuals cannot be a response to moral hazard, since a person's eligibility for old age security is based on date of birth, which cannot be changed. The moral hazard associated with child benefits, the incentive to have more children, is not severe. Phipps (1995: 26) estimates that child benefits offset less than 5 percent of the cost of children, while US studies suggest that benefits have little effect on birth rates (Moffitt, 1992). Moreover, moral hazard ceases to be a "problem" if birth rates are seen as too low. Indeed, child benefits may be used as part of pro-natalist policies, as has been done in Quebec.

The clawback of child benefits from better-off families may be efficient if moral hazard is more severe, or the gains to insurance due to risk aversion smaller, at high income levels than at low income levels. Yet even if this is the case, clawback of benefits itself creates moral hazard. When greater benefits go to low income families, families have an incentive to lower their income. Indeed, this form of moral hazard is what drives our argument for the efficiency of universality.

7. Alternative Preferences

We have assumed that preferences are of the "subtractive" form: $U(C_i - N_p L_i)$. An individual-specific need requires a fixed, exogenous cost above and beyond the consumption costs incurred by

the non-need, and this cost is independent of income. Subtractive preferences are appropriate for some needs, like the need for medical care, or the cost of legally mandated expenditures, such as sending children to school or having children vaccinated. Other forms of need, however, alter the cost of achieving a given level of utility, that is, they effectively change prices.¹¹ For example, discrimination may mean that certain groups in the population pay higher prices for goods and services. Consumption technology means that the relationship between income and consumption depends on household size. Larger households enjoy economies of scale in consumption, but face higher prices in that, for example, it costs more for four people to eat at a restaurant than it does for one.

Price-altering needs can be modelled using what we call “equivalence scale” preferences: $U=U(C_i/N_i(A_i),L_i)$ where C_i represents consumption expenditure, N_i an equivalence-scale which reflects the effect of needs on prices, and A_i observable attributes. Equivalence scale preferences can be used to adjust for household size using a welfare function of the form $H_i = A_i U(C_i / N_i(A_i), L_i)$. Here A_i is the number of persons in the household, N_i an equivalence scale based on household size, C_i is total household consumption expenditures, and L_i a vector of household labour supplies. Utility is multiplied by A_i because the government presumably cares about everybody’s consumption and welfare, and weights each person’s utility equally in its social welfare function. Policy makers implicitly invoke this household welfare function whenever they use an equivalence scale to adjust for household size when measuring poverty or income inequality. Because this household welfare function is so widely used, it is interesting to consider its implications for the universality debate.¹²

If $A_i=N_i$, so there are no economies of scale in household consumption, and assuming separability in consumption and labour, then equalizing marginal utility of consumption across

households of varying sizes then implies equalizing consumption per person across households. The optimal tax/transfer system then, ignoring the caveats like differing labour supply elasticities, moral hazard etc. mentioned above) would make benefits proportional to disposable income. The income-splitting method used in the personal income tax systems of countries such as France, which effectively uses an equivalence scale that is equal to the number of persons in the family (including minors) exactly achieves this objective. A more economically sound approach would be to permit more limited income splitting, based on an equivalence scale which reflects scale economies in consumption of larger families and the value of household production. The generosity of income splitting to high income families may render it politically infeasible. In this case, tax exemptions for dependent children partially accomplish the same objectives in as much as they are worth more in higher income brackets.

A full examination of alternative assumptions about how needs enter the utility function is worthwhile but is beyond the scope of this paper. But it is not correct to infer that if we modify our assumptions we inevitably find that it is optimal to give lower benefits to the rich than to the poor. On the contrary, exchanging our assumption of subtractive preferences for an arguably more plausible assumption seems instead to justify giving *higher* benefits to the rich.

8. Conclusions and Policy Implications

We have shown that if individuals' utility functions are of the form $U(C_i - N_i, L_i)$, where the individual-specific needs N_i are exogenous, distributed independently of other relevant characteristics, and are observable to the government, then a government maximizing a Social Welfare Function defined over individual utilities would implement a tax/transfer system separable in needs and monetary income, which is how we define "universality". Clawback of benefits from

the rich is undesirable both because it violates parallel equity, taking from the high-need rich to give to the low-need rich, and because it violates efficiency, giving the high-need a higher net marginal tax rate than the low-need. By modifying our assumptions it is possible to justify deviations from universality, but these deviations are as likely *a priori* to justify clawforward (paying bigger benefits to the rich) as clawback (paying smaller benefits to the rich). Universality is the benchmark case.

The appropriate treatment of needs has long been neglected in the taxation literature. This paper is the first step towards a theoretical foundation for the tax treatment of needs, and for a much needed re-examination of the recognition of needs in Canada's tax/transfer system.

References

- Akerlof, George A (1978) "The Economics of "Tagging" as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning" American Economic Review 68, 8 - 19.
- Atkinson, Anthony (1993) "On Targeting Social Security: Theory and Western Experience with Family Benefits" Discussion Paper Number WSP/99, Welfare State Programme, London School of Economics.
- Besley, Timothy (1990) "Means Testing Versus Universal Provision in Poverty Alleviation Programmes" Economica 57, 119-29.
- Caballero, Ricardo J. (1990) "Consumption Puzzles and Precautionary Savings", Journal of Monetary Economics, 25, 113-136.
- Creedy, John (1996) "Comparing Tax and Transfer Systems: Poverty, Inequality and Target Efficiency" Economica 63 supplement, S163-74.
- Dehejia, V and N. Rowe (1995) "The Laffer Curve and Precautionary Taxation: A Rationale for Paying Down the National Debt" Australian Economics Papers 14, 322-331.
- Garfinkel, Irwin, ed., (1979) Income-Tested Transfer Programs Academic Press.

Keen, M. (1992) "Needs and Targeting" Economic Journal 102, 67-79.

Kesselman, Jonathan R. and Irwin Garfinkel (1978) "Professor Friedman, Meet Lady Rhys-Williams: NIT vs. CIT" Journal of Public Economics 10, 179-216.

Kesselman, Jonathan R. (1993) "The Child Tax Benefit: Simple, Fair, Responsive?" Canadian Public Policy 19, 109-132.

Mirrlees, J. A. (1971) "An Exploration in the Theory of Optimum Income Taxation" Review of Economic Studies 38, 175-208.

Moffitt, R. (1992) "Incentive Effects of the U.S. Welfare System: A Review" Journal of Economic Literature 30, 1-61.

Phipps, Shelley (1995) "Canadian Child Benefits: Behavioural Consequences and Income Adequacy" Canadian Public Policy 21, 20-30.

Endnotes

*.Department of Economics, Carleton University, 1125 Colonel By Drive, Ottawa, Canada K1S 5B6. E-mail: nrowe@ccs.carleton.ca and fwoolley@ccs.carleton.ca. Both authors have children and wear glasses. We thank two anonymous referees, Jon Kesselman, Krishna Pendakur, Andrew Clark, Greg Dow, David Long and seminar participants at Carleton University, Simon Fraser University, and the 1997 CEA meetings for helpful comments. This research was supported through a SSHRC GR6 grant

1. Some programs pay benefits in cash, others in kind. We assume that the two are equivalent.
2. An exception is Kesselman (1993: 118), who describes clearly how the move away from universality affects tax rates. "The Child Benefit scheme in effect finances much of its gains for lower-income families by increasing taxes only on upper-income units with children".
3. An early critique of target efficiency can be found in Kesselman and Garfinkel (1978). For the arguments for and against universality see Atkinson (1993) and Garfinkel (1979). Both Besley (1990) and Creedy (1996) compare universal with targeted benefits, but they assume that individuals differ only in terms of income.
4. We discuss the implications of imperfectly observable needs in section 5 below.
5. By "marginal tax rate" we mean "marginal taxes net of transfers".
6. Another possible aspect is universality of benefits across, say, geographical regions. A similar but stricter definition of universality is that given by Michael Keen (1992: 71), whereby a tax/transfer function is defined as universal if $F(Y, A_i) = Y + S_i$, where S_i is a poll subsidy, constant within but possibly different between groups.
7. FA ceased to be universal in 1974, when FA payments were included in taxable income. Because the tax system is progressive, lower after tax benefits are received at higher income

levels. Universality was further reduced in the mid-1980s, when FA benefits were clawed back from higher income individuals.

8. The government may deliver this insurance better than private firms because many conditions against which one wants to insure are revealed prior to the age at which individuals can sign contracts. Universal social programs then implement the insurance policies that individuals would rationally choose from behind the Rawls/Harsanyi Veil of Ignorance.

9. Imagine Robin Hood, the optimal tax theorist who takes from the rich to give to the poor, has a brother Random Hood, who takes from some people randomly and gives to other people randomly. When Robin sees that Random has upset his carefully-chosen optimal tax system, he restores it by simply reversing Random's transfers.

10. See the literature on precautionary savings (Caballero, 1990) and precautionary taxation, (Dehejia and Rowe, 1995).

11. Social norms may mean that the cost of meeting a given need increases with income. For example, clothing which is perfectly acceptable for a graduate student may be inappropriate if worn by the department chair. A referee suggests we also consider preferences of the form $U(C-NC, L)$, to capture these cases. But if preferences are given by $U = \log(C-NC)$, for example, it can be shown that a government equalizing marginal utility across individuals would give zero benefits for glasses.

12. There are a number of problems with this approach. First, assuming a single household utility function hides household decision making and intra-household inequality. Second, it could be argued that, because household composition is the result of optimizing behaviour on the part of its adult members, the social planner cannot infer anything about relative individual welfare levels from

observing household size or composition. However, the ability of individuals to optimize over household size is severely constrained. Moreover, the recognition of children or other dependants in almost every country's tax/transfer system reveals that social planners care about family composition. It is, therefore, interesting to discuss the consequences of different ways of accounting for family composition.