

Optimal Control of Externalities in the Presence of Income Taxation

Louis Kaplow*

Abstract

A substantial literature examines second-best environmental policy, focusing particularly on how the Pigouvian directive that marginal taxes should equal marginal external harms needs to be modified in light of the preexisting distortion due to labor income taxation. Additional literature is motivated by the possibility that distributive concerns should amend the internalization prescription. It is demonstrated, however, that simple first-best rules – unmodified for labor supply distortion or distribution – are correct in a natural, basic formulation of the problem. Specifically, setting all commodity taxes equal to marginal harms (and subsidies equal to marginal benefits) can generate a Pareto improvement. Likewise, a marginal reform in the direction of the first-best can yield a Pareto improvement. Qualifications and explanations for differences from previous work are also presented.

JEL Classes D61, D62, D63, H21, H23, K32

*Harvard University and National Bureau of Economic Research. I am grateful to Steven Shavell, Roberton Williams, and workshop participants at Harvard University and the NBER for comments and to the John M. Olin Center for Law, Economics, and Business at Harvard University for financial support. This article is related to the larger project, Kaplow (2008).

1. Introduction

The control of externalities is a complex second-best problem. The first-best prescription, following Pigou (1920), is to set marginal taxes and subsidies equal to marginal external harms and benefits. Because of the preexisting distortion due to labor income taxation, however, researchers have explored how the Pigouvian rule needs to be modified on account of the interaction between environmental regulation and the income tax, particularly concerning the need to raise revenue and the indirect effects of environmental policy on labor supply. Work began in the 1970s with Sandmo (1975) and others, following the growing interest in optimal taxation more generally. Subsequently, researchers became interested in the possibility of a double dividend – that corrective taxation might both enhance welfare by internalizing externalities and also raise revenue that would allow a reduction in distortionary income taxation. See, for example, Ballard and Medema (1993), Cordes, Nicholson, and Sammartino (1990), and Pearce (1991). This work, in turn, led to an extensive modern literature on environmental regulation, much of which suggests that there is no double dividend and, instead, that optimal environmental control may well fall short of the Pigouvian first best. For a survey and a collection of literature, see respectively Bovenberg and Goulder (2002) and Goulder (2002).

An additional second-best problem involves distribution. See, for example, Casler and Rafiqui (1993) and West (2004). Because many proposed correctives, such as heightened taxation of gasoline, are believed to be regressive and, moreover, environmental benefits may have values that rise (perhaps disproportionately) with income, distributive concerns seem potentially important. Indeed, Pigou (1928) himself argued that distributive effects needed to be taken into account in assessing corrective taxes and subsidies.

This article seeks to advance our understanding of the regulation of externalities by examining the problem in what seems to be a natural and foundational setting in light of work on taxation more generally: namely, in a model in which there can be taxes or subsidies on each commodity and a nonlinear income tax, and in which individuals differ in their earning abilities. Specifically, the question addressed is, in a world in which there may be externalities, what reforms of commodity taxes and subsidies can generate Pareto improvements, beginning with any arbitrary initial system. By including a preexisting income tax – with no restriction on how it is set initially – interactions involving tax revenue and labor supply are addressed. And by focusing on Pareto improvements, concerns about distribution are taken into account.

The principle result is that, in a standard, simple setting (specifically, in which utility is weakly separable in labor), it is possible to move from any set of commodity taxes and subsidies to first-best Pigouvian taxes and subsidies – wherein each tax or subsidy equals marginal harm or benefit – in a manner that generates a Pareto improvement. In particular, this result arises when the income tax is adjusted such that the reform as a whole is distribution neutral. An extension shows that, subject to some additional assumptions, any marginal change in commodity taxes and subsidies that is proportionally in the direction of the first best produces a Pareto improvement if a distribution-neutral income tax adjustment is employed. In addition, a necessary and sufficient condition is offered that indicates which other commodity tax reforms

make possible Pareto improvements. The article also addresses how relaxing distribution neutrality and weak separability affects the results and discusses how these assumptions explain differences from prior work.

Before proceeding, it is useful to relate the present analysis and results to two additional strands of literature. The first pertains to the method of proof employed here, which as mentioned involves an adjustment to the preexisting income tax such that the reform as a whole is distribution neutral. This approach follows Hylland and Zeckhauser (1979), Kaplow (1996), and some subsequent work that focuses on public goods. Kaplow (1996, 2004, 2008) has discussed informally the application of this approach to the regulation of externalities; the contribution of the present article is to analyze the problem formally, offer extensions, and relate the results to existing literature on environmental regulation. When distribution-neutral income tax adjustments are utilized, distributive concerns are obviously taken into account, and it becomes possible to characterize Pareto improvements. Furthermore, as will be explained, when reforms are implemented in this distribution-neutral manner, labor supply is unaffected in a baseline case; hence, concerns about second-best interactions with the preexisting income tax also become moot. Accordingly, there is an important sense in which the concerns about labor supply and distribution that have occupied much of the second-best literature on externalities are independent of the question of how best to control externalities.

The present investigation can also be related to the literature on optimal commodity taxation in the presence of an income tax. The seminal analysis of this model – for the case in which there are no externalities – is provided by Atkinson and Stiglitz (1976), who examine optimal commodity taxation when the income tax is also set optimally. They find that, with weak labor separability, uniform commodity taxation (equivalent to no commodity taxation) is optimal. A way to understand their result is to realize that, with the separability assumption, differential commodity taxation cannot help to reduce the labor-leisure distortion due to labor income taxation; hence, its only effect is to distort consumption allocations. Accordingly, it is optimal for commodity price ratios to equal the corresponding ratios of social resource costs, which in their model consist solely of production costs. In the present setting, social resource costs include externalities, so the natural extension of their finding is that price ratios should equal the ratios of the sum of production costs and external effects, which is precisely what obtains if one follows Pigou's (1920) prescription by setting commodity taxes and subsidies equal to marginal external harms and benefits.

The article formalizes this intuition, thereby extending Atkinson and Stiglitz's (1976) important result to settings with externalities. Furthermore, the analysis here is not restricted to cases in which the income tax is optimal and is not limited to commodity tax reforms in the neighborhood of their optimum.¹ In these two respects, the present results are most analogous to

¹For preliminary analyses of the case with externalities, see Pirttilä and Tuomala (1997) and Cremer, Gahvari, and Ladoux (1998), each of which examines formula for optimal commodity taxes (not permitting analysis of reforms away from the optimum) for the case in which the income tax is assumed to be optimal. Specifically, Cremer, Gahvari, and Ladoux (1998) find that, when the only effect of an externality regarding a single good is additively separable, a term involving the externality appears only in the first-order condition involving that good, and Pirttilä and Tuomala (1997)

those in Kaplow's (2006) recent article on commodity taxation. See also Laroque (2005) and Kaplow (2008). Nevertheless, the bulk of the analysis here is different precisely because of the presence of externalities. Furthermore, it is difficult to view the present results – however intuitively plausible they may seem on reflection – as obvious when most existing literature on externalities appears to advance qualitatively different prescriptions. Formal analysis is necessary to understand when and why the conclusions from various models diverge and to assess the implications of the difference in approaches.

2. Optimality of First-Best Pigouvian Taxes and Subsidies

2.1. Model

Individuals earn income y , which is equal to wl , the product of their wage w and their labor effort l . Wages (types) are distributed according to the density function $f(w)$. The government imposes a (nonlinear) income tax $T(y)$. As is standard, the government observes income but does not separately observe individuals' wages or choices of labor effort.

Individuals spend their after-income-tax income, $y - T(y)$, on n commodities, x_1, \dots, x_n . Commodity prices for the x_i (which equal production costs) are denoted p_i and are assumed to be positive constants. There are also commodity taxes on each x_i of τ_i (which may be subsidies, in which case they are negative). Individuals thus face net prices of $p_i + \tau_i$.

An individual of type w 's choice of labor l is equivalent to a choice of income y that will be denoted $y(w)$. For the chosen level of income, and taking the tax system (and the level of externalities, see below) as given, the individual chooses how much to expend on each commodity, $x_i(y(w))$. The individual's budget constraint can thus be written as

$$(1) \sum (p_i + \tau_i)x_i(y(w)) = y(w) - T(y(w)),$$

where summations throughout are from i equals 1 to n .

Each commodity x_i is associated with an externality e_i , which denotes the total consumption by all individuals of that commodity:

$$(2) e_i = \int x_i(y(w))f(w)dw.$$

When individuals choose $y(w)$ and the $x_i(y(w))$'s, they are assumed to take the e_i as given (which is appropriate since there is a continuum of individuals, each of measure zero).

Individuals' utility functions are $u(v(x_1, \dots, x_n, e_1, \dots, e_n), l)$, where v is a subutility

find that the commodity tax formula for the dirty good in their model has an additional term which, under a similar separability assumption, involves only the externality.

function. This utility function is assumed to be continuously differentiable, strictly concave, increasing in commodities, and decreasing in labor effort. Utility may have any relationship to the levels of the e_i 's; that is, the external effect due to each of the commodities may be positive, negative, or nonexistent. This form of the utility function, with the subutility function v , entails what is referred to as weak separability of labor (or leisure): For a given level of after-income-tax income, individuals will allocate their disposable income among commodities in the same manner regardless of the level of labor effort required to earn that level of income; furthermore, in choosing labor effort, it does not matter what levels of consumption and externalities combine to produce a given level of subutility v .

In the analysis that follows, use will be made of the indirect subutility function $V(\boldsymbol{\tau}, T, y)$, which is the value of $v(x_1, \dots, x_n, e_1, \dots, e_n)$, maximized over the x_i 's, where the commodity tax vector $\boldsymbol{\tau}$, the income tax schedule T , and before-tax income y (as well as the e_i 's) are taken as given. Observe that since v depends only on the x_i 's (and the e_i 's), and since the budget constraint (1) depends only on the x_i 's, $\boldsymbol{\tau}$, T , and y – and not on w or l independently – the indirect subutility function V is the same for all individuals, regardless of their type w . It will also be useful to use this indirect subutility function to define $U(\boldsymbol{\tau}, T, w, l) = u(V(\boldsymbol{\tau}, T, y), l)$.

The government's budget constraint is

$$(3) \int [T(y(w)) + \sum \tau_i x_i(y(w))] f(w) dw = R,$$

where R is a given revenue requirement.

Finally, the marginal external harm associated with any commodity is

$$(4) h_i = - \int \frac{u_v(w) v_{e_i}(w)}{\lambda(w)} f(w) dw,$$

where subscripts indicate partial derivatives, the notation (w) indicates evaluation for an individual of type w , and $\lambda(w)$ refers to the Lagrange multiplier (for individuals of type w), which signifies the marginal utility of disposable income. The fraction in the integrand, therefore, is the marginal effect of the externality on utility divided by the marginal utility of disposable income, which denotes the marginal external effect for a given type w , measured in dollars. Note that for positive externalities, $h_i < 0$. Now we can define:

First-best Pigouvian taxes and subsidies. A commodity tax vector $\{\tau_1, \dots, \tau_n\}$ consists of first-best Pigouvian taxes and subsidies if and only if $(p_i + \tau_i)/(p_j + \tau_j) = (p_i + h_i)/(p_j + h_j)$, for all i, j .

Notice that the definition does not require that $\tau_i = h_i$, for all i . The reason has to do with normalization: If all commodity taxes are raised or lowered in such a manner as to leave all consumer price ratios unchanged, individuals' behavior will be unaffected – if the level of the income tax is also adjusted to produce the same effective disposable income. (This point is

analogous to the familiar idea that uniform commodity taxes are equivalent to a proportional adjustment to the income tax.)

2.2. Analysis

The approach is to begin with any regime that does not consist of first-best Pigouvian taxes and subsidies and then to construct a regime with first-best Pigouvian taxes and subsidies that makes everyone better off. As an intermediate step in analyzing any reform (not necessarily a move to the social optimum), say, from tax regime $\{\tau, T\}$ to regime $\{\tau^*, T^*\}$, it is useful to examine the intermediate income tax schedule T° that has the property that each individual's utility under $\{\tau^*, T^\circ\}$ is the same as that under $\{\tau, T\}$ if the individual continues to choose the same level of labor effort l . (Whether individuals will choose the same labor effort under this intermediate regime is the subject of Lemma 1, below.)

To be more precise, define $T^\circ(y)$ such that $V(\tau, T, y) = V(\tau^*, T^\circ, y)$ for all y . In general, a reform from τ to τ^* will, for a given y , change the value of subutility V . For each y , the tax schedule $T^\circ(y)$ is set at the level that restores the original level of subutility (by changing after-tax income for the stipulated level of before-tax income). This tax adjustment, $T^\circ(y) - T(y)$, is simply the schedule of utility-compensating changes in disposable income. Given how the intermediate income tax schedule is constructed, it is possible to establish the following result.

Lemma 1: Every type of individual w chooses the same level of labor effort under $\{\tau^*, T^\circ\}$ as under $\{\tau, T\}$.

Proof: It is straightforward to establish that $U(\tau, T, w, l) = u(V(\tau, T, y), l) = u(V(\tau^*, T^\circ, y), l) = U(\tau^*, T^\circ, w, l)$, for all w, l . The first equality follows by the definition of U . The second equality follows because $T^\circ(y)$ is constructed such that $V(\tau, T, y) = V(\tau^*, T^\circ, y)$ for all y . And the third equality also follows from the definition of U . Therefore, $U(\tau, T, w, l) = U(\tau^*, T^\circ, w, l)$, for all w, l . Because, for any type w , this equality holds for all l , the level of utility an individual achieves for each possible choice of l is the same in each of the two regimes. Therefore, for each type w , whatever l maximizes U in the initial regime $\{\tau, T\}$ must be the l that maximizes U in the intermediate regime $\{\tau^*, T^\circ\}$. ■

The intuition is as follows: The first part of the lemma merely shows the obvious point that, if the subutility V is unchanged for every y and thus for every choice of l (which is true by construction), then indeed U is unchanged for every choice of l . The second part of the lemma simply observes that since U as a function of l is unchanged when moving from $\{\tau, T\}$ to $\{\tau^*, T^\circ\}$, the choice of l that maximizes U is likewise unchanged. This result that labor supply is unaffected greatly simplifies the analysis of any tax reform, including those of particular interest here.

To begin the construction of a Pareto-improving tax reform, start with an initial regime $\{\tau, T\}$ that is not a Pigouvian first best (as defined above). For simplicity, choose from among the multitude of equivalent first-best Pigouvian tax systems the one for which $\tau_i^* = h_i$, for all i .

For purposes of the argument to follow, it will be assumed that the h_i 's are constant and thus independent of the x_i 's and the e_i 's. (The more general case will be commented on below.) Moving to this new commodity tax vector will tend to change individuals' utilities on account of three effects: changes in the payment and receipt of commodity taxes and subsidies, changes in consumption due to the new relative price vector, and changes in externalities on account of others' changes in consumption. Whatever is the net effect on utility for any type w and given labor effort for that type, we will consider the intermediate income tax schedule $T^\circ(y)$, defined above, that offsets the net effect on utility at each income level.

We know that from Lemma 1 that moving to this intermediate regime leaves labor supply unchanged. The next question to consider is how revenue compares between the initial regime and the intermediate regime.

Lemma 2: If the h_i 's are constant, then regime $\{\tau^*, T^\circ\}$ (with first-best Pigouvian taxes and subsidies $\tau_i^* = h_i$, for all i) raises more revenue than does regime $\{\tau, T\}$ (without first-best Pigouvian taxes and subsidies).

Proof: To begin, we can see from expression (3) that the change in revenue is given by

$$(5) \int \left[(T^\circ(y(w)) - T(y(w))) + \sum (\tau_i^* - \tau_i) x_i(y(w)) + \sum \tau_i^* (x_i^\circ(y(w)) - x_i(y(w))) \right] f(w) dw.$$

The first term in the integrand is the change in income tax revenue for an individual of type w , who earns $y(w)$, from moving to the intermediate regime. The next two terms are the change in commodity tax revenue for the individual, which is decomposed for convenience: The second term is the revenue change from moving to the first-best Pigouvian commodity tax vector, holding constant the original level of consumption, and the third term is the revenue change from the individual's adjustment in consumption, evaluated using the new commodity tax vector, \mathbf{x}° . These three terms are integrated over the population to yield the total change in revenue.

In analyzing expression (5), it is useful to decompose the first term, the change in income tax revenue, by imagining that the shift to the intermediate regime occurs in three stages:

1. The commodity tax reform is introduced but the x_i 's are imagined to be held constant, and the $T(y)$ schedule is adjusted to hold individuals at the same utility levels taking into account that commodity tax payments are changed – the second term in expression (5). Let $T_1^\circ(y)$ denote the resulting income tax schedule.
2. Individuals adjust their x_i 's but the e_i 's are imagined to be held constant, and the $T(y)$ schedule is further modified, to $T_2^\circ(y)$, for the effect of these adjustments on individuals' own levels of utility.
3. The e_i 's are allowed to reflect the new level of the x_i 's (which in turn may be further adjusted because of the change in the levels of the e_i 's, ultimately yielding the x_i° 's), and the $T(y)$ schedule is further modified, now to $T^\circ(y)$.

Using this notation, expression (5) can be rewritten as

$$(6) \int \left[(T^\circ(y(w)) - T_2^\circ(y(w))) + (T_2^\circ(y(w)) - T_1^\circ(y(w))) + (T_1^\circ(y(w)) - T(y(w))) \right] f(w) dw$$

$$+ \int \left[\sum (\tau_i^* - \tau_i) x_i(y(w)) + \sum \tau_i^* (x_i^\circ(y(w)) - x_i(y(w))) \right] f(w) dw.$$

Stage 1: Beginning with the first stage, in which the x_i 's are held constant, the income tax schedule $T(y)$ for each type w is raised (lowered), to the schedule $T_1^\circ(y)$, by an amount equal to the negative of the fourth term (the first term in line 2) in expression (6), which denotes the increase (or reduction) in expenditures on commodities necessary to finance a type- w individual's original consumption bundle. Note that, after this adjustment to the income tax, each type of individual will just be able to afford the original level of consumption. Because the increase (decrease) in income tax revenue for each type of individual, $T_1^\circ(y) - T(y)$, precisely equals the decrease (increase) in commodity tax and subsidy revenue for each, it follows that the third and fourth terms in (6) are precisely offsetting at each y and thus also when each is integrated over the population. Therefore, the net revenue effect with regard to this first stage is nil.

Stage 2: Now we move to the second stage and allow individuals to adjust their x_i 's, but we imagine that the e_i 's are still held constant. It will now be demonstrated that there is a positive revenue effect on account of the change in $T(y)$ necessary to offset the direct effect on individuals' utility from their own changes in consumption. That is, $T_2^\circ(y) - T_1^\circ(y) > 0$ for all y . This arises because, in this hypothetical setting, all individuals would wish to change their consumption, which they would do only if they stood to gain.

To show this result, begin with the fact that an individual's optimal consumption vector is determined by standard first-order conditions, which in the initial regime are $(\partial v / \partial x_i) / (\partial v / \partial x_j) = (p_i + \tau_i) / (p_j + \tau_j)$, for all i, j . Given that the commodity taxes in the original regime are not first-best Pigouvian taxes and subsidies, there exists i, j such that $(p_i + \tau_i) / (p_j + \tau_j) \neq (p_i + h_i) / (p_j + h_j)$. However, the latter term equals $(p_i + \tau_i^*) / (p_j + \tau_j^*)$. Because the first-order conditions are all satisfied in the initial regime, at least one must therefore be violated after the commodity tax reform if the consumption bundle is unchanged. As a consequence, the consumption vector that was optimal under the original regime cannot be optimal under the hypothesized intermediate regime. Now, for individuals thus adjusting their consumption on account of the change in price ratios, it follows that their utility must be higher under the intermediate regime unless the income tax schedule is adjusted upwards to reduce their disposable income. Of course, precisely this adjustment must be done in the intermediate regime because $T_2^\circ(y)$ is, by definition, constructed so as to hold everyone's utility constant at the end of stage 2. Accordingly, with regard to this second stage, it is indeed true that $T_2^\circ(y) - T_1^\circ(y) > 0$ for all y ; that is, the second term in (6) is positive. The effect of changes in consumption on commodity tax revenue is considered in combination with stage 3.

Stage 3: Finally, in the third stage, we allow the e_i 's to reflect the new level of the x_i 's – which in turn may be further adjusted on account of the change in the levels of the e_i 's,

ultimately yielding the x_i° 's – and we examine what further change in the $T(y)$ schedule, from $T_2^\circ(y)$ to $T^\circ(y)$, is necessary to compensate for this change in the levels of externalities. Furthermore, we can compare this income tax revenue effect to the change in commodity tax and subsidy revenue due to the shifts in consumption, the final term in expression (6). It will now be demonstrated that these two effects must offset because the commodity taxes are being used as corrective taxes (and subsidies), set at the level that internalizes any externalities.

The compensation for all changes in the level of externalities is, as stated, accomplished by moving the income tax schedule from $T_2^\circ(y)$ to $T^\circ(y)$. The aggregate effect on income tax revenues is:

$$(7) \int (T^\circ(y(w)) - T_2^\circ(y(w)))f(w)dw$$

$$= \sum \int \left[\frac{u_v(w)v_{e_i}(w)}{\lambda(w)} \left(\int (x_i^\circ(y(w)) - x_i(y(w)))f(w)dw \right) \right] f(w)dw = -\sum h_i(e_i^\circ - e_i).$$

The left side of the second line of expression (7) is the aggregate effect on individuals' utility (measured in dollars) due to changes in the level of the externality. The integrand of the internal integral is the change in the levels of consumption; integrated over all individuals, this gives the change in the levels of the externalities, as recorded on the right side. For a unit change in the level of any particular externality e_i , the integral of the effect on utility, measured in dollars, over the population equals the negative of the marginal external harm, h_i , as stated in expression (4).

Let us compare this revenue effect due to the compensatory adjustment in the income tax schedule to the final term in expression (6), the change in revenue through the commodity taxes and subsidies on account of individuals' changes in consumption.

$$(8) \int \left[\sum \tau_i^* (x_i^\circ(y(w)) - x_i(y(w))) \right] f(w)dw = \sum h_i \int [(x_i^\circ(y(w)) - x_i(y(w)))f(w)dw]$$

$$= \sum h_i (e_i^\circ - e_i).$$

The left side of (8) shows that the change in total commodity tax (and subsidy) collections on account of consumption adjustments equals the integral over the population of each individual's change in commodity tax payments due to the shift in consumption in moving from the original regime to the intermediate regime. On the right side of (8) in the top line, h_i is substituted for τ_i^* because the τ_i^* 's are defined to be first-best Pigouvian taxes and subsidies (normalized to equal external harm), and the orders of the summation and integration are reversed for convenience. Finally, the move to the second line of (8) reflects, as with expression (7), that the levels of the externalities are defined as the integral of consumption over the population as a whole.

Comparing expressions (7) and (8), it is apparent that the amount of revenue lost (gained) through this third stage in the income tax adjustment – the compensation for externalities – precisely equals the revenue gain (loss) through commodity taxes – which are now set at first-best Pigouvian levels – on account of individuals adjusting their levels of consumption. Therefore, the first and last terms in expression (6) offset.

Summarizing, from the third stage, the first and last terms in (6) offset; from the first stage, the third and fourth terms in (6) offset; and, from the second stage, the second term in (6) is positive. Therefore, it has been demonstrated that expression (6) is positive. ■

The intuition behind this proof is as follows. For stage 1, the only effects are on taxes (behavior is held constant); hence, to keep each individual at the same level of utility, it must be that the first component of the intermediate income tax adjustment just offsets the effect of the commodity tax reform, at each level of income. (For example, if the reform makes individuals at a given income y spend more because taxes on commodities causing negative externalities are higher, it must be that income taxes at income y are reduced accordingly so that individuals can afford the same consumption bundle.) This stage obviously is revenue neutral. In stage 2, individuals are adjusting their x_i 's and everything else is held constant, except for the second component of the income tax adjustment. Because individuals necessarily increase their utility when they voluntarily choose to change their consumption, the income tax adjustment that holds utility constant must reduce their disposable income and therefore raise revenue. Finally, the stage 3 income tax adjustment that compensates for externalities just offsets the change in commodity tax revenue due to changes in consumption. The basic reason is that commodity taxes now equal marginal external effects, so they collect (expend) an amount equal to marginal harm (benefit) from consumption adjustments, which in turn equals the amount that income tax collections must fall (rise) to offset the effect of the change in externalities on individuals' utilities.

To complete the argument that a Pareto improvement is possible in moving to first-best Pigouvian taxes and subsidies, construct $T^*(y)$ from $T^o(y)$ by performing two operations simultaneously: Gradually rebate the surplus pro rata, and, as individuals adjust their levels of consumption, adjust the income tax schedule in order to offset the effects on utility of changes in the levels of externalities, as described in the proof of Lemma 2 regarding stage 3. As with stage 3, this latter operation is revenue neutral when one combines the revenue effects from the income tax adjustment with those from consumption adjustments changing the level of commodity tax receipts. The former operation (raising the rebate) will ultimately exhaust the surplus. Moreover – and of central importance to the present argument – the effect from this operation is that all individuals' utility will increase.²

At this point, we have constructed a new income tax schedule, to accompany the move to first-best Pigouvian taxes and subsidies, such that every type of individual is strictly better off,

²Note that, as the income tax schedule is reduced, individuals may reduce labor supply and thus tax revenue may fall, but as long as there initially is a surplus and aggregate behavior is continuous, some net reduction in everyone's income tax payment through the former operation will be possible.

thereby establishing the following result:

Proposition 1: If the h_i 's are constant, then for any tax system $\{\tau, T\}$ for which commodity taxes and subsidies are not at first-best Pigouvian levels, there exists a tax system $\{\tau^*, T^*\}$ with first-best Pigouvian taxes and subsidies that is strictly Pareto superior – i.e., $U(\tau^*, T^*, w, l(w)) > U(\tau, T, w, l(w))$, for all w .

The role of the assumption that the h_i 's are constant and thus independent of the x_i 's and the e_i 's in the proof pertains to stage 3 in Lemma 2 and to the rebate process. These arguments rely on the fact that – with τ_i 's that have been set equal to the h_i 's – changes in commodity tax and subsidy payments due to changes in consumption are necessarily equal to changes in the level of external harm and benefit due to those same changes in consumption. If one allows the h_i 's to be endogenous, one would need to modify the argument by employing a continuous adjustment process, roughly as follows. Initially, reform the τ_i 's to equal the h_i 's at the existing level of the x_i 's. Then, (hypothetically) allow the x_i 's to adjust gradually. One can then perform all of the pertinent operations used in the proof of Lemma 2 and in the rebate process, but in each case with the level of the τ_i 's being continuously adjusted as the x_i 's change so that, at every point, the τ_i 's equal the h_i 's. For each increment of these marginal changes, the preceding analysis applies. Hence, once the adjustment considered in Lemma 2 is complete, the net effect on revenue will be the integral of positive marginal changes due to the stage 2 income tax adjustments – all other effects netting to zero as before – so revenue will necessarily increase. Likewise, the rebate process can be implemented as before, subject to further adjustment of the τ_i 's. Also note that, because in all phases of the construction the τ_i 's are set equal to the h_i 's, when the construction is complete the system $\{\tau^*, T^*\}$ will indeed be one that has first-best Pigouvian taxes and subsidies. (Finally, observe that the endogeneity of the h_i 's would be moot if we were considering only local changes near the Pigouvian first best.)

It is useful to remark on the intuition underlying Proposition 1. Any direct effects of the move to first-best Pigouvian taxes and subsidies on individuals' budget constraints can be offset by an income tax adjustment, after which individuals can continue to afford their original consumption bundle. Given that price ratios have changed, individuals will want to change their levels of consumption, and the direct effect of such adjustments on the individuals who make them is utility increasing (for otherwise they would not have chosen to alter their consumption). External effects of these changes in consumption are fully compensated through further adjustments to the income tax schedule, and this compensation is fully funded by changes in receipts from the commodity taxes and subsidies, each of which equals the pertinent marginal external effect. Hence, just as in a world in which there is fixed labor supply – or simply endowments of goods and no labor decision – and no redistributive, distortionary income tax, all individuals can be made better off by the complete internalization of externalities.

Distribution and labor supply effects do not interfere with the argument because the income tax schedule is adjusted to keep them fixed. The intermediate tax schedule (before the rebate, which makes everyone better off) keeps everyone's utility constant; hence, there are no distributive effects to consider. Moreover, as explained in proving Lemma 1, when this is done, the return to any choice of labor effort, for individuals of any type w , is the same, so individuals

will be induced to keep their labor supply unchanged. Thus, in the present construction, the income tax problem – involving the tradeoff of distribution and labor supply distortion – is independent of the problem of controlling externalities, so the solution to the latter problem has its familiar first-best character.

3. Extensions

3.1. Reforms other than to First-Best Pigouvian Taxes and Subsidies

When reformed commodity taxes and subsidies do not all equal the level of marginal external harm or benefit, the change in commodity tax revenue due to shifts in consumption need not be sufficient to finance the compensation for the change in the levels of the externalities provided by the stage 3 income tax adjustment. Consider first reforms that move proportionately toward the first-best commodity tax regime. Specifically, beginning with an income tax $T(y)$ and an initial set of τ_i 's that are not first-best Pigouvian taxes and subsidies, such a reform can be described formally as a move to $\tau_i(\alpha) = \alpha h_i + (1-\alpha)\tau_i$, for all i , where $\alpha \in [0, 1]$. Furthermore, because the h_i are themselves endogenous, being influenced by changes in α , consider the marginal reform $d\tau_i(\alpha)/d\alpha$, for all i , evaluated at $\alpha = 0$. For the income tax adjustment, we can define the income tax schedule $T(y, \alpha)$ that changes as a function of α so as to keep utility constant at every level of income and thus for every type.

Because increasing α moves in the direction of the first best, changes in consumption tend to be desirable. Individuals will tend to consume less of commodities that are underpriced (for example, those with undercorrected negative externalities) because those prices will rise, and individuals will tend to consume more of overpriced commodities (overcorrected negative externalities and undercorrected positive externalities) because those prices will fall. There is, however, a qualification due to possible cross-effects. To illustrate, suppose that initially all commodity taxes are zero and only two goods involve externalities. Flowers produce a positive externality (to those driving by), and driving, corresponding, say, to the commodity gasoline, produces a negative externality (pollution). The reform would provide a small subsidy to flowers and impose a small tax on gasoline. Suppose, however, that driving and the positive externality produced by flowers are highly complementary. Then the increase in flowers might induce individuals to increase driving – that is, the effect due to the complementarity might exceed the contrary effect due to the tax on gasoline. Additionally, if this increase is sufficiently great, it is possible that the reform will be inefficient. A sufficient condition that rules out this possibility is $\partial^2 v / \partial e_j \partial x_i = 0$, for all i, j .

One further complication is that the pro rata rebate of the revenue surplus could actually make individuals worse off rather than better off when there are uncorrected negative externalities. (In such a situation, it would tend to be optimal for the government to impose a tax, such as a uniform lump-sum tax, and destroy the proceeds.) To set this problem aside, assume that “more is better for everyone,” specifically, that if the government finds itself with a surplus, there exists some manner of rebating it to the population (not necessarily pro rata, given that externalities may affect individuals earning different levels of income differently) that produces a Pareto improvement.

Although the analysis of this case turns out to be more involved than that which establishes Proposition 1, it is possible to demonstrate the following result (which is proved in the Appendix):

Proposition 2: If $\partial^2 v / \partial e_j \partial x_i = 0$, for all i, j , and if “more is better for everyone,” then for any tax system $\{\tau, T\}$ for which commodity taxes and subsidies are not at first-best Pigouvian levels, for a marginal increase in α , evaluated at $\alpha = 0$, where $\tau_i(\alpha) = \alpha h_i + (1 - \alpha)\tau_i$, for all i , there exists $T(y, \alpha)$ such that the marginal reform is strictly Pareto improving – i.e., $dU(\tau(\alpha), T(y(w), \alpha), w, l(w)) / d\alpha|_{\alpha=0} > 0$, for all w .

It is more difficult to characterize reforms that do not move proportionally toward the first best. Taking a simple reform, it is familiar that moving a single tax or subsidy in the direction of the first-best need not be efficient. For example, raising a tax that was previously too low may cause substitution to another good whose tax was also too low, perhaps to an even greater extent. Nevertheless, it turns out that there is a conceptually simple, if not always easy to apply, test of when a Pareto improvement is possible. This test involves a narrow, traditional concept of efficiency.

Efficiency-increasing commodity tax reform. For any tax system $\{\tau, T\}$, a commodity tax reform τ^* is efficiency increasing if, when combined with the intermediate income tax schedule $T^\circ(y)$ as defined previously,

$$(9) \sum p_i \int x_i^\circ(y(w)) f(w) dw < \sum p_i \int x_i(y(w)) f(w) dw.$$

Expression (9) states that the total real resource cost of everyone’s consumption vectors in the intermediate regime is less than the total real resource cost in the initial regime. Because everyone’s utility is the same in these two regimes, this condition indicates that the intermediate regime is more efficient with regard to consumption choices in a narrow, conventional sense – i.e., when concerns with the labor-leisure distortion and distribution are ignored. Note that this condition makes no explicit reference to externalities; externalities are relevant, however, because the extent of positive and negative externalities will affect the level of consumption necessary to bring individuals to the same level of utility as in the initial regime. It is straightforward to demonstrate the following (which is proved in the Appendix):

Proposition 3: If “more is better for everyone,” then beginning with any tax system $\{\tau, T\}$, for any efficiency-increasing commodity tax reform τ^* , there exists $T^*(y)$ such that the reform regime is strictly Pareto superior – i.e., $U(\tau^*, T^*, w, l(w)) > U(\tau, T, w, l(w))$, for all w .

In essence, Proposition 3 states that, if a commodity tax reform increases efficiency in a traditional sense – that is, if it increases surplus in a world in which labor supply is constant, tantamount to a world with fixed labor supply or simply one in which initial wealth endowments are given – then the reform will be desirable, indeed strictly Pareto improving, when combined

with an appropriate income tax adjustment, even in a world in which labor supply is not constant and there exists a distortionary labor income tax. This final proposition reinforces the sense in which the environmental tax problem and the income tax problem can be viewed as independent (given the assumption of weak separability).

3.2. Non-Distribution-Neutral Income Tax Adjustment

The foregoing results assume that the intermediate income tax adjustment holds utility fixed, which necessarily entails that it is distribution neutral (in terms of utilities). This subsection considers the implications of relaxing this assumption and also comments on how deviations from distribution neutrality explain some (but not all) of the difference between the present results and those in the pertinent literature.

Distributive effects may have three sources. The first is the incidence of commodity tax changes. The direct effect in terms of commodity taxes paid and subsidies received may depend on income (and thus type and utility) due to differing aggregate consumption and differing relative consumption; likewise, the utility effect from adjustments to private consumption may vary with income. Second, the benefits of controlling externalities may vary with income. (As elaborated further below, note that even if these benefits are additively separable in utility, the benefits when measured in dollars will rise with income – and, in general, nonlinearly, depending on the rate of change in the marginal utility of consumption.) Third, the income tax adjustment can in principle take any form. This article constructs the third to just offset the first two in terms of effects on utility. When this is not done, the overall result will not be distribution neutral.³

Fortunately, the analysis of non-distribution-neutral income tax adjustments is conceptually straightforward, which can be seen by the following decomposition: (1) Implement the commodity tax reform using the distribution-neutral income tax adjustment specified previously. (2) Instantly thereafter, implement a further reform of the income tax, from the distribution-neutral adjustment of step 1 to whatever (non-distribution-neutral) income tax adjustment one wishes to consider.

The analysis of step 1 is that presented previously. To complete the analysis of the non-distribution-neutral reform, one needs to examine step 2. Observe that this second step is a purely redistributive adjustment to the income tax, and accordingly it would be analyzed as such. For example, if the initial regime is close to the redistributive optimum and the reform is small, then it is approximately true that no further adjustment to the welfare analysis should be required because at the optimum the distribution-distortion tradeoff is a wash. If instead the reform, say, increased redistribution and the status quo redistributed too little, then there would be a further

³Absent some specific public choice argument, the most natural conjecture, if one were required, seems to be that such reforms will, on average, tend to leave the preexisting political equilibrium regarding the extent of redistribution unaltered. In that case, distribution-neutral analysis would not merely be a useful conceptual benchmark; it would also be a plausible approximation of political reality.

welfare gain, whereas if the status quo redistributed too much, there would be an associated welfare loss. Note that in these latter cases, there would still be partial offsets: For example, if redistribution becomes even more excessive, the incremental distortion would overstate the welfare loss because there would still be some welfare gain from the additional redistribution.

This discussion illuminates some of the existing literature. The explanation may not be immediately apparent because much work uses representative-individual models, in which distributive effects are literally moot. Yet a concern for distribution is what motivates the common stipulation that proportional income tax adjustments rather than adjustments to a uniform grant are the marginal source of finance. Accordingly, it is natural to ask what the distributive effects would be if one superimposed the reform and stipulated income tax adjustment on a model like that examined here in which individuals do vary in income-earning ability. In that case, distributive effects may well be present. Furthermore, the second step of the decomposition makes apparent that, if the contemplated reforms implicitly increase (decrease) redistribution, they would be associated with additional (reduced) distortion, which would make the seemingly optimal commodity tax rate below (above) the Pigouvian level.

Upon examination, some of the prior literature does consider reforms that implicitly increase redistribution and thus should be expected to involve distortion due to that effect. For example, work on pollution quotas (surveyed in Bovenberg and Goulder, 2002) compares environmental regimes that themselves have different effects on revenue, with any gap made up by changing marginal income tax rates. The regimes that raise more revenue are thus less distortionary, but the distortion reduction is a consequence of reduced use of redistributive income taxation. In those models, the other revenue differences have a lump-sum character; hence, regimes with lower revenue are equivalent to modifications of a linear income tax that increase the uniform grant financed by an increase in the marginal tax rate, a combination that obviously raises redistribution as well as distortion. This characterization is reinforced by Fullerton's (1997, p. 248 n.8) explanation of the second-best literature's weak double-dividend claim: "if an uncorrected externality is subjected to initial taxation, then welfare is higher if the revenue is used to reduce other distorting taxes than if it is returned to consumers lump sum." To take a different type of example, Goulder et al. (1999) use a tax adjustment that holds constant the real purchasing power of the income transfer component of the tax-transfer system. This property implies that the full cost of improved environmental protection is borne by richer individuals, so increased environmental protection is implicitly linked to greater redistribution, which in turn produces labor supply distortion.

A more broadly applicable yet subtle effect is due to the fact that the curvature of (sub)utility as a function of consumption is simultaneously related to the slope of the labor supply curve (see Chetty 2006) and also to the incidence of commodity tax reforms as well as that of the accompanying benefits from controlling externalities. Although this connection is not widely appreciated, the implication is that explicit assumptions about labor supply – typically, that the labor supply curve is upward sloping – entail implicit assumptions about the distributive incidence of reforms.

To illustrate this elusive point, consider the following additively separable utility function: $u(x_c, x_d, e, l) = (x_c^{-\rho} x_d^{-\rho})^{1-\rho} / (1-\rho) - b(e) - z(l)$, where x_c is a clean good (no externality), x_d is a dirty good (negative externality), and e is total consumption of x_d . (This utility function exhibits constant relative risk aversion in consumption of ρ ; in the limit as ρ approaches 1, it is appropriate to substitute the log of the value in the first parentheses.) Suppose further that there is a single commodity tax τ_d on x_d and that there is a linear income tax with rate t and a zero intercept (i.e., disposable income is $(1-t)wl$). The policy experiment is to raise τ_d while simultaneously reducing t to maintain budget balance. It is possible to show that $\rho < 1$ ($\rho > 1$) is associated with all three of the following: labor supply is upward (downward) sloping; the experiment reduces (increases) labor supply; and the benefit from controlling the externality, although equal for all types *in utils*, rises less (more) than proportionately with income when measured *in dollars*, which implies that the reform's overall effect is to increase (reduce) redistribution.⁴ If $\rho = 1$, labor supply is independent of the wage, the experiment has no effect on labor supply, and the reform is distribution neutral. Although merely a simple, special case, this model is close to common ones employed in the literature. Furthermore, it brings to light a central point of this article that is not immediately obvious, namely, that distributive effects from a reform as a whole can arise in subtle ways – even when the benefit of controlling externalities is additively separable – and that net distributive effects are implicitly related to effects on labor supply in precisely the manner suggested by step 2 of the decomposition. That is, additional labor supply distortion goes hand in hand with greater redistribution. Indeed, with weak separability of labor, Lemma 1 demonstrates that this is a general, hard-wired relationship.⁵

3.3. Nonseparability

The model also assumes that utility is weakly separable in labor, with the implication that neither the commodity mix nor the level of particular externalities affects labor effort – only the total subutility due to commodities and externalities matters. As is familiar from Corlett and Hague (1953), if some commodities are substitutes for or complements with labor (leisure), differential taxes or subsidies tend to be optimal. In the present setting, the benchmark is not uniformity but rather Pigouvian differentials. To illustrate the practical significance of nonseparability, West and Williams (2007) provide evidence that gasoline consumption is a leisure complement; accordingly, it is optimal to tax gasoline at a rate above the Pigouvian first-best level. (Observe that, even if gasoline consumption involved no negative externalities, it would still be optimal to tax gasoline above the Pigouvian level (zero) – that is, to impose a positive differential tax on gasoline – because of its complementarity with leisure.)

It is also possible that externalities themselves interact with labor effort. For example, reducing some negative externalities may make outdoor recreation more attractive, in which case it would be optimal to stop short of the Pigouvian first best. If, instead, central city ambience is

⁴The demonstration is omitted because it is somewhat tedious.

⁵Further corroboration of the often hidden but nevertheless real relationship between distribution and distortion is provided by Allgood and Snow (1998), who show that much of the disparity in estimates of the marginal cost of funds for public projects and of redistribution is due to implicit differences in the extent of redistribution that is assumed to take place.

improved, which in turn makes work more appealing, then going beyond full internalization would raise welfare.

Much prior literature on optimal internalization of externalities, like the present paper, assumes at least weak separability with regard to environmental benefits but, in contrast, does not assume that private consumption is weakly separable. Accordingly, some differences in results are due to this distinction. As explained, relaxing separability could change the optimum in either direction, depending on whether, say, a negative-externality-causing good was a leisure complement (as in West and Williams's study of gasoline consumption) or a leisure substitute. Since prior work more often suggests that optimal taxes are below Pigouvian levels, this difference could be due in part to an implicit assumption that dirty goods are leisure substitutes.⁶

4. Conclusion

This article examines the problem of controlling externalities in a model with commodity taxes and subsidies as the policy instrument, an income tax (initially set arbitrarily), and a continuum of individuals whose earning abilities differ. It considers consumption tax reforms that involve adjustments to the income tax schedule that in aggregate are distribution neutral (taking into account both the direct effects of reforms, including on commodity tax revenue and externalities, and the effects of the income tax adjustments). Such reforms allow one to set aside distributive concerns and focus on Pareto improvements. Moreover, as demonstrated, there is also no effect on labor supply in the baseline case with separable labor. Which reforms are optimal, therefore, depends on what might be viewed as narrow, traditional considerations of efficiency.

The main result is that setting commodity taxes and subsidies equal to marginal harms and benefits is optimal. Distribution-neutral implementation results in a Pareto improvement. An extension shows that, with some additional assumptions, marginal reforms in the direction of the naive first-best also result in Pareto improvements (with distribution-neutral implementation). These results suggest that basic first-best principles of internalization provide a useful focal point that clarifies thinking and assists in the examination of reforms.

The analysis here employs a number of different types of assumptions. The additional assumptions for partial reforms pertain to second-best problems that can arise when externalities are not fully internalized. (Thus, they were unnecessary in demonstrating Proposition 1.) These qualifications are orthogonal to the concerns about distribution and about labor supply distortion due to the income tax that have been the focus of most second-best work on the control of externalities: In a world with individualized lump-sum taxes that achieve the desired distribution without any distortion, the pertinent qualifications indicate that moving commodity taxes in the direction of (but not all the way to) the first best might not raise welfare.

The assumption of distribution-neutral implementation was examined at length. It was

⁶On the possible interaction between consumption as a whole and labor supply, see Chetty (2006).

explained that one can analyze the merits of non-distribution-neutral reforms in two steps: first, with distribution held constant, and then, separately, with regard to redistribution considered in isolation from the change in the regulation of externalities. The costs and benefits of changing redistribution do not depend on the particular reform that instigated such change. As a consequence, it probably does not make sense for every policy assessment involving externalities to incorporate an independent analysis of the equity-efficiency tradeoff that results when the extent of redistributive income taxation is adjusted. However, as explained, some existing work analyzes reforms that implicitly involve increased redistribution, the distortionary effect of which is not usually distinguished from the welfare consequences inherently associated with the regulation of externalities.

The inquiry at hand is also related to the literature on optimal commodity taxation in the presence of an optimal income tax (but without externalities). That research has identified numerous other qualifications to Atkinson and Stiglitz's (1976) uniformity result.⁷ One that was examined above involved the assumption of weak separability of labor. For the most part, these qualifications are independent of whether the setting involves externalities, and thus most do not go to the core of the second-best literature on the regulation of externalities in the presence of income taxation. The analysis of the present paper thus can be understood as demonstrating a substantial unity between the pure commodity tax problem and that involving externalities, one that can be taken advantage of in subsequent research, including that on qualifications to first-best prescriptions.

⁷See, for example, Cremer, Pestieau, and Rochet (2001), Marchand, Pestieau, and Racionero (2003), Mirrlees (1976), Naito (1999), and Saez (2002, 2004). Additional qualifications are addressed in the literature on the optimal provision of public goods. See, for example, Ng (2000) and Slemrod and Yitzhaki (2001).

Appendix

Proof of Proposition 2: Begin with an income tax $T(y)$ and an initial set of τ_i 's that are not first-best Pigouvian taxes and subsidies. Recall that the commodity tax reform to be considered is a move to $\tau_i(\alpha) = \alpha h_i + (1-\alpha)\tau_i$, for all i , where $\alpha \in [0, 1]$. That is, all commodity taxes and subsidies are moved the same fraction of the distance toward the simple first-best Pigouvian scheme under which $\tau_i = h_i$, for all i .⁸ Furthermore, we are considering a reform $d\tau_i(\alpha)/d\alpha$, for all i , evaluated at $\alpha = 0$. And the income tax schedule is $T(y, \alpha)$.

The logic of Lemma 1 depends on the fact that T is adjusted to keep utility constant at every income level (and on weak labor separability), but not on the particular sort of reform under consideration, so we can proceed straightaway to an analogue of Lemma 2. Let $R(\alpha)$ be the revenue raised under commodity tax regime $\tau_i(\alpha)$ where the income tax schedule $T(y, \alpha)$ is adjusted as described – i.e., $R(\alpha)$ is equal to the left side of expression (3) making the appropriate substitutions for the commodity tax vector and the income tax schedule.

Lemma 2': If $\partial^2 v / \partial e_j \partial x_i = 0$, for all i, j , then $dR(\alpha)/d\alpha|_{\alpha=0} > 0$.

Proof: Begin by using the budget constraint (1) to form the Lagrangian for each type of individual.

$$(A1) \quad L = u + \lambda \left[y(w) - T(y(w), \alpha) - \sum (p_i + \tau_i(\alpha)) x_i(y(w)) \right].$$

Because $T(y(w), \alpha)$ is defined to keep utility constant as α changes, it follows that $dL/d\alpha = 0$ for each type. Furthermore, from the analogue to Lemma 1, $dl/d\alpha = 0$. Therefore,

$$(A2) \quad \sum u_v v_{x_i} \frac{dx_i}{d\alpha} + \sum u_v v_{e_i} \frac{de_i}{d\alpha} = \lambda \left[T_\alpha + \sum (p_i + \tau_i(\alpha)) \frac{dx_i}{d\alpha} + \sum x_i \frac{d\tau_i(\alpha)}{d\alpha} \right].$$

Rearranging terms, expression (A2) is equivalent to

$$(A3) \quad T_\alpha + \sum x_i \frac{d\tau_i(\alpha)}{d\alpha} + \sum \tau_i(\alpha) \frac{dx_i}{d\alpha} = \sum \frac{u_v v_{x_i}}{\lambda} \frac{dx_i}{d\alpha} + \sum \frac{u_v v_{e_i}}{\lambda} \frac{de_i}{d\alpha} - \sum p_i \frac{dx_i}{d\alpha}.$$

Observe that the left side of expression (A3) is the contribution to the change in revenue by a given type of individual, a point to which we will return momentarily. For the moment, we will concentrate on the right side.

⁸No particular normalization is assumed for the initial commodity tax vector, although a particular one is chosen for the vector toward which the reform moves. Any resulting change in the overall level of commodity taxes, however, is immaterial because it is compensated by the income tax adjustment (implicitly, as part of the stage 1 income tax adjustment, following the presentation in Lemma 2).

An individual's first-order conditions for the choice of the x_i 's are $(u_v/\lambda)(\partial v/\partial x_i) = p_i + \tau_i(\alpha)$, for all i . Use this to substitute in the first term on the right side of (A3), then combine that term with the third term, to find that the right side equals

$$(A4) \sum \tau_i(\alpha) \frac{dx_i}{d\alpha} + \sum \frac{u_v v_{e_i}}{\lambda} \frac{de_i}{d\alpha}.$$

At this point, we can integrate both sides of expression (A3) over the population – first replacing the right side with expression (A4) – and evaluate at $\alpha = 0$ to yield

$$(A5) \left. \frac{dR}{d\alpha} \right|_{\alpha=0} = \sum (\tau_i - h_i) \frac{de_i}{d\alpha} = - \sum \frac{d\tau_i(\alpha)}{d\alpha} \frac{de_i}{d\alpha}.$$

To explain expression (A5), first, for the left side, simply compare the left side of expression (A3), when integrated, to the left side of expression (3) – or to expression (5), which gives the change in revenue for a discrete reform. After the first equal sign, the first component in moving from (A4) makes use of the definition in expression (2) of the e_i 's as the integral of the x_i 's over the population, and the second component in moving from (A4) uses the definition of the h_i 's from expression (4). Finally, in moving to the right side of the second equal sign, the substitution follows from the definition of the proportional commodity tax reform, $\tau_i(\alpha) = \alpha h_i + (1-\alpha)\tau_i$, for all i , evaluated at $\alpha = 0$.

The right side of expression (A5) must be positive, on account of what is referred to as the (Hicksian) compensated law of demand, which follows directly from application of the weak axiom of revealed preference. To explain how this property applies to expression (A5), first note that $d\tau_i(\alpha)/d\alpha$ in the present setting is the change in the price for commodity i faced by consumers (since producer prices are taken to be constant). Next, observe that $de_i/d\alpha$ is the change in aggregate consumption for commodity i . The reform under consideration, however, differs from the pure textbook case of a change in prices because of the presence of externalities. Specifically, as noted in subsection 3.1, it is possible in principle that the change in levels of externalities could affect individuals' consumption choices. However, we have assumed that $\partial^2 v/\partial e_j \partial x_i = 0$, for all i, j . Finally, the compensated law of demand requires that the change in price vector under consideration be utility-compensated, but we have defined $T(wl, \alpha)$ to accomplish just that. Therefore, the compensated law of demand holds for each individual and thus it holds when aggregated over the population. This, in turn, implies that the right side of expression (A5) is positive.⁹ ■

⁹To elaborate, consider the case of a discrete change. The right side of expression (A5) equals (in discrete form) $-\Delta \mathbf{p} \Delta \mathbf{x}$, where \mathbf{p} and \mathbf{x} are the consumer price vector and aggregate consumption vector, respectively. The weak axiom of revealed preference can be applied twice to each individual: first to show that the bundle chosen in the original regime cannot be afforded in the reform regime, and second to show that the bundle chosen in the reform regime cannot be afforded in the initial regime. Summing these inequalities implies that $\Delta \mathbf{p} \Delta \mathbf{x}(y(w)) < 0$ for each individual, where $\Delta \mathbf{x}(y(w))$ is the change in the consumption vector of an individual with income $y(w)$.

The intuition underlying the last step in the argument of Lemma 2' is related to the discussion in subsection 3.1, which suggested that individuals' shifts in consumption would be toward those commodities for which taxes fall and away from those for which taxes rise. Although this phenomenon need not be true for every commodity (because commodities may be substitutes for or complements with other commodities), it must be true in the aggregate for each individual.

Further intuition may be gleaned from the intermediate expression in (A5). It states that the net revenue effect constitutes the difference between the commodity tax rate and marginal external harm, times the change in quantity, for each commodity.¹⁰ The latter, marginal harm, will equal the amount of the compensatory income tax adjustment on account of externalities. Thus, not surprisingly, the net revenue effect is the difference between what is collected on account of each commodity by the commodity tax on that commodity and what must be paid through the income tax due to changes in external harm associated with that commodity. This effect can be unambiguously signed because of the nature of the particular reform under consideration: The definition of the reform is used in the proof of the lemma in moving from the middle to the right side of expression (A5), the latter of which, as explained, must be positive as a consequence of individuals' maximizing behavior.

To complete the argument, all that remains is to rebate the surplus to produce a Pareto improvement. This step, however, involves the further complication noted in the text, requiring that we further assume that "more is better for everyone." At this point, we have established Proposition 2.

Proof of Proposition 3: Begin the analysis with any non-first-best commodity tax and subsidy system $\{\tau_1, \dots, \tau_n\}$, $T(y)$ and consider any commodity tax reform τ_i^* . As in the proofs of Propositions 1 and 2, define an intermediate regime with income tax schedule $T^\circ(y)$ that has the property that, if all individuals (of every type w) continue to choose the same level of labor effort as under the initial tax system, then their utility will be unchanged. Lemma 1, stating that labor effort is indeed unchanged, once again holds because, as previously noted, the analysis depends only on the manner in which $T^\circ(y)$ is constructed (and weak separability).

It is now straightforward to show that a version of Lemma 2 holds.

Lemma 2'': A commodity tax reform $\{\tau_1^*, \dots, \tau_n^*\}$, $T^\circ(y)$ raises more revenue than regime $\{\tau_1, \dots, \tau_n\}$, $T(y)$ if and only if it is an efficiency-increasing commodity tax reform (i.e., expression (9) holds).

Proof: For each of the budget constraints (1) for these two regimes, integrate them over the population of types, subtract one from the other, and rearrange terms, to yield:

¹⁰There is no revenue effect analogous to that in stage 2 of the proof of Lemma 2 because here we consider a marginal change, and the envelope theorem applies to individuals' consumption adjustments.

$$\begin{aligned}
 (A6) \left[\int T^\circ(y(w))f(w)dw + \sum \tau_i^* \int x_i^\circ(y(w))f(w)dw \right] \\
 - \left[\int T(y(w))f(w)dw + \sum \tau_i \int x_i(y(w))f(w)dw \right] = \\
 \sum p_i \int x_i(y(w))f(w)dw - \sum p_i \int x_i^\circ(y(w))f(w)dw.
 \end{aligned}$$

The first bracketed term on the left side of expression (A6) is total revenue under the intermediate regime, and the second bracketed term is total revenue under the initial regime. The right side is positive if and only if expression (9), the definition of an efficiency-increasing commodity tax reform, holds. ■

Accordingly, we can complete the argument as before, which establishes Proposition 3.

References

- Allgood, Sam and Arthur Snow. 1998. "The Marginal Cost of Raising Tax Revenue and Redistributing Income." *Journal of Political Economy*. 106:6, pp. 1246-73.
- Atkinson, Anthony B. and Joseph E. Stiglitz. 1976. "The Design of Tax Structure: Direct Versus Indirect Taxation." *Journal of Public Economics*. 6:1-2, pp. 55-75.
- Ballard, Charles L. and Steven G. Medema. 1993. "The Marginal Efficiency Effects of Taxes and Subsidies in the Presence of Externalities: A Computational General Equilibrium Approach." *Journal of Public Economics*. 52:2, pp. 199-216.
- Bovenberg, A. Lans and Lawrence H. Goulder. 2002. "Environmental Taxation and Regulation," in *Handbook of Public Economics, vol. 3*. Alan J. Auerbach and Martin Feldstein, eds. Amsterdam: Elsevier, pp. 1471-545.
- Casler, Stephen D. and Aisha Rafiqi. 1993. "Evaluating Fuel Tax Equity: Direct and Indirect Distributional Effects." *National Tax Journal*. 46:2, pp. 197-205.
- Chetty, Raj. 2006. "A New Method of Estimating Risk Aversion." *American Economic Review*. 96:5, pp. 1821-34
- Cordes, Joseph J., Eric M. Nicholson, and Frank J. Sammartino. 1990. "Raising Revenue by Taxing Activities with Social Costs." *National Tax Journal*. 43:3, pp. 343-56.
- Corlett, W.J. and D.C. Hague. 1953. "Complementarity and the Excess Burden of Taxation." *Review of Economic Studies* 21:1, pp. 21-30.
- Cremer, Helmuth, Firouz Gahvari, and Norbert Ladoux. 1998. "Externalities and Optimal Taxation." *Journal of Public Economics*. 70, pp. 343-64.
- Cremer, Helmuth, Pierre Pestieau, and Jean-Charles Rochet. 2001. "Direct versus Indirect Taxation: The Design of the Tax Structure Revisited." *International Economic Review*. 42, pp. 781-99.
- Fullerton, Don. 1997. "Environmental Levies and Distortionary Taxation: Comment." *American Economic Review*. 87:1, pp. 245-51.
- Goulder, Lawrence H. 2002. *Environmental Policy Making in Economies with Prior Tax Distortions*. Cheltenham, UK: Edward Elgar.
- Goulder, Lawrence H., Ian W.H. Parry, Robertson C. Williams III, and Dallas Burtraw. 1999. "The Cost-Effectiveness of Alternative Instruments for Environmental Protection in a Second-Best Setting." *Journal of Public Economics*. 72:3, pp. 329-60.
- Hylland, Aanund and Richard Zeckhauser. 1979. "Distributional Objectives Should Affect Taxes But Not Program Choice or Design." *Scandinavian Journal of Economics*. 81:2, pp. 264-84.
- Kaplow, Louis. 1996. "The Optimal Supply of Public Goods and the Distortionary Cost of Taxation." *National Tax Journal*. 49:4, pp. 513-33.
- Kaplow, Louis. 2004. "On the (Ir)relevance of Distribution and Labor Supply Distortion to Government Policy." *Journal of Economic Perspectives*. 18:4, pp. 159-75.
- Kaplow, Louis. 2006. "On the Undesirability of Commodity Taxation Even When Income Taxation Is Not Optimal." *Journal of Public Economics*. 90:6-7, pp. 1235-50.
- Kaplow, Louis. 2008. *The Theory of Taxation and Public Economics*. Princeton: Princeton University Press.

- Laroque, Guy. 2005. "Indirect Taxation Is Superfluous under Separability and Taste Homogeneity: A Simple Proof." *Economics Letters*. 87, pp. 141-44.
- Marchand, Maurice, Pierre Pestieau, and Maria Racionero. 2003. "Optimal Redistribution when Different Workers Are Indistinguishable." *Canadian Journal of Economics*. 36, pp. 911-22.
- Mirrlees, J.A. 1976. "Optimal tax theory: A synthesis." *Journal of Public Economics*. 6, pp. 327-58.
- Naito, Hisahiro. 1999. "Re-examination of Uniform Commodity Taxes under a Non-linear Income Tax System and Its Implication for Production Efficiency." *Journal of Public Economics*. 71:2, pp. 165-88.
- Ng, Yew-Kwang. 2000. "The Optimal Size of Public Spending and the Distortionary Cost of Taxation." *National Tax Journal*. 53:2, pp. 253-72.
- Pearce, David. 1991. "The Role of Carbon Taxes in Adjusting to Global Warming." *Economic Journal*. 101:407, pp. 938-48.
- Pigou, A.C. 1920. *The Economics of Welfare*. London: Macmillan.
- Pigou, A.C. 1928. *A Study in Public Finance*. London: Macmillan.
- Pirttilä, Jukka, and Matti Tuomala. 1997. "Income Tax, Commodity Tax and Environmental Policy." *International Tax and Public Finance*. 4, pp. 379-93.
- Saez, Emmanuel. 2002. "The Desirability of Commodity Taxation under Non-linear Income Taxation and Heterogeneous Tastes." *Journal of Public Economics*. 83:2, pp. 217-30.
- Saez, Emmanuel. 2004. "Direct or Indirect Tax Instruments for Redistribution: Short-run versus Long-run." *Journal of Public Economics*. 88:3-4, pp. 503-18.
- Sandmo, Agnar. 1975. "Optimal Taxation in the Presence of Externalities." *Swedish Journal of Economics*. 77:1, pp. 86-98.
- Slemrod, Joel and Shlomo Yitzhaki. 2001. "Integrating Expenditure and Tax Decisions: The Marginal Cost of Funds and the Marginal Benefit of Projects." *National Tax Journal*. 54:2, pp. 189-201.
- West, Sarah E. 2004. "Distributional Effects of Alternative Vehicle Pollution Control Policies." *Journal of Public Economics*. 88:3-4, pp. 735-57.
- West, Sarah and Roberton C. Williams III. 2007. "Optimal Taxation and Cross-price Effects on Labor Supply: Estimates of the Optimal Gas Tax." *Journal of Public Economics*. 91:3-4, pp. 593-617.