

ISSN 1045-6333

HARVARD

JOHN M. OLIN CENTER FOR LAW, ECONOMICS, AND BUSINESS

OPTIMAL CONTROL OF EXTERNALITIES IN THE PRESENCE OF INCOME TAXATION

Louis Kaplow

Discussion Paper No. 547

06/2006

Harvard Law School
Cambridge, MA 02138

This paper can be downloaded without charge from:

The Harvard John M. Olin Discussion Paper Series:
http://www.law.harvard.edu/programs/olin_center/

The Social Science Research Network Electronic Paper Collection:
http://papers.ssrn.com/abstract_id=#####

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. For other reforms, a simple efficiency test characterizing when a Pareto improvement is possible is offered. Qualifications and explanations for the substantial departure from results in previous work are also elaborated.

First Draft: June 2004; Revised Draft: June 2006

*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.

Optimal Control of Externalities in the Presence of Income Taxation

Louis Kaplow

© 2006 Louis Kaplow. All Rights Reserved.

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 1970's 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 may 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.

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 will be taken into account. And by focusing on Pareto improvements, concerns about distribution will be addressed.

Three principal results are presented in sections 2 through 4 respectively. First, 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. Second, subject to some additional assumptions, any marginal change in commodity taxes and subsidies that is proportionally in the direction of the first best can be implemented in a manner that produces a Pareto improvement. Third, a conceptually simple necessary and sufficient condition is offered that indicates which other commodity tax reforms make possible Pareto improvements. Section 5 discusses some of the principal qualifications and explains why the present results diverge so substantially from those in most prior work. Section 6 offers concluding remarks.

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 involves an adjustment to the preexisting income tax such that the reform as a whole (combining the income tax adjustment to the proposed modification of commodity taxes) is distribution neutral. This follows Hylland and Zeckhauser (1979), Kaplow (1996), and some subsequent work that focuses on public goods. (Although Kaplow (1996, 2004) has discussed the application of this approach to the regulation of externalities, he did not formally analyze the problem.) An important feature of this method is that, when a distribution-neutral income tax adjustment is 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 be optimal and is not limited to commodity tax reforms in the

neighborhood of their optimum.¹ In these two respects, the results are most analogous to those in Kaplow's (2006) recent article on commodity taxation. The analysis here, however, is substantially more involved because of the presence of externalities, as will be apparent in the demonstration of the two main propositions below.

2. Optimality of First-Best Pigouvian Taxes and Subsidies

2.1. Model

There are n commodities, x_1, \dots, x_n . Corresponding variables e_1, \dots, e_n denote the total consumption of each commodity by all individuals (as defined precisely below). Individuals choose levels of consumption and labor effort l to maximize the utility function $u(v(x_1, \dots, x_n, e_1, \dots, e_n), l)$, where v is a subutility function. The 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 .

Individuals earn income wl that depends on their wage (type) w , which has density $f(w)$. Commodity prices for goods x_i (which equal production costs measured in units of income) are p_i , assumed to be constant and greater than zero. There are a (nonlinear) income tax schedule $T(wl)$ and commodity taxes on each good 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 budget constraint can be written as

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

where summations throughout are from i equals 1 to n and the notation $x_i(wl)$ denotes the level of x_i chosen by an individual of type w and thus income of wl , where l implicitly refers to the labor effort of an individual of type w . Using the foregoing notation, we can now state that, for any i , $e_i = \int x_i(wl)f(w)dw$. Note that, when individuals choose the $x_i(wl)$ to maximize their utility, the levels of each of the x_i may depend on the levels of any of the e_i , which individuals take as given.

In the analysis that follows, use will be made of the indirect subutility function

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.

$V(\boldsymbol{\tau}, T, wl)$, 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 wl (as well as the e_i 's) are taken as given. That is, V is the maximized value of v subject to the budget constraint (1). Observe that since v depends only on the x_i 's (and the e_i 's), and since the constraint (1) depends only on the x_i 's, $\boldsymbol{\tau}$, T , and wl – 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 $U(\boldsymbol{\tau}, T, w, l) = u(V(\boldsymbol{\tau}, T, wl), l)$.

The government's budget constraint is

$$(2) \int \left[T(wl) + \sum \tau_i x_i(wl) \right] f(w) dw = R,$$

where R is a given revenue requirement.

Finally, it is useful to define the marginal external harm associated with any commodity as

$$(3) h_i = - \int \frac{u_v(w) v_{e_i}(w)}{\lambda(w)} f(w) dw,$$

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 , as one might expect. 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 system that does not consist of first-best Pigouvian taxes and subsidies and then to construct a system with first-best Pigouvian taxes and subsidies that makes everyone better off. To begin the construction of a Pareto-improving tax reform, start with an initial regime $\{\tau_1, \dots, \tau_n\}$, $T(wl)$ that is not a Pigouvian first best. 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 . (Throughout this article, “*” will be used to denote values associated with the reform under consideration, which need not involve moves to the optimum.) 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 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 can now define an intermediate income tax schedule $T^\circ(wl)$ at each income level so as to offset the net effect on utility. That is, we will examine an income tax schedule $T^\circ(wl)$ 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.² (Whether individuals will choose the same labor effort under this intermediate regime is the subject of Lemma 1, below.)

To be more precise, define $T^\circ(wl)$ such that $V(\tau, T, wl) = V(\tau^*, T^\circ, wl)$ for all wl (where, as noted, all individuals are assumed to exert the same level of labor effort). As just suggested, the reform from τ to τ^* will, for a given wl , change the value of subutility V . For each wl , one can set the tax schedule $T^\circ(wl)$ at the level – directly changing after-tax income for the stipulated level of before-tax income – that restores the original level of subutility. This tax adjustment, $T^\circ(wl) - T(wl)$, is simply the utility-compensating change 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 $\{\tau_1^*, \dots, \tau_n^*\}$, $T^\circ(wl)$ as under $\{\tau_1, \dots, \tau_n\}$, $T(wl)$.

Proof: It is straightforward to establish that $U(\tau, T, w, l) = u(V(\tau, T, wl), l) = u(V(\tau^*, T^\circ, wl), l) = U(\tau^*, T^\circ, w, l)$, for all w, l . The first equality follows by the definition of U . The second equality follows because $T^\circ(wl)$ is constructed such that $V(\tau, T, wl) = V(\tau^*, T^\circ, wl)$ for all wl . 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 (τ, T) must be the l that maximizes U in the intermediate regime (τ^*, T°) . ■

The next question is how revenue compares between the initial regime and the intermediate regime.

²It is not asserted at this point that the tax schedule $T^\circ(wl)$ is feasible; in Lemma 2, it will in fact be shown to generate a surplus. For purposes of the analysis, it is simply a hypothetical, intermediate construct. Only the final schedule, $T^*(wl)$, needs to be feasible.

Lemma 2: Regime $\{\tau_1^*, \dots, \tau_n^*\}$, $T^\circ(wl)$ (with first-best Pigouvian taxes and subsidies $\tau_i^* = h_i$, for all i) raises more revenue than does regime $\{\tau_1, \dots, \tau_n\}$, $T(wl)$.

Proof: From expression (2), the change in revenue is given by

$$(4) \int \left[(T^\circ(wl) - T(wl)) + \sum (\tau_i^* - \tau_i) x_i(wl) + \sum \tau_i^* (x_i^\circ(wl) - x_i(wl)) \right] f(w) dw.$$

The first term in the integrand is the change in income tax revenue for an individual of type w , who earns wl , from moving to the intermediate regime. The next two terms are the change in commodity tax revenue for the individual, which it is convenient to decompose: 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 (4), 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 $T(wl)$ 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 (4).
2. Individuals adjust their x_i 's but the e_i 's are imagined to be held constant, and the $T(wl)$ schedule is further modified 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 $T(wl)$ schedule is further modified.

Stage 1: Beginning with the first stage in which the x_i 's are held constant, the income tax schedule $T(wl)$ for each type w is raised (lowered) by an amount equal to the negative of the second term in expression (4), 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 precisely equals the decrease (increase) in commodity tax and subsidy revenue for each, it is clear that, when both are integrated over the population, 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 $T(wl)$ necessary to offset the direct effect on individuals' utility from their own changes in consumption. 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.

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. 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 $T(w)$ is adjusted upwards to reduce their disposable income. Of course, precisely this adjustment must be done in the intermediate regime because $T^\circ(w)$ is, by definition, constructed so as to hold everyone's utility constant. Accordingly, with regard to this second stage, it must be that the intermediate regime's income tax schedule $T^\circ(w)$ raises more revenue from individuals, on account of their utility benefit from adjusting consumption. (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 $T(w)$ schedule 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 third term in expression (4). 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 accomplished by further adjusting the $T(w)$ schedule. The aggregate effect on income tax revenues, summed over all commodities and integrated over all types, is:

$$(5) \sum \int \left[\frac{u_v(w)v_{e_i}(w)}{\lambda(w)} \left(\int (x_i^\circ(w) - x_i(w)) f(w) dw \right) \right] f(w) dw = - \sum h_i(e_i^\circ - e_i).$$

The left side of (5) 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 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 (3). Thus, income tax revenue changes in stage 3 by the amount given in expression (5).

Let us compare this revenue effect due to the compensatory adjustment in the income tax schedule to the third term in expression (4), the change in revenue through the commodity taxes and subsidies on account of individuals' changes in consumption.

$$\begin{aligned}
(6) \int \left[\sum \tau_i^* (x_i^\circ(wl) - x_i(wl)) \right] f(w) dw &= \sum h_i \int \left[(x_i^\circ(wl) - x_i(wl)) f(w) dw \right] \\
&= \sum h_i (e_i^\circ - e_i).
\end{aligned}$$

The left side of (6) 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 (6), 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 (6) reflects, as with expression (5), that the levels of the externalities are defined as the integral of consumption over the population as a whole.

Comparing expressions (5) and (6), 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.

Summarizing, the stage 1 income tax adjustment that compensates for changes in commodity tax payments at initial levels of consumption just offsets the corresponding change in commodity tax revenue, the stage 2 income tax adjustment that absorbs utility gains from adjustments in levels of consumption has a positive effect on revenue, and the stage 3 income tax adjustment that compensates for externalities just offsets the change in commodity tax revenue due to changes in consumption. Therefore, the reform of moving to the intermediate regime raises total revenue. ■

To complete the argument that a Pareto improvement is possible in moving to first-best Pigouvian taxes and subsidies, construct $T^*(wl)$ from $T^\circ(wl)$ 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 before, 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 (or expenditures). The former operation (raising the rebate) will ultimately exhaust the surplus; moreover, the effect from this operation is that all individuals' utility will increase. (Note 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.)

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

Proposition 1: If the h_i 's are constant, then for any tax system $\{\tau_1, \dots, \tau_n\}$, $T(w)$ for which commodity taxes and subsidies are not at first-best Pigouvian levels, there exists a tax system $\{\tau_1^*, \dots, \tau_n^*\}$, $T^*(w)$ with first-best Pigouvian taxes and subsidies that is strictly Pareto superior – i.e., $u^*(w) > u(w)$, 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 a step 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 could obtain the same conclusion as in Proposition 1 through 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 moment, the τ_i 's equal the h_i 's. For each increment of these marginal changes, the preceding arguments go through. 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 $\{\tau_1^*, \dots, \tau_n^*\}$, $T^*(w)$ will indeed be one that has first-best Pigouvian taxes and subsidies.³

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.

³The argument sketched in the text demonstrates the proposition – that in moving to first-best Pigouvian taxes and subsidies a Pareto improvement is possible – but it does not guarantee that the resulting commodity tax vector is a global optimum unless further restrictions are imposed on the subutility function v (such as a monotonicity requirement on the h_i 's). The analysis does establish, however, that every local optimum involves first-best Pigouvian taxes and subsidies and thus that the global optimum must as well.

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. Optimality of Proportional Reform toward First-Best Pigouvian Taxes and Subsidies

3.1. Introduction

Partial commodity tax reforms, unlike a move all the way to Pigouvian first-best commodity taxes and subsidies, raise a number of complications that will be considered in the next two sections. As will be elaborated momentarily, the problems all involve stage 3: When the new 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.

To consider the applicability of the foregoing analysis to a broader range of commodity tax reforms, we can review the steps of Proposition 1. To begin, the proof of Lemma 1 – that labor supply does not change in moving to the intermediate regime – has only to do with the manner in which the income tax schedule $T^\circ(w/l)$ is constructed. Hence, this lemma holds for any commodity tax reform.

In proving Lemma 2 – that the intermediate regime raises more revenue – stages 1 and 2 can also proceed as before. In stage 1, the income tax schedule is adjusted to offset the effect of the commodity tax reform, holding constant consumption and thus also the level of externalities. This stage clearly would be revenue neutral for any commodity tax reform. In stage 2, the income tax schedule is further adjusted to absorb the utility gain from individuals' adjusting their consumption vectors (holding the level of externalities constant). This too would produce a revenue gain for any commodity tax reform that changes some relative prices.

Stage 3, however, in which the income tax adjustment that fully compensates for externalities just equals the change in commodity tax revenue due to changes in consumption vectors, depends as noted on the fact that the new commodity tax vector just equals marginal harm for all commodities. For a partial reform, the new commodity tax vector will, by definition, not equal marginal harm for at least some commodities.

In this section, we will consider reforms that move proportionately toward the first-best commodity tax regime, thus eliminating certain sorts of substitution among commodities due to selective commodity tax reform (see section 4). It is natural to inquire whether a modified proof

of stage 3 can be constructed, based on the following intuition: Consider a good producing a negative externality for which the original commodity tax was too low relative to the first-best. The good becomes relatively more expensive as a consequence of the reform, which tends to discourage consumption. Hence, individuals will benefit from less exposure to the negative externality. This benefit will be absorbed by the income tax adjustment. However, the loss in commodity tax revenue will be less than the reduction in harm – because even after the reform the tax rate is below the level of the marginal harm – so the loss in commodity tax revenue will be less than the increase in income tax revenue from the stage 3 adjustment that absorbs the entire change in harm. Therefore, there is a net revenue gain on this account. Similar logic extends to the case in which the original tax was too high (the reduced tax, still above marginal harm, will induce increases in consumption of the good and thus of the corresponding externality, the harm from which, compensated by the income tax adjustment, will be more than covered by the rise in commodity tax revenue since the post-reform commodity tax rate exceeds the marginal harm). By the same reasoning, partial reforms would also tend to produce net revenue gains with regard to subsidies and positive externalities.

Taken together, it may appear that a proportionate commodity tax reform will tend to produce a net revenue increase at stage 3 of the construction of the proof of Lemma 2, which combined with the revenue gain from the income tax adjustment at stage 2 (and the nil aggregate revenue effect at stage 1) demonstrates that a Pareto improvement is possible. There are, however, possible complications due to cross-effects among commodities. Nevertheless, the foregoing intuition when viewed in aggregate turns out to be essentially correct, as will be shown momentarily, subject to one caveat: When externalities adjust in stage 3, these adjustments may produce further changes in the x_i 's, and these further changes need not adhere to the pattern just identified.

To illustrate how this problem can arise, 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). A partial, proportionate reform would provide a subsidy to flowers and a 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 complementary might exceed the contrary effect due to the tax on gasoline. Additionally, if this increase is sufficiently great, it is possible that the uncompensated increase in pollution (recall that the gasoline tax is assumed to be less than marginal harm) will involve a total cost that exceeds both the incremental benefit from more flowers and the stage 2 utility gain from individuals' adjusting their own consumption. To rule out this possibility in the analysis that follows, it will be assumed that individuals' consumption vectors are unaffected by the levels of externalities. That is, $\partial^2 v / \partial e_j \partial x_i = 0$, for all i, j .

3.2. Analysis

Begin with an income tax $T(wl)$ and an initial set of τ_i 's that are not first-best Pigouvian taxes and subsidies. 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 .⁴

As it turns out, some subtle complications arise due to the fact that the levels of the externalities differ at the two points of comparison when a discrete reform is under consideration. Hence, the analysis to follow will consider marginal reforms. Specifically, we will consider a 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 $T(wl, \alpha)$, which changes as a function of α so as to keep utility constant at every level of income and thus for every type. As explained in subsection 3.1, 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 reform $\tau_i(\alpha)$ where the income tax schedule $T(wl, \alpha)$ is adjusted as described – i.e., $R(\alpha)$ is equal to the left side of expression (2) making the appropriate substitutions for the commodity tax vector and the income tax schedule.

Lemma 2': If $\partial^2 v / \partial e_i \partial x_i = 0$, 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.

$$(7) \quad L = u + \lambda \left[wl - T(wl, \alpha) - \sum (p_i + \tau_i(\alpha))x_i(wl) \right].$$

Because $T(wl, \alpha)$ 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,

$$(8) \quad \sum u_v v_x \frac{dx_i}{d\alpha} + \sum u_v v_e \frac{de_i}{d\alpha} = \lambda \left[T_\alpha + \sum (p_i + \tau_i(\alpha)) \frac{dx_i}{d\alpha} + \sum x_i \frac{d\tau_i(\alpha)}{d\alpha} \right].$$

Rearranging terms, expression (8) is equivalent to

$$(9) \quad T_\alpha + \sum x_i \frac{d\tau_i(\alpha)}{d\alpha} + \sum \tau_i(\alpha) \frac{dx_i}{d\alpha} = \sum \frac{u_v v_x}{\lambda} \frac{dx_i}{d\alpha} + \sum \frac{u_v v_e}{\lambda} \frac{de_i}{d\alpha} - \sum p_i \frac{dx_i}{d\alpha}.$$

⁴No 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).

Observe that the left side of expression (9) 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.

An individual's first-order conditions for the choice of the x_i 's are $(u_i/\lambda)(\partial v/\partial x_i) = p_i + \tau_i(\alpha)$, for all i . Use this to substitute in the first term on the right side of (9), then combine that term with the third term, to find that the right side equals

$$(10) \sum \tau_i(\alpha) \frac{dx_i}{d\alpha} + \sum \frac{u_i v_e}{\lambda} \frac{de_i}{d\alpha}$$

At this point, we can integrate both sides of expression (9) over the population – first replacing the right side with expression (10) – and evaluate at $\alpha = 0$ to yield

$$(11) \left. \frac{dR}{d\alpha} \right|_{\alpha=0} = \sum (\tau_i - h_i) \frac{de_i}{d\alpha} = - \sum \frac{d\tau_i(\alpha)}{d\alpha} \frac{de_i}{d\alpha}$$

To explain expression (11), first, for the left side, simply compare the left side of expression (9), when integrated, to the left side of expression (2) – or to expression (4), which gives the change in revenue for a discrete reform. After the first equal sign, the first component in moving from (10) makes use of the definition of the e_i 's as the integral of the x_i 's over the population, and the second component in moving from (10) uses the definition of the h_i 's from expression (3). 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 (11) 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 (11), 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 $T(w/l, \alpha)$ to accomplish just that. Therefore, the 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 (11)

is positive.⁵ ■

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 (11). 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.⁶ 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 (11), 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 a further complication. Because externalities are not fully internalized, even after the marginal reform, individuals receiving rebates may increase consumption, among other things, of commodities that produce marginal harm in excess of the commodity tax payment. If the uncompensated externality is sufficiently great, individuals actually could be made worse off when everyone is given a rebate. This might be referred to as a "manna problem" – i.e., if manna falls from heaven to produce a government surplus, there may not exist a way to dispose of the surplus that benefits everyone. (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.) Rather than pursuing the subtleties of this possibility further, let us simply assume that "more is better for everyone," specifically, that if the government finds itself with a

⁵To elaborate, consider the case of a discrete change. The right side of expression (11) 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 $\Delta\mathbf{p}\Delta\mathbf{x}(w_l) < 0$ for each individual, where $\Delta\mathbf{x}(w_l)$ is the change in the consumption vector of an individual with income w_l .

⁶There 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.

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. Accordingly, we have established:

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 $\{\tau_1, \dots, \tau_n\}$, $T(w)$ 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 $T^*(w, \alpha)$ such that the marginal reform is strictly Pareto improving – i.e., $du(w, \alpha) / d\alpha|_{\alpha=0} > 0$, for all w .

The intuition behind Proposition 2 – the centerpiece being Lemma 2' – is as stated in subsection 3.1. For each individual (and thus in aggregate), consumption shifts are dictated, on average, by price changes. Thus, individuals will tend to purchase more of commodities for which marginal taxes and thus consumer prices fall – and these will be on goods for which negative externalities are presently overtaxed and goods for which positive externalities are undersubsidized. In each instance, the effect on others' utilities, which must be compensated for by the income tax adjustment, will be more than made up for by contrary changes in commodity tax revenue. Likewise, individuals will tend to purchase less of commodities for which marginal taxes and thus consumer prices rise, and here too there will be a net surplus due to the difference between changes in commodity tax revenue and compensation through income tax adjustments. Though this pattern need not hold for every particular commodity – e.g., an undertaxed commodity, facing a higher price, may still experience an increase in demand if it is, say, a complement to another commodity that is more in demand due to a tax decrease – it must hold for consumption as a whole. Thus, even though a partial proportional reform does not fully internalize externalities, it still produces shifts that, on average, are in a welfare-increasing direction.

Proposition 2 applies only to marginal proportional commodity tax reforms, but it gives a prescription for discrete reforms. As commodity taxes are adjusted, the levels of externalities will change, and hence so will the marginal harms, the h_i 's. Thus, a sequence of marginal proportional reforms will need to chase a moving target. But as long as the direction of reform continues to point toward the first best, changes will continue to be Pareto improving. Such improvements will no longer be possible only when the first best is reached, as indicated by Proposition 1.⁷

4. Other Reforms of Environmental Taxes and Subsidies

The environmental tax reforms encompassed by Proposition 2 move proportionally toward first-best Pigouvian taxes and subsidies. It is more difficult to characterize other reforms. Notably, moving a single tax or subsidy in the direction of the first-best need not be efficient.

⁷ It might appear that Proposition 2 implies Proposition 1, but this is not true because Proposition 2 requires two additional assumptions, both made necessary by the fact that the post-reform commodity taxes were not assumed to equal marginal harm.

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. The test involves a narrow, traditional concept of efficiency.

Efficiency-increasing commodity tax reform. For any tax system $\{\tau_1, \dots, \tau_n\}$, $T(wl)$, a commodity tax reform $\{\tau_1^*, \dots, \tau_n^*\}$ is efficiency increasing if, when combined with the income tax schedule $T^\circ(wl)$:

$$(12) \sum p_i \int x_i^\circ(wl) f(w) dw < \sum p_i \int x_i(wl) f(w) dw.$$

Expression (12) 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. Moreover, as will be elaborated further below, 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.

Begin the analysis with any non-first-best commodity tax and subsidy system $\{\tau_1, \dots, \tau_n\}$, $T(wl)$ 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 $T^\circ(wl)$ 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 $T^\circ(wl)$ is constructed (and weak separability).

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

Lemma 2'': A commodity tax reform $\{\tau_1^*, \dots, \tau_n^*\}$, $T^\circ(wl)$ raises more revenue than regime $\{\tau_1, \dots, \tau_n\}$, $T(wl)$ if and only if it is an efficiency-increasing commodity tax reform (i.e., expression (12) 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:

$$(13) \left[\int T^{\circ}(wl) f(w) dw + \sum \tau_i^* \int x_i^{\circ}(wl) f(w) dw \right] \\ - \left[\int T(wl) f(w) dw + \sum \tau_i \int x_i(wl) f(w) dw \right] = \\ \sum p_i \int x_i(wl) f(w) dw - \sum p_i \int x_i^{\circ}(wl) f(w) dw.$$

The first bracketed term on the left side of expression (13) 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 (12), the definition of an efficiency-increasing commodity tax reform, holds. ■

Accordingly, we can complete the argument as before, which establishes:

Proposition 3: If “more is better for everyone,” then beginning with any tax system $\{\tau_1, \dots, \tau_n\}$, $T(wl)$, for any efficiency-increasing commodity tax reform τ_i^* , there exists $T^*(wl)$ such that the reform regime is strictly Pareto superior – i.e., $u^*(w) > u(w)$, for all w .

In essence, Proposition 3 states that if a commodity tax reform increases efficiency in a traditional sense – i.e., 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.

The intuition behind this result parallels the analysis just presented: If a commodity tax reform increases efficiency in the intermediate regime, this means that fewer resources are needed for individuals to achieve their initial levels of utility. Because individuals thus do not need to spend as much (aside from commodity taxes) in the hypothesized intermediate regime, total tax collections must be greater for given income levels; the resulting surplus can be distributed in a manner that yields a Pareto improvement. Furthermore, when the income tax is adjusted in a manner that accomplishes this, labor supply effects do not interfere with the argument. 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).⁸

⁸It might appear that Proposition 3 is more general than the first two propositions and therefore subsumes them; however, this is not the case. Lemma 2'' is indeed easy to establish for an efficiency-increasing commodity tax reform, but the full analysis of Lemma 2 or Lemma 2', as the case may be, is necessary to show that expression (12) holds for reforms that move to, or in the direction of, first-best Pigouvian taxes and subsidies. In essence, demonstrating that total tax revenue is greater under the intermediate regime and demonstrating that fewer productive resources are utilized amount to the same thing. (One might describe expression (13) as a sort of

Unfortunately, Proposition 3 provides only modest illumination because, as the introduction to this section suggests, it will often be difficult to determine when the efficiency condition (12) holds when externalities are present. Nevertheless, the condition does provide some indication of what sorts of reforms are likely to be efficient, and these will tend to be reforms that improve (in a simple, direct sense) the treatment of externalities. Suppose, for example, that significant negative externalities are produced by some group of goods that are not (relatively) taxed nearly as much as first-best principles dictate. Raising taxes on these goods (or taxing their complements and subsidizing clean substitutes) will reduce their consumption. Although the direct effect is to reduce consumers' utilities, by assumption individuals will gain even more from the resulting reduction in externalities. Accordingly, it is plausible that individuals can be made as well off as before with fewer resources devoted to consumption.

The key point, as already noted, is that none of this analysis – however simple or complex it turns out to be in a given instance – depends on considerations of distribution or labor supply distortion if one can adjust the income tax schedule to hold distribution constant. It is this result about the independence of the problem of controlling externalities from that of using the income tax to achieve a desired tradeoff of distribution and labor supply distortion that remains true for a wide class of environmental tax reforms, subject to standard qualifications noted in the conclusion to follow.

5. Discussion

5.1. *Qualifications*

In characterizing partial reforms, two qualifications pertaining to externalities were introduced: (1) Changes in externalities do not affect consumers' allocations of their budgets across commodities (to avoid cases such as that in which subsidizing flowers increases the level of a positive externality that in turn increases driving, which involves a not-fully-internalized negative externality); and (2) a surplus is desirable, *ceteris paribus* (to avoid the “manna” problem that arises when individuals' increased consumption produces uncompensated externalities that exceed the direct benefits of the consumption increase). These assumptions both address second-best problems that can arise when externalities are not fully internalized. (Thus, they were unnecessary in demonstrating Proposition 1.) The practical significance of these complications is an empirical question that may be of some importance for particular reforms. It should be emphasized, however, that 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.

These results are also subject to other types of qualifications. Most notably, the assumption of weak separability of labor was employed in demonstrating that labor supply is unaffected by distribution-neutral implementation. The effect of relaxing this assumption is the

accounting identity when a regime is constructed with $T^\circ(wl)$, as in the proof of all three propositions.)

same as in other settings, notably that of the optimal commodity taxation problem with no externalities. 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. If, for example, internalizing a negative externality affecting beaches makes leisure more attractive, stopping short of the Pigouvian first-best would be optimal (or, in the alternative, taxing beach attendance may be attractive). If, instead, central city ambience is improved, which in turn makes work more attractive, then going beyond full internalization would raise welfare.

This understanding immediately suggests that the separability assumption, although potentially important for particular policy interventions, is entirely orthogonal to the two main ideas in the existing second-best literature. First, nonseparability has nothing to do with whether or not individuals' valuations are disproportional to income and, if so, in which direction, and thus does not parallel the concern about distribution. Regarding labor supply, the distortions identified in the literature, as discussed in the next subsection, are due primarily to the extent of redistribution implicit in the reforms analyzed.

Furthermore, this qualification has nothing in particular to do with externalities. For example, although West and Williams (2004) provide evidence that gasoline consumption is a leisure complement and hence should be taxed above the Pigouvian first-best level, their result would be equally true if gasoline consumption involved no 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.

The literature on optimal commodity taxation in the presence of an optimal income tax has identified numerous other qualifications to Atkinson and Stiglitz's (1976) uniformity result.⁹ For the most part, these qualifications also are orthogonal both to the second-best literature's arguments about distribution and labor supply and to whether the setting involves externalities.¹⁰

⁹See, 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).

¹⁰For example, if individuals' preferences differ, it may not be possible to identify Pareto improvements. To see this, suppose that there was a negative externality to chocolate ice cream but not vanilla. When internalizing this externality, the corresponding distribution-neutral income tax adjustment could succeed in keeping average utility constant at each income level, but chocolate lovers may still be losers (if their relative utility loss is less than the efficiency gain reflected in the rebate at the final step in the analysis), and likewise for vanilla lovers if the externality pertained to vanilla. When making actual policy changes in a world with millions of individuals, Pareto improvements rarely if ever exist. Nevertheless, economists focus on Pareto improvements in performing analysis because identifying Pareto-efficient reforms in models is likely to be useful (even if such subtle distributive effects as on chocolate versus vanilla lovers are of some relevance). In any event, the nonexistence of true Pareto improvements in practice is unrelated to the second-best literature's concerns with the overall distribution of income (i.e.,

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.

5.2. Prior Literature

As the foregoing discussion suggests, the substantial divergence between the present results and those in much of the second-best literature on externalities is not due to the simplifying assumptions employed here (some of which are shared with that literature in any event). Instead, the problems that the prior literature identifies arise in large part because that literature does not employ distribution-neutral finance. Some literature, notably concerning distributive effects, does not specify how the income tax might be adjusted; obviously, if distribution-neutral adjustments had been considered, distributive effects would have vanished.

Much of the literature that focuses on labor supply distortion does specify tax adjustments, most commonly proportional adjustments to a linear income tax (that is, the single marginal tax rate is adjusted to balance the budget). As a general matter, the result could be more or less redistribution, depending on the distributive incidence of the policy under consideration. More redistribution would generally be associated with increased distortion, and less redistribution with reduced distortion.

A coincidence of two factors explains the typical results in the literature. First, it is common to use representative-individual models. Literally, there are no distributive effects. Yet the concern for distribution is what motivates the stipulation that an income tax adjustment rather than a nondistortionary uniform lump-sum tax is the 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.

Second, the actual reforms considered have predominantly been ones that implicitly involve increased redistribution, which can explain why it is so commonly found that there will be additional distortion as environmental protection is tightened (which finding is used to state that the optimum involves pollution taxes below the first-best Pigouvian level). To illustrate this practice, note that some literature, such as that 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 reliance on 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 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

between the rich and poor) and with labor supply.

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.”

Other work, following Bovenberg and de Mooij’s (1994) seminal contribution, assumes a separable public benefit and upward sloping labor supply. But if one examines a simple case in which private consumption, the public benefit from environmental protection, and labor are additively separable, it can be shown that upward sloping labor supply is precisely the condition for increased redistribution to result from the proportional income tax increase that is stipulated in such work.¹¹ An additional source of increased redistribution appears in studies such as Goulder et al. (1999) that hold 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.

¹¹This point can be illustrated in a simplified variant of the present model. Suppose that individuals’ utility is given by $u(c,g,l) = v(c) + b(g) - z(l)$, where g is government expenditures to correct an externality, and individuals’ budget constraints are given by $c = wl(1-t)$. If one differentiates individuals’ first-order conditions for labor supply with respect to w and g , the results are

$$l_w = (1-R) \frac{(1-t)v'}{-\partial^2 u / \partial^2}$$

$$l_g = (R-1) \frac{wt_g v'}{-\partial^2 u / \partial^2},$$

where t_g is the change in t required to finance the increase in g and R is the coefficient of relative risk aversion, $-cv''/v'$. Because the second components of both expressions are positive, the necessary and sufficient condition for upward-sloping labor supply in such a model, $R < 1$, is the same as the condition for labor supply to fall as a consequence of the reform. Finally, it is straightforward to show that this is also the (necessary and sufficient) condition for redistribution to arise. Specifically, the value of the public benefit measured in dollars is b'/v' . To determine how this changes as a fraction of disposable income c , one takes the derivative of $(b'/v')/c$ with respect to c , which equals $(R-1)(b'/v'c^2)$. This result shows that upward-sloping labor supply, falling labor supply when the public good is increased, and a distribution of the public benefit that is less favorable to the rich than the additional tax they pay go hand in hand. The intuition behind this relationship between the slope of the labor supply curve and redistribution is that the benefit of the project, measured in dollars, rises less than proportionately with disposable income when individuals’ marginal utility does not fall very rapidly, and that property also indicates that income effects are not very large relative to substitution effects, so the labor supply curve is rising (and conversely). (See Chetty (forthcoming) for a more general exploration of the relationship between labor supply and the curvature of individuals’ utility functions.)

As this discussion of the literature documents, it is indeed the case that additional redistribution is implicitly stipulated to accompany the policies under consideration. This feature, in turn, can explain the common finding that increased environmental protection is associated with additional labor supply distortion.¹² Accordingly, it should not be surprising that, if one holds distribution constant as in the present analysis, the typically identified labor supply effects do not arise. Then one obtains the result that reforms involving the regulation of externalities will, in the basic case, generate a Pareto improvement if and only if they move in the direction of greater efficiency, narrowly construed.¹³

6. 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 reforms that are accompanied by 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. Which reforms are optimal, therefore, depends on what might be viewed as narrow, traditional considerations of efficiency.

The first result is that setting commodity taxes and subsidies equal to marginal harms and benefits is optimal. Distribution-neutral implementation results in a Pareto improvement.¹⁴

¹²In similar fashion, Allgood and Snow (1998) 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.

¹³In reacting to prior drafts of this article, some have hypothesized that separability assumptions might actually be responsible for most of the differences in the literature. Specifically, it has been suggested that one consider a model, like that sometimes employed in past work, in which the public benefit is fully separable but labor supply and consumption are not. This possibility can be analyzed using the model from the note 11. Utility would instead be expressed as $u(c, g, l) = v(c, l) + b(g)$. If one proceeds as before, the expressions for l_w and l_g each have an additional component of lv_{21} , which corresponds to the Corlett and Hague (1953) correction for possible substitutability or complementarity with leisure that was described in subsection 5.1. As explained there, this factor is orthogonal to the rest of the problem. Additionally, since it may go in either direction, it cannot readily explain how prior literature systematically produces results exhibiting increased distortion (whereas these results are explained by the implicit increase in redistribution, which in fact is present in such work).

¹⁴Proposition 1 per se only establishes that such a reform is a Pareto improvement. However, it also directly implies that it is the optimal commodity tax reform (aside from the problem of multiple local optima mentioned in note 3): If there were a candidate optimum that did not involve first-best Pigouvian taxes and subsidies, one could implement a reform that moves to the first-best taxes and subsidies and generates a Pareto improvement, contradicting

Second, marginal reforms in the direction of the naive first-best also result in Pareto improvements (with distribution-neutral implementation). Third, any general reform can produce a Pareto improvement if and only if it satisfies a simple efficiency condition, that the same level of utility can be achieved through less resources being expended to produce commodities.

These results suggest that basic first-best principles of internalization should constitute the benchmark for the analysis of externalities, despite the findings of the existing second-best literature. A benchmark in this setting means a focal point that clarifies thinking and guides analysis. As discussed in section 5, there are qualifications, the direction and empirical importance of which depend on the reform under consideration. These qualifications, however, are largely orthogonal to the central problem and thus can be understood and measured independently.

The same point holds true for the most fundamental assumption employed in the present analysis, that the reforms under consideration are accompanied by a distribution-neutral income tax adjustment. To see this, suppose instead, for example, that a particular reform under consideration is likely to be accompanied by increased redistribution. In that case, one could still analyze the reforms' merits without regard to distributive effects and then separately report the extent of redistribution. The increase in redistribution will be accompanied by additional distortion, but how much distortion thereby occurs does not as a first approximation depend on the particular reform that produced it. Nor does the appropriate social weight to be given to the increase in redistribution itself. 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. As explained in section 5, much existing work differs in that it does in essence incorporate a particular form of such an analysis: Specifically, increased redistribution is implicitly assumed to occur, the distortionary effects of that are measured and counted in assessing the underlying policy regarding externalities, and the effect of the change in redistribution on social welfare is effectively ignored. The present analysis suggests that this approach is both unnecessarily complex and misleading.¹⁵

that one was at an optimum.

¹⁵It also should be noted that the implicit assumption that increased environmental protection will as a general matter in fact be accompanied by increased redistribution does not seem particularly plausible. (Nor does the opposite assumption.) Absent some specific public choice argument, the most natural conjecture, if one were required, appears 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 benchmark; it would also be a more accurate characterization of political reality.

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 Ruud A. De Mooij. 1994. "Environmental Levies and Distortionary Taxation." *American Economic Review*. 84:4, pp. 1086-89.
- 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. Forthcoming. "A Bound on Risk Aversion Using Labor Supply Elasticities." *American Economic Review*.
- 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, Roberton 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.
- 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.
- 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. 2004. “Empirical Estimates for Environmental Policy Making in a Second-Best Setting.” NBER Working Paper 10330.