ECONSTOR
Make Your Publications Visible.

A Service of
ZBW
Leibniz-Informationszentrum Wirtschaft
Leibniz Information Centre for Economics

Scheuer, Florian; Werning, Iván

Working Paper

Mirrlees Meets Diamond-Mirrlees

CESifo Working Paper, No. 5830

Provided in Cooperation with:
Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Scheuer, Florian; Werning, Iván (2016) : Mirrlees Meets Diamond-Mirrlees, CESifo Working Paper, No. 5830, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:
https://hdl.handle.net/10419/130452

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

WWW.ECONSTOR.EU

Mitglied der Leibniz-Gemeinschaft

The international platform of Ludwig-Maximilians University's Center for Economic Studies and the Ifo Institute

CESifo Working Papers

www.cesifo.org/wp

Mirrlees Meets Diamond-Mirrlees

Florian Scheuer
Iván Werning

CESIFO WORKING PAPER NO. 5830
CATEGORY 1: PUBLIC FINANCE
MARCH 2016

An electronic version of the paper may be downloaded
- *from the SSRN website: www.SSRN.com*
- *from the RePEc website: www.RePEc.org*
- *from the CESifo website: www.CESifo-group.org/wp*

ISSN 2364-1428

CESifo
Center for Economic Studies & Ifo Institute

CESifo *Working Paper No. 5830*

Mirrlees Meets Diamond-Mirrlees

Abstract

We show that the Diamond and Mirrlees (1971) linear tax model contains the Mirrlees (1971) nonlinear tax model as a special case. In this sense, the Mirrlees model is an application of Diamond-Mirrlees. We also derive the optimal tax formula in Mirrlees from the Diamond-Mirrlees formula. In the Mirrlees model, the relevant compensated crossprice elasticities are zero, providing a situation where an inverse elasticity rule holds. We provide four extensions that illustrate the power and ease of our approach, based on Diamond-Mirrlees, to study nonlinear taxation. First, we consider annual taxation in a lifecycle context. Second, we include human capital investments. Third, we incorporate more general forms of heterogeneity into the basic Mirrlees model. Fourth, we consider an extensive margin labor force participation decision, alongside the intensive margin choice. In all these cases, the relevant optimality condition is easily obtained as an application of the general Diamond-Mirrlees tax formula.

Florian Scheuer
Department of Economics
Stanford University
USA - 94305-6072 Stanford CA
scheuer@stanford.edu

Iván Werning
MIT Department of Economics
50 Memorial Drive
USA - Cambridge MA 02139
iwerning@mit.edu

March 2016
For helpful comments and discussions we thank Jim Poterba, Casey Rothschild and Emmanuel Saez.

1 Introduction

The Mirrlees (1971) model is a milestone in the study of optimal nonlinear taxation of labor earnings. The Diamond and Mirrlees (1971) model is a milestone in the study of optimal linear commodity taxation. Here we show that the Diamond-Mirrlees model, suitably adapted to allow for a continuum of goods, is strictly more general than the Mirrlees model. In this sense, the Mirrlees model is an application of the Diamond-Mirrlees model.

This connection is obtained by reinterpreting and expanding the commodity space in the Diamond-Mirrlees model. Although only linear taxation of each good is allowed, nonlinear taxation can be mimicked by treating each consumption level as a different sub-good. The tax rate within each sub-good then determines the tax for each consumption level, which is equivalent to a nonlinear tax. The only slight complication with this approach is that it requires working with a continuum of goods and a continuum of agents. In particular, in the Mirrlees model, there is a nonlinear tax on the efficiency units y of labor supplied, which equals pre-tax earnings. Instead of treating y as the quantity for a single good, we model y as indexing the characteristic of a separate sub-good. Since any positive supply y is allowed, the set of sub-goods allowed is the positive real line.[1]

Both Mirrlees (1971) and Diamond and Mirrlees (1971) provided optimal tax formulas that have been amply studied, interpreted and employed. They provide intuition into the optimum and suggest the relevant empirical counterparts, or sufficient statistics, to the theory. In the case of Mirrlees (1971), the tax formula was employed and reinterpreted by Diamond (1998) and Saez (2001). In the case of Diamond and Mirrlees (1971), one can point to Mirrlees (1975) and especially Diamond (1975), who provided a many-person Ramsey tax formula.

We provide a connection by showing that the formula in Mirrlees can be derived from the formula in Diamond-Mirrlees in a straightforward manner. In particular, we start with a version of the general Diamond-Mirrlees formula, as provided by Diamond (1975), and show that it specializes to the Mirrlees tax formula in its integral form, as provided by Diamond (1998), Saez (2001) and others. A connection between the two formulas is natural given that we have shown that Diamond-Mirrlees' framework nests Mirrlees'. However, moving from the Diamond-Mirrlees formula to the Mirrlees formula is not entirely immediate because the optimality conditions in Mirrlees were developed for a continuous model and are therefore of a somewhat different nature. Fortunately, after a convenient change in variables, the connection between the two formulas is greatly simplified.[2]

[1] Piketty (1997) and Saez (2002b) consider a discrete "job" model with a finite number of jobs and associated earnings levels, deriving discrete optimal tax formulas, but they do not provide a connection with the Diamond-Mirrlees linear tax framework.

[2] Diamond and Mirrlees (1971) also briefly extend their analysis to consider parametric nonlinear tax sys-

One major benefit of demonstrating the connection between the two formulas is to offer a common economic interpretation. The Diamond-Mirrlees formula (as seen through the lens of the interpretation offered by Diamond, 1975) equates two sides, each with a simple economic interpretation. One side of the equation involves compensated cross-price elasticities, used to compute the change in compensated demand for a particular commodity when all taxes are increased proportionally across the board by an infinitesimal amount. The other side involves the demands for this particular commodity for all agents weighted by their respective social marginal utilities of income—which in turn combine welfare weights, marginal utilities of consumption, and income effects to account for fiscal externalities from income transfers.

The Mirrlees model, on the other hand, has at center stage two elements: the local compensated elasticity of labor and the local shape of the skill distribution or earnings distribution. It also involves social marginal utilities and income effects.

We show that the Diamond-Mirrlees formula reduces to the Mirrlees formula for two reasons. First, the cross-price derivatives for compensated demands in the Diamond-Mirrlees formula turn out to be zero, drastically simplifying one side of the equation. Thus, the Mirrlees model and its formula, when seen through the lens of Diamond-Mirrlees, constitutes the rare "diagonal" case where an exact "inverse elasticity rule" applies. Second, in our formulation, the commodity space is already specified as a choice over cumulative distribution functions for labor supply. As a result, the Diamond-Mirrlees formula directly involves the distribution of labor. In the basic Mirrlees model, this translates directly to the distribution of earnings.

Our results highlight a deep connection between two canonical models in public finance and provide an alternative interpretation of the Mirrlees formula in terms of the Diamond-Mirrlees formula. Another benefit of attacking the nonlinear tax problem this way is that the Mirrlees formula is shown to hold under weaker conditions than commonly imposed. For example, a general, possibly nonlinear, production function is a key feature of the Diamond-Mirrlees model, whereas the baseline Mirrlees setup involves a simple linear technology. Finally, our approach provides a powerful and simple tool to explore extensions of the standard Mirrlees model. We consider four such extensions.

The original Mirrlees model is cast in a one-shot static setting, with a single consumption and labor supply decision. Thus, the model abstracts from dynamic considerations as well as uncertainty. Our first extension shows how to incorporate lifecycle features. In particular, each individual faces a time-varying productivity profile, but pays taxes based on current

tems and derive an optimality condition. However, due to its abstract nature, they do not develop it in detail, and this parametric approach has not been followed up by the literature. This is not our starting point, nor our ending point. We work with the Diamond-Mirrlees linear tax formula, which has been developed and applied in detail, and use it to derive the Mirrlees non-parametric optimal tax formula.

income. This is in line with present practice, where taxes are assessed annually, despite individuals' earnings varying significantly over their lifecycle.

Due to the lack of age and history dependence in taxation, the optimal annual income tax schedule solves a severely constrained—and hence complex—planning problem under the standard mechanism design approach. The connection to the Diamond-Mirrlees model, however, allows us to derive a formula for the optimal annual tax that is similar to the standard static one with two differences: it features a local Frisch elasticity of labor supply, which plays a similar role as the compensated elasticity in the static Mirrlees model, and a new additional term that captures lifetime effects. In the simple case with quasilinear preferences, the latter effect vanishes and the tax formula coincides in format with that of the static model. However, even in this special case, our analysis highlights that the application of this formula requires taking into account that welfare weights are a function of lifetime differences in earnings, rather than current annual differences in earnings. Since inequality in lifetime earnings is smaller than inequality in annual earnings, the benefits for redistribution are smaller, for a given welfare function.

Second, we incorporate human capital investment into the lifecycle framework. Notably, individuals can choose an education level that will affect their lifetime productivity profile, and individuals may differ in both their costs of this investment and its effect on productivities. We show that the formula for the optimal annual tax from the lifecycle model with exogenous productivity profiles extends to this case, with the only difference that the novel effect now also captures the effect of taxes on human capital. More generally, the novel effect can be interpreted as a "catch all" for any additional margins that affect individuals' lifetime productivity profiles and budget constraints.

Third, the Diamond-Mirrlees model allows for general differences across agents. In contrast, the benchmark Mirrlees model adopts a single dimension of heterogeneity satisfying a single-crossing assumption. Using our approach, we show how one can easily extend the Mirrlees analysis to allow for rich multi-dimensional forms of heterogeneity. Our results show that the standard formula holds using simple averages of the usual sufficient statistics, elasticities and marginal social utilities. This generalizes Saez (2001), who allowed for heterogeneity in his perturbation analysis of the asymptotic top marginal tax rate, and Jacquet and Lehmann (2015), who obtain a result under additively separable preferences based on an extended mechanism design approach that incorporates the constraint that a single income tax schedule cannot fully separate agents when there are multiple dimensions of heterogeneity. The Diamond-Mirrlees approach provides a very straightforward way of dealing with rich forms of heterogeneity.

Fourth, the Mirrlees model only considers an intensive margin of choice for labor supply. Other analyses have incorporated an extensive participation margin, following the

seminal contribution by Diamond (1980). We show that the Diamond-Mirrlees model also nests these models, including the pure extensive-margin model in Diamond (1980) and the hybrid intensive-extensive models considered in Saez (2002b) and Jacquet, Lehmann and Van der Linden (2013). Indeed, we consider a slightly more general specification and use the Diamond-Mirrlees approach to obtain the relevant tax formula. As in the lifecycle extensions, the demand system with both an intensive and extensive margin is no longer diagonal with zero cross-elasticities, and optimal tax formulas are no longer an application of the "inverse elasticity rule." Despite this fact, the demand system still retains an elementary structure and, thus, delivers relatively simple and easily interpretable tax formulas.

Of the four extensions we offer, we believe the first two to be the most significant, in the sense that, to the best of our knowledge, they have no precedent in the literature. Moreover, a mechanism design approach, while probably feasible, would be relatively contrived in these contexts. Our other two extensions, adding heterogeneity and the extensive margin, have clear precedents in the literature, as already mentioned. Although our assumptions and results differ in details, we believe the main benefit of covering these two extensions is to illustrate the benefits of revisiting them with a Diamond-Mirrlees perspective. Indeed, our method is able to handle these extensions with ease while highlighting the economics of each extension, summarized by the impact different assumptions have on the resulting demand system.

This paper is organized as follows. Section 2 introduces both the Diamond and Mirrlees (1971) and the Mirrlees (1971) model, and Section 3 shows how the Mirrlees model can be understood as a special case of Diamond-Mirrlees. Section 4 presents the optimal tax formulas from both models and Section 5 shows how to obtain the Mirrlees formula directly from the one in Diamond-Mirrlees. All the extensions are collected in Section 6 and Section 7 concludes. Most formal derivations are relegated to the appendix.

2 Diamond-Mirrlees and Mirrlees Models

We begin by briefly describing both frameworks, starting with the Diamond and Mirrlees (1971) linear tax model and then turning to the nonlinear tax model in Mirrlees (1971). To make the two models comparable, we extend Diamond-Mirrlees to a case with a continuum of goods and agents.

2.1 Diamond-Mirrlees

A set of agents is indexed by $h \in H$. Agent h has utility

$$u^h(x^h)$$

over net demands $x \in X$. Technology is represented by

$$G(\bar{x}) \leq 0, \tag{1}$$

where $\bar{x}$ is the aggregate of x^h over H. Agents face a linear budget constraint

$$B(x^h, q) = I$$

with consumer prices q. Diamond-Mirrlees consider both the case where one allows a nonzero lump-sum tax or transfer, $I \neq 0$, as well as the case where it is ruled out, by imposing $I = 0$. We shall be more interested in the natural case where the lump-sum tax is permitted.

The objective of the planner is to maximize a social welfare function

$$W(\{u^h\}),$$

where $\{u^h\}$ collects the utilities obtained by each agent $h \in H$.

Under the simplest interpretation in Diamond-Mirrlees, all production is controlled by the planner. The planner sets prices q and possibly the transfer I (if I is not required to be zero) and agents select their net demands x^h to maximize utility subject to their budget constraint. The planner is constrained by the fact that these demands must be consistent with the technological constraint (1).

As is well understood, whenever technology is convex and has constant returns to scale, this planning problem can be reinterpreted as allowing private production by firms to maximize profits at some producer prices $p \neq q$. In other words, one can implement the previous planning problem by allowing decentralized private production. Taxes are then equal to the difference between consumer and producer prices, $t = q - p$.

Finite agents and goods. In Diamond and Mirrlees (1971), there is a finite population $H = \{1, 2, \ldots, M\}$, so we can write

$$\bar{x} = \sum_{h=1}^{M} x^h.$$

There is a finite set of goods indexed by $i \in \{1, 2, \ldots, N\}$, so that $x^h = (x_1^h, x_2^h, \ldots, x_N^h)$. The budget constraints are then

$$q \cdot x^h = \sum_{i=1}^{N} q_i x_i^h = I,$$

where $q = (q_1, q_2, \ldots, q_N)$. Note that some elements of the vector x may be positive while others negative, with the interpretation that negative entries represent a surplus or supply (i.e. selling in the market), while positive entries represent deficits or demand (i.e. buying in the market).

Continuum of agents and goods. A simple extension to allow for a continuum of agents and commodities is as follows. Let there be a measure of agents μ_h over a set H. The set of goods is allowed to be infinite. Each agent h consumes a signed measure $x^h \in X$ over these goods and is subject to a linear budget constraint $B(x, q) = I$ as before, where q are consumer prices. This is a natural generalization. With a finite set of goods, choosing a measure is equivalent to selecting the quantity of each good.

2.2 Mirrlees

Agents are indexed by their productivity θ with c.d.f. $F(\theta)$ on support Θ. They have utility function

$$U(c, y; \theta),$$

over consumption c and effective labor effort y with the single-crossing condition that the marginal rate of substitution function

$$MRS(c, y; \theta) = -\frac{U_y(c, y; \theta)}{U_c(c, y; \theta)}$$

is strictly decreasing in θ (so higher θ types find it less costly to provide y). The canonical specification in Mirrlees (1971) is $U(c, y; \theta) = u(c, y/\theta)$ for some utility function over c and actual effort y/θ. Agents are subject to the budget constraint

$$c(\theta) \leq y(\theta) - T(y(\theta)) \equiv R(y(\theta)).$$

where T is a nonlinear income tax schedule and R is the associated retention function. The tax on consumption is normalized to zero without loss of generality.

Technology is defined by the resource constraint

$$\int_{\Theta} c(\theta)\, dF(\theta) \leq \int_{\Theta} y(\theta)\, dF(\theta).$$

Thus, the different efficiency units of labor are perfect substitutes.

We will consider a generalization of technology to allow for imperfect substitution. Any choice over $y(\theta)$ induces a distribution over y which we denote by its associated cumulative distribution function (c.d.f.) $H(y)$. We consider the resource constraint to be

$$\int_{\Theta} c(\theta)\, dF(\theta) \leq G(H), \tag{2}$$

for some production function G. The special case mentioned earlier is when

$$G(H) = \int_0^{\infty} y\, dH(y) = \int_0^{\infty} (1 - H(y))\, dy$$

where the second expression follows by integration by parts.

The goal is to maximize a social welfare function $W(\{U(c(\theta), y(\theta); \theta)\})$. The planner sets a tax function T or, equivalently, a retention function R, and agents then select $c(\theta), y(\theta)$ to maximize utility subject to their budget constraint. The planner is constrained by the fact that these demands must be consistent with the technological constraint (2). Once again, under the simplest interpretation, all production is controlled by the planner. But the optimum can be decentralized with private production by firms under the usual conditions.

3 Mirrlees as a Special Case of Diamond-Mirrlees

The main difference between the Diamond-Mirrlees model and the Mirrlees model is that taxation is linear in the former, while it is allowed to be nonlinear in the latter. We will argue that this difference is only apparent: The Diamond-Mirrlees framework can accommodate nonlinear taxation and nest the Mirrlees model.

We present two ways of mapping one model into the other. The first is more straightforward and works directly with prices and taxes in levels. The second entails a change of variables to rewrite things in terms of marginal prices and taxes. This reformulation is more convenient to work with and is instrumental in relating the optimal tax formulas for both models in Section 5.

3.1 Levels Formulation

We now describe an economy in Diamond-Mirrlees that captures the Mirrlees problem. Agents are indexed by their skill type, so that $h = \theta$ and μ_h is defined by the c.d.f. over skills F. The commodity space is comprised of a single consumption and a continuum of

labor varieties indexed by $y \geq 0$.[3] Agent θ chooses a level for consumption $c \geq 0$ as well as a measure over labor varieties which can be summarized by a c.d.f. $H_\theta(y)$.

Technology is given by

$$\int_\Theta c(\theta) dF(\theta) \leq G(H)$$

where $H(y) = \int H_\theta(y) dF(\theta)$ is the aggregate c.d.f. over y.

Each agent faces a budget constraint

$$c \leq \int_0^\infty q(y) dH_\theta(y) + I, \tag{3}$$

where we have normalized the price of the consumption good c to unity. In the Diamond-Mirrlees notation and nomenclature, the tax on consumption has been normalized to zero, while the tax on variety y is given by $q(y) - p(y)$ for some $\{p(y)\}$ representing the derivatives of the production function G; in the standard Mirrlees model with linear technology $p(y) = y$.

Finally, we assume that agents must put full mass of unity on a particular value for y. This is a restriction on preferences, that is, on the space over which the utility function is defined. Specifically, we assume agents attain utility $U(c, y; \theta)$ when they consume c and put full mass on y; they would obtain $-\infty$ if they attempted to distribute mass over various points or put less than measure one. Thus, the measure is a c.d.f. $H_\theta(y)$ that is increasing and a step function, jumping from 0 to 1 at the chosen $y(\theta)$. This implies that the budget constraint specializes to

$$c \leq q(y) + I,$$

so that the $q(y)$-schedule is effectively the retention function in the Mirrlees model.

This completes the description of a particular Diamond-Mirrlees economy that nests the Mirrlees model. Under this formulation, the agents choose a measure $H_\theta(y)$ over y and a consumption level, subject to a budget constraint that is linear in these objects. Thus, standard consumer demand theory applies, with the price of good y as $q(y)$.

The only complication is that the natural quantities in this formulation are densities. In particular, if H_θ admits a density h_θ then the budget constraint becomes $c \leq \int_0^\infty q(y) h_\theta(y) dy + I$. However, in our Mirrlees formulation, we are actually ensuring that H_θ has no density representation.

A related point is that a small change in the price schedule can have discontinuous effects on demand. For example, suppose the production function is linear—so that $p(y) = y$—and start with no taxation—so that $q(y) = p(y) = y$. If the skill distribution is continuous, the

[3] See Section 5.4 for how this can be generalized to multiple consumption goods.

economy produces a density over y in aggregate. However, if one reduces $q(y_0)$ at a point y_0, by any positive amount, then a mass of agents shift towards y_0 (from the neighborhood around y_0). Conversely, if we raise $q(y_0)$ at y_0, then the density of agents at this point drops discontinuously to zero. Thus, aggregate demand behaves discontinuously with respect to these forms of price changes. To overcome both problems, we next reformulate the model using a change of variables.

3.2 A Reformulation

We have cast the Mirrlees model into the Diamond-Mirrlees framework. In this formulation, consumers face prices $q(y)$ and the planner can be seen as controlling taxes $t(y) = q(y) - p(y)$. We now discuss a simple reformulation in terms of the *marginal* price $q'(y)$ and *marginal* taxes $t'(y) = q'(y) - p'(y)$.

Integrating the budget constraint (3) by parts gives

$$c \leq \int_0^\infty q'(y)(1 - H_\theta(y))dy + \tilde{I} \tag{4}$$

where $\tilde{I} = q(0)(1 - H_\theta(0)) + I$.

Under this formulation, we can interpret $q'(y)$ and $1 - H_\theta(y)$ as the price and quantity, respectively, for good y. Agent θ chooses the quantity of each of these goods to maximize utility, taking into account any restriction dictated by preferences (his consumption feasibility set). Because the budget constraint is linear, standard consumer theory applies.

This overcomes the two problems discussed above. First, quantities are now always well-defined, even when the c.d.f. $H_\theta(y)$ admits no density representation. In particular, the demand by household θ for good y is

$$1 - H_\theta(y) = \mathbb{I}(y \leq y(\theta)),$$

where $y(\theta)$ is θ's preferred level of y. For later use, we will also denote by $1 - H_\theta^c(y)$ the compensated demand, i.e. holding utility unchanged for agent θ. Second, aggregate demand for good y

$$1 - H(y) = \int_0^\infty (1 - H_\theta(y))dF(\theta)$$

is no longer discontinuous with respect to focalized changes in the price schedule $q'(y)$. In addition to overcoming these two problems, this formulation in terms of marginal prices is more natural to link to the Mirrlees formula, which is expressed in terms of marginal tax rates.

4 Tax Formulas: Diamond-Mirrlees and Mirrlees

Here we briefly review the optimal tax formulas offered by both models. These formulas crystalize the main results from these theories, offer intuition and provide the starting points for empirical applications. Readers familiar with these formulas can skip this section.

4.1 Diamond-Mirrlees

The first-order optimality conditions for the Diamond-Mirrlees model can be expressed in various useful and insightful ways. There are several different expressions, depending on whether or not one expands the effects of tax changes on tax revenues, whether one uses the compensated or uncompensated demands, and how one groups the different terms. The one we find most useful is due to Diamond (1975) (and the related analysis in Mirrlees, 1976).

In the case of finite goods and agents, the formula for good i is

$$\frac{\partial}{\partial \tau}\left(\sum_{h=1}^{M} x_i^{c,h}(q+\tau t)\right)\Bigg|_{\tau=0} = \sum_{h=1}^{M} \hat{\beta}^h x_i^h. \tag{5}$$

The left-hand side is the change in the demand for good i due to a compensated change in prices in the form of a proportional increase in all taxes.[4] This left-hand side (or the same expression divided by aggregate demand for the good) is often interpreted as an index of "discouragement," which measures by how much the tax system lowers the demand for the good, captured by substitution effects of compensated demands.

The right-hand side is the demand weighted by "social marginal utilities from income," defined as

$$\hat{\beta}^h = \beta^h - 1 + \frac{\partial}{\partial I}\left(\sum_{j=1}^{N} t_j x_j^h(q, I)\right), \tag{6}$$

Here, β^h is the marginal social benefit of increasing consumption for agent h. The next term, -1, captures the resource cost of providing the extra income to increase consumption in the absence of taxes. The final term corrects the latter for fiscal externalities due to the presence of taxes: when transferring income to agent h, this agent will spend the income on goods that are taxed, and income may flow back to the government. Thus, when the last term is

[4] The left-hand side is often written more explicitly as $\sum_h \sum_j t_j \frac{\partial}{\partial q_j} x_i^{c,h}$. However, this format is one step removed from its economic interpretation, i.e. the aggregate change in good i when all taxes rise proportionally and agents are compensated. In addition, this explicit format is specific to the finite good case, since the derivatives $\frac{\partial}{\partial q_j} x_i^{c,h}$ are not immediately well-defined with a continuum of goods, or requires reinterpretation. In contrast, the expression $\frac{\partial}{\partial \tau}\left(\sum_{h=1}^{M} x_i^{c,h}(q+\tau t)\right)\Big|_{\tau=0}$ is closer to the interpretation and carries over immediately to the continuum case.

positive, the social cost is less than 1. Overall, the social marginal utility of income may be positive or negative. Indeed, when the poll tax I is available, then the optimality condition for I implies that the average of the social marginal utilities of income across agents must be zero: $\sum_{h=1}^{M} \hat{\beta}^h = 0$.

Thus, this version of the Diamond-Mirrlees optimal linear tax formula states that the discouragement (or encouragement) of a good through the tax system should be in proportion to the welfare-weighted level of that good. Goods that are consumed more by those with high $\hat{\beta}^h$ should be encouraged and vice versa. In the context of labor supply (a negative entry in the x-vector), if agents who work and earn more have lower $\hat{\beta}^h$ then labor should be discouraged.

4.2 Mirrlees

Just as in the case of Diamond-Mirrlees, the Mirrlees optimality conditions can be expressed in a number of equivalent forms. There are two main choices. First, the conditions can be expressed in differential or in integral form. Second, they can be expressed using the primitive skill distribution or using the implied distribution of earnings. Finally, one can derive the optimality conditions by various methods: applying the Principle of Optimality by setting up a Hamiltonian, setting up a Lagrangian and taking first-order conditions, or using local perturbation arguments. For concreteness, we shall focus on the version of the optimality condition that is expressed in integral form and using the earnings distribution, rather than the skill distribution, as in Saez (2001). However, we show in Appendix A how to connect to other versions.

We first introduce the relevant elasticities that play a role in the formula. Consider the agent problem

$$y(\xi, I) \in \arg\max_{y} U(q(y) - \xi y + I, y; \theta),$$

which allows us to measure the behavioral effect of a small increase in the marginal tax rate (captured by ξ) and income effects (in response to I) starting from a given schedule $q(y)$. Then we define the uncompensated tax elasticity and the income effect by

$$\varepsilon^u(y) = -\left.\frac{\partial y}{\partial \xi}\right|_{\xi=I=0} \frac{q'(y)}{y} \qquad \text{and} \qquad \eta(y) = -\left.\frac{\partial y}{\partial I}\right|_{\xi=I=0} q'(y), \tag{7}$$

with the compensated elasticity obeying the Slutsky relation

$$\varepsilon^c(y) = \varepsilon^u(y) + \eta(y).$$

Note that $\varepsilon^c \geq 0$; moreover, $\eta \geq 0$ if "leisure" $-y$ is a normal good. We will assume that the

initial schedule q is such that the optimum is continuous in τ and I. This is equivalent to assuming that the agent's optimum is unique.

The optimality condition in the Mirrlees model can then be expressed as

$$\frac{T'(y)}{1-T'(y)}\varepsilon^c(y)yh(y) = \int_y^\infty (1-\beta_{\tilde{y}})\, dH(\tilde{y}) + \int_y^\infty \frac{T'(\tilde{y})}{1-T'(\tilde{y})}\eta(\tilde{y})dH(\tilde{y}), \tag{8}$$

at all points where no bunching takes place.[5] Here, H denotes the c.d.f. for labor supply y, h is its associated density, and β_y is the social marginal utility from consumption. Equation (8) must be supplemented with a boundary condition, stating that the right-hand side of (8) is equal zero at the lower bound of the support for $H(y)$.

A version of equation (8) was derived in Saez (2001, equation (19), p. 218) employing a perturbation argument where, starting from the optimal tax schedule, marginal tax rates are increased by a small amount $d\tau$ in the small interval $[y, y+dy]$. Then the left-hand side of condition (8) corresponds to the substitution effect of those individuals in $[y, y+dy]$ due to the increase in the marginal tax rate in this interval. The first term on the right-hand side captures the mechanical effect net of welfare loss from the reform, because increasing the marginal tax rate in $[y, y+dy]$ implies that everyone above y pays $d\tau dy$ in additional taxes, each unit of which is valued by the government $1-\beta_y$. Finally, the second term on the right-hand side captures the income effect of this additional tax payment for everyone above y. Setting the sum of the substitution, mechanical and income effects equal to zero at the optimum yields equation (8).

One minor difference is that our definitions for the elasticities capture changes starting from a baseline where the agent faces a nonlinear price schedule q; the nonlinearity could be due to a nonlinear tax, $t(y)$, or a nonlinear producer price, $p(y)$, or both. In particular, the compensated elasticity is affected by the local curvature of q. These definitions seem natural in a nonlinear taxation context and help streamline optimal tax formulas. Indeed, our formula (8) involves the actual distribution of earnings, while the one in Saez (2001) uses instead a modified "virtual density," which is affected by the local curvature in the tax schedule.

Equation (8) can be interpreted as a first-order differential equation that implicitly characterizes the optimal tax schedule. Solving it yields the well-known ABC-formula

$$\frac{T'(y)}{1-T'(y)} = A(y)B(y)C(y) \qquad \text{with} \tag{9}$$

[5] However, somewhat informally, one can still interpret this equation as holding since dividing by $h(y)$ and noting that $\varepsilon^c = 0$ and $h(y) = \infty$ when there is bunching, the equation holds for any $T'(y)$.

$$A(y) = \frac{1}{\varepsilon^c(y)}, \quad B(y) = \frac{1-H(y)}{yh(y)}, \quad \text{and} \quad C(y) = \int_y^\infty (1-\beta_{\tilde{y}}) \exp\left(\int_y^{\tilde{y}} \frac{\eta(z)}{\varepsilon^c(z)} \frac{dz}{z}\right) \frac{dH(\tilde{y})}{1-H(y)}.$$

Both formulas, (8) and (9), are identical when there are no income effects, $\eta = 0$, as in the related formulas derived by Diamond (1998).[6]

5 Tax Formulas: From Diamond-Mirrlees to Mirrlees

We now show how to reach the Mirrlees formulas (8)–(9) starting from the Diamond-Mirrlees formulas (5)–(6). We do so by first translating the left-hand side of (5) into the Mirrleesian reinterpretation laid out in Section 3.2, followed by the right-hand side.

5.1 Left-Hand Side in Diamond-Mirrlees

The proportional change in all taxes underlying the left-hand side of equation (5) corresponds to changing the marginal consumer price schedule such that

$$q'(y;\tau) = q'(y) + \tau t'(y) \tag{10}$$

for all y, where $t'(y) = q'(y) - p'(y)$. Then the left-hand side of equation (5) is simply equal to

$$\frac{\partial}{\partial \tau}(1 - H^c(y;\tau))\Big|_{\tau=0}, \tag{11}$$

where $H^c(y;\tau)$ is the aggregate distribution of y under price schedule $q'(y;\tau)$, and the superscript c indicates that the compensated responses are required when we vary τ.

When τ is increased infinitesimally from zero, the compensated response for each agent is, by (10) and the definition of the compensated elasticity, a decrease in y equal to

$$-t'(y)\,\frac{\partial y^c(\theta;\tau)}{\partial \tau}\Big|_{\tau=0} = t'(y)\frac{\varepsilon^c(y)y}{q'(y)}.$$

Since each agent lowers y, this produces a shift in the distribution H to the left. At a particular point $(y, H(y))$, the horizontal shift equals precisely $\frac{t'(y)}{q'(y)}\varepsilon^c(y)y$. Equation (11), however, demands the implied vertical shift. To translate the horizontal shift into the vertical shift requires multiplying by the slope of H, that is, the density h. We conclude that the left-hand side of (5) equals

$$\frac{\partial}{\partial \tau}(1 - H^c(y;\tau))\Big|_{\tau=0} = \frac{t'(y)}{q'(y)}\varepsilon^c(y)yh(y). \tag{12}$$

[6] Diamond (1998) expressed the formula as a function and the primitive skill distribution, rather than the implied earnings distribution.

The formal derivation is contained in Appendix A.

Equation (12) reveals that the left-hand side of the Diamond-Mirrlees formula simplifies drastically when applied to the Mirrlees setting: the relevant response for y only depends on the variation in the marginal tax rate $t'(y)$ at y, and not on the variation in the entire schedule in (10). In other words, the Mirrlees model constitutes the rare diagonal case where compensated cross-price elasticities of demand are zero and only the own-price elasticity matters. This coveted case is often highlighted in the commodity tax literature for it implies an "inverse elasticity rule."

5.2 Right-Hand Side in Diamond-Mirrlees

The analog, with a continuum of goods, of the right-hand side of equation (6) in conjunction with (5) is

$$\begin{aligned}\int_0^\infty (1-H_\theta(y))\left(\beta_\theta - 1 - \frac{\partial}{\partial I}\int_0^\infty t'(z)\,(1-H_\theta(z;I))\,dz\right)dF(\theta)\\ = \int_{\theta(y)}^\infty \left(\beta_\theta - 1 - \frac{\partial}{\partial I}\int_0^{y(\theta;I)} t'(z)dz\right)dF(\theta)\\ = \int_{\theta(y)}^\infty \left(\beta_\theta - 1 - t'(y(\theta))\frac{\partial y(\theta;I)}{\partial I}\right)dF(\theta), \quad (13)\end{aligned}$$

where $\theta(y)$ denotes the inverse of $y(\theta)$.[7] Substituting $\partial y(\theta;I)/\partial I = -\eta(y)/q'(y)$ into (13) and changing variables from θ to $y = y(\theta)$ yields

$$-\int_y^\infty (1-\beta_{\tilde{y}})dH(\tilde{y}) + \int_y^\infty \frac{t'(\tilde{y})}{q'(\tilde{y})}\eta(\tilde{y})dH(\tilde{y}), \quad (14)$$

with a slight abuse of notation to write β_y for $\beta_{\theta(y)}$.

5.3 Putting it Together

Equating (12) and (14) yields

$$-\frac{t'(y)}{q'(y)}\varepsilon^c(y)yh(y) = \int_y^\infty (1-\beta_{\tilde{y}})dH(\tilde{y}) - \int_y^\infty \frac{t'(\tilde{y})}{q'(\tilde{y})}\eta(\tilde{y})dH(\tilde{y}). \quad (15)$$

[7] Recall that $y(\theta)$ is monotone increasing by the single-crossing assumption. Moreover, note that the effect of I on $1-H_\theta$ has the opposite sign in (4) compared to (6), which is why the income effects enter with a minus sign in (13).

To translate this into the Mirrlees model with a nonlinear tax over pre-tax earnings $p(y)$, we set $q(y) = p(y) - T(p(y))$ and recall that $t(y) = q(y) - p(y)$, so that $t(y) = -T(p(y))$ and

$$\frac{t'(y)}{q'(y)} = -\frac{T'(p(y))}{1 - T'(p(y))}, \tag{16}$$

which upon substitution gives precisely the Mirrlees formula (8). Thus, applying the Diamond-Mirrlees formula (5) to the Mirrlees model immediately delivers the condition which Saez (2001) obtained from his perturbation argument, equating, at the optimum, the substitution effect on the left-hand side with the sum of the mechanical effects net of welfare loss and the income effects on the right-hand side.

In Appendix A, we show how to solve (15) to obtain the ABC-formula

$$-\frac{t'(y)}{q'(y)} = \frac{1}{\varepsilon^c(y)} \frac{1 - H(y)}{yh(y)} \int_y^\infty (1 - \beta_{\tilde{y}}) \exp\left(\int_y^{\tilde{y}} \frac{\eta(z)}{\varepsilon^c(z)} \frac{dz}{z} \right) \frac{dH(\tilde{y})}{1 - H(y)}, \tag{17}$$

which upon the same substitution of the relationship (16) delivers (9). This concludes the derivation of the Mirrlees formulas (8)–(9) from the Diamond-Mirrlees formulas (5)–(6).

5.4 Discussion

Tax formula in terms of skills. Mirrlees (1971) expresses the optimal tax formula in terms of the primitive skill distribution instead of the implied distribution of labor supply. In Appendix A, we show that there is a direct link between the two, and we demonstrate how to rewrite formulas (8) and (9) as a function of $F(\theta)$ rather than $H(y)$. The rewritten formulas characterize the marginal tax rate $\tau(\theta) = t'(y(\theta))$ faced by type θ. We emphasize though that this is one step removed from the formulas that naturally result from an application of the Diamond-Mirrlees framework, which are in terms of y. In particular, when expressing them in terms of elasticities, the formulas in terms of θ require the use of different elasticity concepts in general.[8]

Technology and tax instruments. An advantage of deriving the Mirrlees optimal tax formula from the Diamond-Mirrlees formula is that it allows for a general structure of the production side of the economy. A general, possibly nonlinear, production function is a key feature of the Diamond-Mirrlees model. In contrast, the baseline Mirrlees setup involves a simple linear technology. As our derivation makes clear, the Mirrlees optimal tax formula

[8] It is also possible to rewrite the optimal tax formulas as a function of the implied distribution of earnings $p(y)$, again requiring different elasticity concepts in general when $p(y) \neq y$ (see for example Scheuer and Werning (2015) for the required elasticity adjustments in the context of superstar effects).

holds for any production function $G(H)$. In other words, the result of Diamond and Mirrlees (1971) that their tax formula is independent of whether technology is linear or not extends to the Mirrlees formula.

We note, however, that this result relies on the availability of sufficiently rich tax instruments. Consider, for instance, a two-sector economy with technology $G(H_1, H_2)$, so a given level of labor supply y can have different effects depending on whether it enters in sector 1 or 2. Then the Diamond-Mirrlees framework, and hence our analysis, requires sector-specific tax instruments $t_1(y)$ and $t_2(y)$. In the absence of this, with only a single tax schedule $t(y)$, the Mirrlees formula generally gets modified to reflect general equilibrium effects on redistribution (Stiglitz, 1982, and Rothschild and Scheuer, 2013). Moreover, production efficiency is not necessarily optimal with restricted tax instruments (Guesnerie, 1998, and Naito, 1999). In our framework with sufficiently targeted instruments, however, these considerations play no role.[9]

Multiple consumption goods. In line with the original Mirrlees model, we have considered a single consumption good. Since the Diamond-Mirrlees model naturally allows for any finite number N of commodities, it is straightforward, however, to extend our analysis to multiple consumption goods. With linear taxes on each of the consumption goods (normalizing one of them to zero), an application of conditions (5) and (6) would immediately deliver formulas for (i) the optimal linear commodity tax rates in the presence of the optimal nonlinear labor income tax schedule and (ii) marginal tax rates of the optimal nonlinear income tax schedule in the presence of the optimal commodity taxes.

Such formulas have been derived in the literature using standard mechanism design (see for example Mirrlees, 1976, and Jacobs and Boadway, 2014) or variational approaches (e.g. Christiansen, 1984, and Saez, 2002a). A crucial feature of these formulas are conditional labor elasticities of the commodity demands, which measure how the demand for a consumption good c_i changes when labor y changes but after-tax income $q(y)$ is held fixed. When these cross-elasticities are zero, which holds under the weakly separable preference specification $U(u(c_1, ..., c_N), y; \theta)$ considered by Atkinson and Stiglitz (1976), it then immediately follows that (i) all commodity taxes are zero at the optimum, and (ii) the formula for the optimal marginal income tax rates is the same as derived here except for additional income effects on commodity demands on the right-hand side of (15).

[9] Saez (2004) links a discrete jobs model to the Diamond-Mirrlees model and discusses the implications of nonlinear production functions, also contrasting the results with Stiglitz (1982) and Naito (1999). However, Saez (2004) does not attempt to derive the Mirrlees optimal tax formula.

6 Four Extensions of the Mirrlees Model

In this section, we briefly consider four extensions that illustrate the power and ease of our approach. First, we extend the Mirrlees model to a lifecycle framework where workers pay an annual income tax, but productivity varies from year to year. Second, we incorporate human capital investments into this lifecycle framework, endogenizing individuals' lifetime productivity profiles. Third, we enrich the static Mirrlees model to allow for additional arbitrary dimensions of heterogeneity, without single-crossing assumptions. Fourth, we incorporate an extensive margin, alongside the intensive margin, for labor supply. The first two of these extensions are novel, and would be rather cumbersome to tackle with the usual mechanism design approach. The third and fourth extensions of the Mirrlees models have precedents in the existing literature.[10] While our assumptions and results are slightly different, the main benefit of our treatment of these extensions is to illustrate the power and ease of our approach based on the Diamond-Mirrlees formula.

6.1 Annual Taxation of Earnings in a Lifecycle Context

The original Mirrlees model is a one-shot static model: there is a single consumption good and a single labor supply choice. We now consider a simple dynamic extension, to incorporate a lifecycle choice for labor supply.

Suppose ex ante heterogeneity is indexed by $\theta \sim F(\theta)$ as before. Each individual faces varying productivities δ over her lifetime with conditional distribution $P(\delta|\theta)$. Individuals choose how much labor to supply for each δ, resulting in a schedule $y(\delta;\theta)$. The government sets a nonlinear income tax schedule, resulting in the retention function $q(y)$ for the income earned at any point in time (i.e. an "annual" tax without age or history dependence, as is the case in practice). Moreover, markets are complete, so individuals smooth consumption over their lifecycle respecting their budget constraint

$$c = \int_0^\infty q(y(\delta;\theta))dP(\delta|\theta).$$

Preferences are

$$U(c, Y; \theta)$$

with

$$Y = \int_0^\infty v(y(\delta;\theta), \delta)dP(\delta|\theta).$$

[10] See, among others, Saez (2001), Hendren (2014) and Jacquet and Lehmann (2015) for the third and Diamond (1980), Saez (2002b), Choné and Laroque (2011), Jacquet et al. (2013), Zoutman et al. (2013) and Hendren (2014) for the fourth extension.

We can interpret $v(y, \delta)$ as a measure of the instantaneous disutility from supplying effective labor y at a moment when productivity is δ, and as usual we assume v is such that we have single-crossing in δ. Then Y captures the total disutility from labor over the individual's lifetime. We do not require assumptions about the nature of ex ante heterogeneity in θ.

As before, we can think of each individual as choosing a distribution $H_\theta(y)$ over y. The only difference to before is that this distribution is no longer degenerate (i.e., no longer a step function). Using this insight, we show in Appendix B that our analysis carries over easily and leads to the following formula for the optimal annual tax $t(y)$:

$$-y\varepsilon^F(y)h(y)\left(\frac{t'(y)}{q'(y)}+\Lambda(y)\right)=\int_y^\infty\left(1-\bar{\beta}_{\tilde{y}}\right)dH(\tilde{y})-\int_y^\infty\bar{\eta}(\tilde{y})\frac{t'(\tilde{y})}{q'(\tilde{y})}dH(\tilde{y}). \tag{18}$$

This is very similar to the static formula (15) except for the following differences.

First, $\varepsilon^F(y)$ is a Frisch elasticity of labor supply that holds fixed $\lambda \equiv -U_c/U_Y$, i.e. the marginal rate of substitution between lifetime consumption and lifetime labor supply. This elasticity is purely local in the sense that it depends only on the local shape of the flow disutility function v and on the local shape of the annual tax schedule at y.

Second, there is a novel term on the left-hand side

$$\Lambda(y)=\int_\Theta\frac{1}{\lambda^c(\tau,\theta)}\frac{\partial\lambda^c(\tau,\theta)}{\partial\tau}\bigg|_{\tau=0}dF(\theta|y),$$

which captures precisely the lifetime effects on the compensated labor supply. In particular (and as explained in more detail in Appendix B), $\lambda^c(\tau,\theta,\bar{U})$ is defined such that

$$y^F(\tau,\lambda^c(\tau,\theta,\bar{U}))=y^c(\tau,\theta,\bar{U}),$$

where y^F is the Frisch labor supply, holding λ fixed, and y^c is the compensated labor supply, holding lifetime utility $\bar{U}$ fixed (and we dropped the argument $\bar{U}$ in λ^c). This captures global effects on labor supply and the interactions of labor supply across different "ages," i.e. across different values of δ. The effect $\Lambda(y)$ will generally depend on the entire tax schedule.

Third, on the right-hand side, $\bar{\eta}(y)$ is the average income effect and $\bar{\beta}_y$ the average social welfare weight at y (across θ).

We emphasize the following implications of the lifecycle framework for the shape of the annual income tax schedule. First, the fundamental welfare weights β_θ only vary with ex ante (i.e., lifetime) heterogeneity θ. Since there can be substantially less lifetime inequality than cross-sectional inequality at any given point in time (which is driven by δ in addition to θ), the average welfare weights at a given income β_y naturally vary less than in the static framework. An extreme case occurs when there is no ex ante heterogeneity, so all income

inequality is driven by the shocks δ. Hence, when viewed over their entire lifetimes, all individuals face the same distribution of these shocks, but the resulting cross-sectional income inequality at any point in time can be arbitrarily large. Then β_y is independent of y and optimal annual taxes are zero.

Second, in the special case of quasilinear lifetime preferences

$$U\left(c - \int_0^\infty v(y(\delta;\theta),\delta)dP(\delta|\theta);\theta\right),$$

the additional term on the left-hand side vanishes: $\Lambda(y) = 0$. Hence, in this case, the standard formula from the static setting fully extends to the annual tax in this much richer lifecycle framework.

6.2 Human Capital

It is easy to incorporate human capital investment in this lifecycle framework. In particular, suppose individuals choose an education level e before entering the labor market, which affects their productivity distribution $P(\delta|\theta, e)$. Their lifetime utility is $U(c, Y; \theta, e)$, which can capture costs of the education investment e in a general form (and note that these costs can differ across θ-types). Otherwise, the framework is identical to the one in the preceding subsection. As before, the government looks for the optimal annual nonlinear income tax schedule, or equivalently $q(y)$.[11]

As we show in Appendix B, all the results from the basic lifecycle framework go through. In particular, the optimal tax formula (18) still applies. The effect $\Lambda(y)$ takes the same form as before, but now also captures the effect of taxes on individuals' human capital choices.[12] In this sense, $\Lambda(y)$ can be interpreted as a "catch all" for any additional margins that affect individuals' lifetime productivity profiles and budget constraints.

6.3 More General Forms of Heterogeneity

An important advantage of approaching the Mirrlees model from the perspective of the Diamond-Mirrlees framework is that we can easily accommodate relatively general forms of heterogeneity, as we now show. General forms of heterogeneity are inherent to the structure in Diamond-Mirrlees. In contrast, the baseline Mirrlees setup allows for only one dimension of heterogeneity satisfying a single-crossing condition.

[11] This assumes that the human capital investment e is not taxed or subsidized directly. In other words, we abstract from exploring its optimal tax treatment (see e.g. Bovenberg and Jacobs, 2005, and Stantcheva, 2016, for recent work on this issue).

[12] In particular, Λ again vanishes if lifetime preferences take the quasilinear form $U(c, Y; \theta, e) = \tilde{U}(c - Y; \theta, e)$.

Returning to the static framework, suppose there are groups, indexed by ϕ and distributed according to c.d.f. $P(\phi)$ (and support Φ) in the population, with preferences

$$U(c, y; \theta, \phi).$$

We only require that the single-crossing property in terms of θ is satisfied among individuals with the same ϕ, i.e. $MRS(c, y; \theta, \phi)$ is strictly decreasing in θ for each ϕ. Apart from that, we can allow for arbitrary preference heterogeneity captured by ϕ. For example, ϕ could be from a finite set or a continuum, and it could be single- or multidimensional. This is in line with the Diamond-Mirrlees model, where h can index arbitrary differences across households.

In Appendix B, we show how to generalize the analysis from Section 5 to such a framework. The Mirrlees optimal tax formulas (8) and (9) go through when replacing the elasticities $\varepsilon^c(y)$ and $\eta(y)$ as well as the marginal social welfare weights β_y by their *averages* conditional on y. For example, $\varepsilon^c(y)$ is simply replaced by

$$\bar{\varepsilon}^c(y) = \mathbb{E}[\varepsilon^c(y, \phi)|y] = \int_\Phi \varepsilon^c(y, \phi) dP(\phi|y),$$

where $P(\phi|y)$ is the distribution of ϕ conditional on y (and analogously for $\bar{\eta}(y)$ and $\bar{\beta}_y$).[13]

6.4 Extensive-Margin Choices

Finally, we demonstrate how the Diamond-Mirrlees setting can easily incorporate extensive margin labor choices, generalizing the environments considered by Diamond (1980), Saez (2002b), Choné and Laroque (2011), and Jacquet et al. (2013) among others. We shall derive the resulting tax formula starting from the Diamond-Mirrlees formulas (5)–(6).

For simplicity, suppose individuals are characterized by two-dimensional heterogeneity (θ, φ) with preferences

$$V(c, y; \theta, \varphi) = \begin{cases} U(c, y; \theta) & \text{if } y > 0 \\ u(c; \theta, \varphi) & \text{if } y = 0. \end{cases}$$

Hence, heterogeneity in the φ-dimension only drives participation decisions but not inten-

[13] Using a perturbation approach, Saez (2001) derives this result for the asymptotic top marginal tax rate. Hendren (2014) provides a formula for the fiscal externality from changes to the nonlinear income tax schedule that depends on average elasticities at each income level, also based on a perturbation approach. Jacquet and Lehmann (2015) consider the same structure of heterogeneity as here and obtain this result for the optimal tax formula for the special case of additively separable preferences based on a mechanism design approach with pooling.

sive margin decisions conditional on θ.[14] In other words, preferences are the same as in Section 5 for strictly positive y but can exhibit a discontinuity at $y = 0$ that can be different across individuals with the same θ. Assuming that u is increasing in φ, this will lead individuals with high values of φ, for any given θ, to stay out of the labor market and choose $y = 0$, consuming the demogrand $q(0)$.

We show in Appendix B that an application of the Diamond-Mirrlees formulas in this case leads to the following simple modification of formula (8):

$$\frac{T'(y)}{1-T'(y)}\varepsilon^c(y)yh(y) = \int_y^\infty \left(1 - \bar{\beta}_{\tilde{y}} + \frac{T'(\tilde{y})}{1-T'(\tilde{y})}\eta(\tilde{y}) - \frac{T(\tilde{y})-T(0)}{q(\tilde{y})}\rho(\tilde{y})\right) dH(\tilde{y}), \tag{19}$$

where $\rho(y)$ is the participation elasticity at y, defined by

$$\rho(y) = \frac{\partial h(y)}{\partial q(y)}\frac{q(y)}{h(y)}\bigg|_{\{y(\theta)\}},$$

which is the percentage change in the density at y when $q(y)$ is increased by one percent, holding fixed the intensive margin choices of all individuals with $y > 0$ (i.e. holding fixed the $y(\theta)$-schedule). Moreover, $\bar{\beta}_y$ is the average social welfare weight on individuals who choose y.

As in the lifecycle extensions, the (compensated) demand system is no longer diagonal with an active extensive margin: The proportional change in all marginal tax rates underlying the left-hand side of (5) affects $1 - H(y)$ not just through the (compensated) intensive-margin response at y, but also through the (compensated) extensive-margin responses of all individuals with labor supply above y. Combining this with the pure income effect on the extensive margin from (6) leads to the additional term on the right-hand side of the optimal tax formula.[15]

A special case arises when only the extensive margin is active (see e.g. Diamond, 1980, and Choné and Laroque, 2011), in which case (19) reduces to

$$\frac{T(y)-T(0)}{q(y)} = \frac{1-\bar{\beta}_y}{\rho(y)},$$

[14] Such further heterogeneity could be easily incorporated as shown in the previous subsection. We focus on the extensive margin here.

[15] Saez (2002b) derives the equivalent of this formula for a discrete type setting and for the special case without income effects using a perturbation approach (the working paper version in Saez (2000) also provides a continuous types analogue). Jacquet et al. (2013) consider preferences with an additively separable participation cost (so $V(c, y; \theta, \varphi) = U(c, y; \theta) - \mathbb{I}(y > 0)\varphi$). For this special case of our environment, they derive the same formula as ours using perturbation and mechanism design approaches. Zoutman et al. (2013) and Hendren (2014) provide related formulas for the fiscal externality in the inverse optimum problem with both intensive and extensive margins.

i.e., an inverse elasticity rule similar to the pure intensive margin model considered so far, but in terms of the average tax rate and the participation elasticity.

7 Conclusion

This paper uncovered a deep connection between two canonical models in public finance and their optimal tax formulas. We find this connection is insightful and, thus, worthwhile in its own right. In addition, this line of attack on the nonlinear tax problem can easily allow for weaker conditions and extensions. We have provided four such extensions to illustrate the appeal of the Diamond-Mirrlees approach. We conjecture that this approach could be usefully applied in other settings as well.

References

Atkinson, A. and J. Stiglitz, "The Design of Tax Structure: Direct Versus Indirect Taxation," *Journal of Public Economics*, 1976, *6*, 55–75. 5.4

Bovenberg, L. and B. Jacobs, "Redistribution and Education Subsidies are Siamese Twins," *Journal of Public Economics*, 2005, *89*, 2005–2035. 11

Choné, P. and G. Laroque, "Optimal Taxation in the Extensive Model," *Journal of Economic Theory*, 2011, *146*, 425–453. 10, 6.4, 6.4

Christiansen, V., "Which Commodity Taxes Should Supplement the Income Tax," *Journal of Public Economics*, 1984, *24*, 195–220. 5.4

Diamond, P., "A Many-Person Ramsey Tax Rule," *Journal of Public Economics*, 1975, *4*, 335–342. 1, 4.1

— , "Income Taxation With Fixed Hours of Work," *Journal of Public Economics*, 1980, *13*, 101–110. 1, 10, 6.4, 6.4

— , "Optimal Income Taxation: An Example with a U-Shaped Pattern of Optimal Marginal Tax Rates," *American Economic Review*, 1998, *88* (1), 83–95. 1, 4.2, 6

— **and J. Mirrlees**, "Optimal Taxation and Public Production II: Tax Rules," *American Economic Review*, June 1971, *61* (3), 261–78. 1, 2, 2, 2.1, 5.4

Guesnerie, R., "Peut-on Toujours Redistribuer les Gains a la Spécialization et a l'Echange? Un Retour Pointille sur Ricardo et Heckscher-Ohlin," *Revue Economique*, 1998, *49*, 555–579. 5.4

Hendren, N., "The Inequality Deflator: Interpersonal Comparisons Without a Social Welfare Function," *NBER Working Paper 20351*, 2014. 10, 13, 15

Jacobs, B. and R. Boadway, "Optimal Linear Commodity Taxation Under Optimal Non-linear Income Taxation," *Journal of Public Economics*, 2014, *117*, 201–210. 5.4

Jacquet, L. and E. Lehmann, "Optimal Income Taxation when Skills and Behavioral Elasticities are Heterogeneous," *CESifo Working Paper 5265*, 2015. 1, 10, 13

_ , _ , and B. Van der Linden, "The Optimal Marginal Tax Rates with both Extensive and Intensive Responses," *Journal of Economic Theory*, 2013, *148*, 1770–1805. 1, 10, 6.4, 15

Mirrlees, J., "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 1971, *38*, 175–208. 1, 2, 2.2, 5.4

_ , "Optimal Commodity Taxation in a Two-Class Economy," *Journal of Public Economics*, 1975, *4*, 27–33. 1

_ , "Optimal Tax Theory: A Synthesis," *Journal of Public Economics*, 1976, *6*, 327–358. 4.1, 5.4

Naito, H., "Re-Examination of Uniform Commodity Taxes Under a Non-linear Income Tax and its Implication for Production Efficiency," *Journal of Public Economics*, 1999, *71* (165-188). 5.4, 9

Piketty, T., "La Redistribution Fiscale Face au Chômage," *Revue française d'économie*, 1997, *12*, 157–201. 1

Rothschild, C. and F. Scheuer, "Redistributive Taxation in the Roy Model," *Quarterly Journal of Economics*, 2013, *128*, 623–668. 5.4

Saez, E., "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses," *NBER Working Paper 7708*, 2000. 15

_ , "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, *68* (1), 205–29. 1, 4.2, 4.2, 5.3, 10, 13

_ , "The Desirability of Commodity Taxation under Non-linear Income Taxation and Heterogeneous Tastes," *Journal of Public Economics*, 2002, *83*, 217–230. 5.4

_ , "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses," *The Quarterly Journal of Economics*, 2002, *117* (3), 1039–1073. 1, 1, 10, 6.4, 15

_ , "Direct Or Indirect Tax Instruments For Redistribution: Short-Run Versus Long-Run," *Journal of Public Economics*, 2004, *88*, 503–518. 9

Scheuer, F. and I. Werning, "The Taxation of Superstars," *NBER Working Paper 21323*, 2015. 8

Stantcheva, S., "Optimal Taxation and Human Capital Policies over the Life Cycle," *Mimeo, Harvard University*, 2016. 11

Stiglitz, J., "Self-Selection and Pareto Efficient Taxation," *Journal of Public Economics*, 1982, *17*, 213–240. 5.4, 9

Zoutman, F., B. Jacobs, and E. Jongen, "Revealed Social Preferences of Dutch Political Parties," *Working Paper*, 2013. 10, 15

A Derivations

A.1 Derivation of Equation (12)

The the left-hand side of (5) with a continuum of goods equals

$$\frac{\partial}{\partial \tau}\int_0^\infty (1-H_\theta^c(y;\tau))dF(\theta)\bigg|_{\tau=0} = \frac{\partial}{\partial \tau}\int_{\theta^c(y;\tau)}^\infty dF(\theta)\bigg|_{\tau=0} = -\frac{\partial \theta^c(y;\tau)}{\partial \tau}\bigg|_{\tau=0} f(\theta(y)) \tag{20}$$

where the superscript c indicates compensated choices, $\theta(y;\tau)$ is the inverse of $y(\theta;\tau)$ with respect to its first argument, and $\theta(y)$ stands short for $\theta(y;0)$. We are using the fact that $y(\theta;\tau)$ is increasing in θ for any τ by the single-crossing condition.

The optimum for agent θ must satisfy the tangency condition

$$MRS(c,y;\theta) = q'(y;\tau) = q'(y) + \tau t'(y). \tag{21}$$

To compute the compensated demand, we use this equation with $c = e(v,y;\theta)$ where e is the inverse of U with respect to its first argument. To compute the uncompensated demand, we use the budget constraint $c = q(y) + I$. Differentiating (21) yields

$$\frac{\partial \theta^c(y;\tau)}{\partial \tau}\bigg|_{\tau=0} = \frac{q'(y)-p'(y)}{MRS_\theta} = \frac{t'(y)/q'(y)}{MRS_\theta/MRS}. \tag{22}$$

Moreover, observe that the density of y is given by $h(y) = f(\theta(y))\theta'(y)$. Again differentiating (21) for $\tau = 0$ yields

$$\theta'(y) = -\frac{MRS_c + \frac{MRS_y}{MRS} - \frac{q''}{q'}}{MRS_\theta/MRS}. \tag{23}$$

Finally, the elasticities defined in (7) can be obtained by differentiating

$$MRS(q(y) - \xi y + I, y; \theta) = q'(y) - \xi.$$

Hence,

$$\varepsilon^u(y) = \frac{-MRS_c + 1/y}{MRS_c + \frac{MRS_y}{MRS} - \frac{q''}{q'}}, \tag{24}$$

$$\eta(y) = \frac{MRS_c}{MRS_c + \frac{MRS_y}{MRS} - \frac{q''}{q'}}, \tag{25}$$

and

$$\varepsilon^c(y) = \varepsilon^u(y) + \eta(y) = \frac{1/y}{MRS_c + \frac{MRS_y}{MRS} - \frac{q''}{q'}}. \tag{26}$$

Using (26) in (23) yields

$$\theta'(y) = -\frac{1}{y\varepsilon^c(y)} \frac{1}{MRS_\theta / MRS}. \tag{27}$$

Substituting all this in (20), we obtain (12).

A.2 Derivation of Equation (17)

Define

$$\mu(y) \equiv -\frac{t'(y)}{q'(y)} \varepsilon^c(y) y h(y)$$

and write equation (15) as

$$\mu(y) = \int_y^\infty (1 - \beta_{\tilde{y}}) dH(\tilde{y}) + \int_y^\infty \frac{\eta(\tilde{y})}{\varepsilon^c(\tilde{y})} \frac{\mu(\tilde{y})}{\tilde{y}} d\tilde{y}.$$

Differentiating this yields

$$\mu'(y) + (1 - \beta_y) h(y) = -\frac{\eta(y)}{\varepsilon^c(y)} \frac{\mu(y)}{y}.$$

Integrating this ordinary first-order differential equation forward to solve for μ and using its definition yields (17).

A.3 Formulas in Terms of the Skill Distribution

Combine (20) and (22) and change variables from θ to $y(\theta)$ to write the left-hand side of (5) as

$$\frac{t'(y(\theta))/q'(y(\theta))}{MRS_\theta/MRS} f(\theta) = \frac{\tau(\theta)}{1-\tau(\theta)} \theta f(\theta)\chi(\theta),$$

where we defined $\tau(\theta) = T'(p(y(\theta)))$ and $\chi(\theta) = -\left(MRS_\theta \theta / MRS\right)^{-1}$. Using this together with (14) yields

$$\frac{\tau(\theta)}{1-\tau(\theta)} \theta f(\theta)\chi(\theta) = \int_\theta^\infty (1-\beta_{\tilde{\theta}}) dF(\tilde{\theta}) + \int_\theta^\infty \frac{\tau(\tilde{\theta})}{1-\tau(\tilde{\theta})} \eta(\tilde{\theta}) dF(\tilde{\theta}), \tag{28}$$

where we slightly abused notation to write $\eta(\theta) = \eta(y(\theta))$. This is the equivalent of (8) written in terms of θ. Defining the left-hand side of equation (28) as $\hat{\mu}(\theta)$, we can write it as

$$\hat{\mu}(\theta) = \int_\theta^\infty (1-\beta_{\tilde{\theta}}) dF(\tilde{\theta}) + \int_\theta^\infty \frac{\eta(\tilde{\theta})}{\tilde{\theta}\chi(\tilde{\theta})} \hat{\mu}(\tilde{\theta}) d\tilde{\theta}.$$

Observe that

$$\frac{\eta(\theta)}{\theta\chi(\theta)} = \frac{\eta(\theta)}{\varepsilon^c(\theta)} \frac{y'(\theta)}{y(\theta)},$$

where we used (25), (26) and (27) (and, again slightly abusing notation, wrote $\varepsilon^c(\theta) = \varepsilon^c(y(\theta))$). Using this and differentiating yields

$$\hat{\mu}'(\theta) + (1-\beta_\theta) f(\theta) = -\frac{\eta(\theta)}{\varepsilon^c(\theta)} \frac{y'(\theta)}{y(\theta)} \hat{\mu}(\theta).$$

Solving this forward yields

$$\frac{\tau(\theta)}{1-\tau(\theta)} = \frac{1}{\chi(\theta)} \frac{1-F(\theta)}{\theta f(\theta)} \int_\theta^\infty (1-\beta_{\tilde{\theta}}) \exp\left(\int_\theta^{\tilde{\theta}} \frac{\eta(s)}{\varepsilon^c(s)} \frac{dy(s)}{y(s)} \right) \frac{dF(\tilde{\theta})}{1-F(\theta)}, \tag{29}$$

which is the equivalent of (9) written in terms of θ.

B Extensions

B.1 Lifecycle Framework

Due to single-crossing in δ, $y(\delta;\theta)$ is increasing in δ, so

$$1 - H_\theta(y) = \int_0^\infty (1 - H_{\delta,\theta}(y)) dP(\delta|\theta)$$

where

$$1 - H_{\delta,\theta}(y) = \mathbb{I}(y \leq y(\delta;\theta)).$$

Hence, the left-hand side of (5) simply becomes

$$\begin{aligned}
\frac{\partial}{\partial\tau}(1 - H^c(y;\tau))\Big|_{\tau=0} &= \frac{\partial}{\partial\tau}\int_\Theta (1 - H^c_\theta(y;\tau))dF(\theta)\Big|_{\tau=0} \\
&= \frac{\partial}{\partial\tau}\int_\Theta\int_0^\infty (1 - H^c_{\delta,\theta}(y;\tau))dP(\delta|\theta)dF(\theta)\Big|_{\tau=0} \\
&= \frac{\partial}{\partial\tau}\int_\Theta\int_{\delta^c(y;\theta,\tau)}^\infty dP(\delta|\theta)dF(\theta)\Big|_{\tau=0} \\
&= -\int_\Theta p(\delta(y;\theta)|\theta)\,\frac{\partial\delta^c(y;\theta,\tau)}{\partial\tau}\Big|_{\tau=0} dF(\theta) \qquad (30)
\end{aligned}$$

where $\delta(y;\theta)$ is the inverse of $y(\delta;\theta)$ with respect to its first argument and $p(\delta|\theta)$ is the density corresponding to $P(\delta|theta)$.

Individuals solve

$$\max_{c(\theta),y(\delta;\theta)} U\left(c(\theta), \int_0^\infty v(y(\delta;\theta),\delta)dP(\delta|\theta);\theta\right)$$

subject to

$$c(\theta) = \int_0^\infty q(y(\delta;\theta))dP(\delta|\theta)$$

with first-order conditions

$$\begin{aligned}
U_c &= \tilde{\lambda} \\
U_Y v_y(y(\delta;\theta),\delta) &= -\tilde{\lambda}q'(y(\delta;\theta)).
\end{aligned}$$

The Frisch labor supply as defined in the main text is thus $y^F(\delta;\lambda,\tau)$ such that

$$v_y(y^F,\delta) = \lambda q'(y^F;\tau)$$

where $\lambda \equiv -\tilde{\lambda}/U_Y$. Note that λ will in general depend on θ.

We can write the compensated labor supply as $y^c(\delta;\theta,\bar{U},\tau)$ such that

$$v_y(y^c,\delta) = \lambda^c(\theta,\bar{U},\tau)q'(y^c;\tau).$$

Dropping the argument $\bar{U}$, this equivalently determines $\delta^c(y;\theta,\tau)$ such that

$$v_y(y,\delta^c) = \lambda^c(\theta,\tau)q'(y;\tau).$$

We are now able to compute

$$\left.\frac{\partial \delta^c(y;\theta,\tau)}{\partial \tau}\right|_{\tau=0} = \frac{\frac{\partial \lambda^c}{\partial \tau} q' + \lambda^c t'}{v_{y\delta}} = \frac{\frac{1}{\lambda^c}\frac{\partial \lambda^c}{\partial \tau} + t'/q'}{v_{y\delta}/v_y}.$$

At $\tau = 0$, we can also compute (for the change of variables from δ to y)

$$\frac{\partial \delta(y;\theta)}{\partial y} \equiv \delta'(y;\theta) = -\frac{v_{yy} - \lambda^c q''}{v_{y\delta}} = -\frac{v_{yy}/v_y - q''/q'}{v_{y\delta}/v_y}.$$

Finally, note that the Frisch elasticity is based on

$$v_y(y^F, \delta) = \lambda \left(q'(y^F) - \xi \right),$$

so

$$\varepsilon^F(y) = -\left.\frac{\partial y^F}{\partial \xi}\right|_{\xi=0} \frac{q'}{y} = \frac{\lambda}{v_{yy} - \lambda q''}\frac{q'}{y} = \frac{1/y}{v_{yy}/v_y - q''/q'}.$$

(Observe that this does not depend on θ.) Using all this, we can write (30) as

$$\begin{aligned}
& \int_\Theta p(\delta(y;\theta)|\theta)\delta'(y;\theta)\frac{\frac{1}{\lambda^c}\frac{\partial \lambda^c}{\partial \tau} + t'/q'}{v_{yy}/v_y - q''/q'} dF(\theta) \\
= & \int_\Theta p(\delta(y;\theta)|\theta)\delta'(y;\theta)\left(\frac{1}{\lambda^c}\frac{\partial \lambda^c}{\partial \tau} + t'/q'\right) y\varepsilon^F(y) dF(\theta) \\
= & \; y\varepsilon^F(y)h(y)\left(\frac{t'(y)}{q'(y)} + \int_\Theta \frac{1}{\lambda^c}\frac{\partial \lambda^c}{\partial \tau} dF(\theta|y)\right).
\end{aligned}$$

For the last step, we noted that $\lambda^c(\theta,\tau)$ depends on θ and we used the fact that

$$\frac{p(\delta(y;\theta)|\theta)\delta'(y;\theta)f(\theta)}{h(y)} = f(\theta|y). \tag{31}$$

To see this, note that, given θ, by monotonicity of δ in y, we have $H(y|\theta) = P(\delta(y;\theta)|\theta)$. Differentiating this, we obtain the density of y conditional on θ: $h(y|\theta) = p(\delta(y;\theta)|\theta)\delta'(y;\theta)$. Multiplying this by the marginal density $f(\theta)$ for θ, we obtain the joint density $h(y,\theta) = p(\delta(y;\theta)|\theta)\delta'(y;\theta)f(\theta)$. By Bayes' Rule, this implies the conditional density $f(\theta|y)$ of θ conditional on y given by (31).

As for the right-hand side, we have

$$\begin{aligned}
& \int_{\Theta}\int_{0}^{\infty}(1-H_{\delta,\theta}(y))\left(\beta_\theta - 1 - \frac{\partial}{\partial I}\int_0^\infty t'(z)\,(1-H_{\delta,\theta}(z;I))\,dz\right)dP(\delta|\theta)dF(\theta) \\
= & \int_{\Theta}\int_{\delta(y;\theta)}^{\infty}\left(\beta_\theta - 1 - \frac{\partial}{\partial I}\int_0^{y(\delta;\theta,I)} t'(z)dz\right)dP(\delta|\theta)dF(\theta) \\
= & \int_{\Theta}\int_{\delta(y;\theta)}^{\infty}\left(\beta_\theta - 1 - \frac{\partial y(\delta;\theta,I)}{\partial I}t'(y(\delta;\theta))\right)dP(\delta|\theta)dF(\theta).
\end{aligned}$$

Using $\partial y(\delta(y;\theta);\theta,I)/\partial I = -\eta(y,\theta)/q'(y)$, this becomes after changing variables in the inner integral

$$\begin{aligned}
& \int_{\Theta}\int_{y}^{\infty}\left(\beta_\theta - 1 + \eta(z,\theta)\frac{t'(z)}{q'(z)}\right)p(\delta(z;\theta)|\theta)\delta'(z;\theta)dzdF(\theta) \\
= & -\int_y^\infty\int_\Theta(1-\beta_\theta)\,dF(\theta|z)dH(z) + \int_y^\infty\int_\Theta \eta(z,\theta)dF(\theta|z)\frac{t'(z)}{q'(z)}dH(z) \\
= & -\int_y^\infty(1-\bar{\beta}_z)\,dH(z) + \int_y^\infty \bar{\eta}(z)\frac{t'(z)}{q'(z)}dH(z), \qquad (32)
\end{aligned}$$

where $\bar{\eta}(y)$ is the average income effect and $\bar{\beta}_y$ the average social welfare weight at y.

B.2 Human Capital

Individuals now solve

$$\max_{c(\theta),y(\delta;\theta),e(\theta)} U\left(c(\theta), \int_0^\infty v(y(\delta;\theta),\delta)dP(\delta|\theta,e(\theta));\theta,e(\theta)\right)$$

subject to

$$c(\theta) = \int_0^\infty q(y(\delta;\theta))dP(\delta|\theta,e(\theta))$$

with first-order conditions

$$\begin{aligned}
U_c &= \tilde{\lambda} \\
U_Y v_y(y(\delta;\theta),\delta) &= -\tilde{\lambda}q'(y(\delta;\theta)) \\
U_Y\int_0^\infty v(y(\delta;\theta),\delta)dP_e(\delta|\theta,e(\theta)) + U_e &= -\tilde{\lambda}\int_0^\infty q(y(\delta;\theta))dP_e(\delta|\theta,e(\theta)).
\end{aligned}$$

Defining the Frisch labor supply as above, holding fixed $\lambda = -\tilde{\lambda}/U_Y = -U_c/U_Y$, all the analysis in Appendix B.1 goes through.

B.3 General Heterogeneity

As for the left-hand side of the Diamond-Mirrlees formula (5), consider the same variation of the price schedule $q'(y)$ as in (10) and let $H_{\theta,\phi}(y;\tau) = \mathbb{I}(y \geq y(\theta;\phi,\tau))$, where $y(\theta;\phi,\tau)$ is the income chosen by θ, ϕ when faced with $q'(y;\tau)$, given by the first-order condition

$$MRS(c, y; \theta, \phi) = q'(y;\tau) \tag{33}$$

Then we can write the left-hand side of (5) as

$$\begin{aligned} &\frac{\partial}{\partial \tau} \int_\Phi \int_0^\infty (1 - H^c_{\theta,\phi}(y;\tau)) dF(\theta|\phi) dP(\phi) \Big|_{\tau=0} = \frac{\partial}{\partial \tau} \int_\Phi \int_{\theta^c(y;\phi,\tau)}^\infty dF(\theta|\phi) dP(\phi) \Big|_{\tau=0} \\ = \; & - \int_\Phi f(\theta(y;\phi)|\phi) \left. \frac{\partial \theta^c(y;\phi,\tau)}{\partial \tau} \right|_{\tau=0} dP(\phi) \end{aligned} \tag{34}$$

where $F(\theta|\phi)$ is the c.d.f. of θ conditional on ϕ and $f(\theta|\phi)$ is the corresponding conditional density. Differentiating (33) yields

$$\left. \frac{\partial \theta^c(y;\phi,\tau)}{\partial \tau} \right|_{\tau=0} = \frac{t'(y)}{MRS_\theta} = \frac{t'(y)/q'(y)}{MRS_\theta / MRS}.$$

Observe that the income distribution is $H(y) = \int_\Phi F(\theta(y;\phi)|\phi) dP(\phi)$ and hence $h(y) = \int_\Phi f(\theta(y;\phi)|\phi)\theta'(y;\phi) dP(\phi)$. We compute

$$\theta'(y;\phi) = -\frac{MRS_c + \frac{MRS_y}{MRS} - \frac{q''(y)}{q'(y)}}{\frac{MRS_\theta}{MRS}} = -\frac{1}{y\varepsilon^c(y,\phi)} \frac{MRS}{MRS_\theta},$$

where we used the expression, analogous to (26), for the compensated elasticity

$$\varepsilon^c(y,\phi) = \frac{1/y}{MRS_c(y, \theta(y;\phi), \phi) + \frac{MRS_y(y,\theta(y;\phi),\phi)}{MRS(y)} - \frac{q''(y)}{q'(y)}}.$$

Using this, (34) becomes

$$\begin{aligned} \int_\Phi f(\theta(y;\phi)|\phi) \frac{t'(y)}{q'(y)} \frac{MRS}{MRS_\theta} dP(\phi) &= yh(y) \frac{t'(y)}{q'(y)} \int_\Phi \frac{f(\theta(y;\phi)|\phi)\theta'(y;\phi)}{h(y)} \varepsilon^c(y,\phi) dP(\phi) \\ &= yh(y) \frac{t'(y)}{q'(y)} \bar{\varepsilon}^c(y), \end{aligned} \tag{35}$$

where $\bar{\varepsilon}^c(y) = \mathbb{E}[\varepsilon^c(y,\phi)|y]$ denotes the average compensated elasticity at y (and we again used the fact that, by the same argument as in Appendix B.1, the weights on the right-hand

side of the first equation in (35) equal the density $p(\phi|y)$ of ϕ conditional on y).

As for the right-hand side of (5), we have

$$\begin{aligned}
&\int_\Phi \int_0^\infty (1 - H_{\theta,\phi}(y)) \left(\beta_{\theta,\phi} - 1 - \frac{\partial}{\partial I} \int_0^\infty t'(z) \left(1 - H_{\theta,\phi}(z; I)\right) dz \right) dF(\theta|\phi) dP(\phi) \\
= &\int_\Phi \int_{\theta(y;\phi)}^\infty \left(\beta_{\theta,\phi} - 1 - \frac{\partial}{\partial I} \int_0^{y(\theta;\phi,I)} t'(z) dz \right) dF(\theta|\phi) dP(\phi) \\
= &\int_\Phi \int_{\theta(y;\phi)}^\infty \left(\beta_{\theta,\phi} - 1 - \frac{\partial y(\theta;\phi,I)}{\partial I} t'(y(\theta;\phi)) \right) dF(\theta|\phi) dP(\phi).
\end{aligned}$$

Again using $\partial y(\theta(y;\phi);\phi,I)/\partial I = -\eta(y,\phi)/q'(y)$, this becomes after changing variables in the inner integral

$$\begin{aligned}
&\int_\Phi \int_y^\infty \left(\beta_{\theta(z;\phi),\phi} - 1 + \eta(z,\phi) \frac{t'(z)}{q'(z)} \right) f(\theta(z;\phi)|\phi) \theta'(z;\phi) dz dP(\phi) \\
= &-\int_y^\infty \int_\Phi \left(1 - \beta_{\theta(z;\phi),\phi}\right) dP(\phi|z) dH(z) + \int_y^\infty \int_\Phi \eta(z,\phi) dP(\phi|z) \frac{t'(z)}{q'(z)} dH(z) \\
= &-\int_y^\infty \left(1 - \bar{\beta}_z\right) dH(z) + \int_y^\infty \bar{\eta}(z) \frac{t'(z)}{q'(z)} dH(z) \qquad (36)
\end{aligned}$$

where $\bar{\eta}(y)$ is the average income effect and $\bar{\beta}_y$ the average social welfare weight at y. Equating (35) and (36) and following the same steps as in Appendix A delivers the results.

B.4 Extensive-Margin Choices

Denoting by $y(\theta)$ the preferred labor supply of an individual of type θ among all $y > 0$, this individual will choose $y(\theta)$ instead of $y = 0$ if and only if

$$\varphi \leq \varphi_\theta(q)$$

where $\varphi_\theta(q)$ is such that

$$U(q(y(\theta)), y(\theta); \theta) = u(q(0); \theta, \varphi_\theta(q)).$$

Let the distribution of φ conditional on θ be given by $\Gamma(\varphi|\theta)$ and denote the corresponding conditional density by $\gamma(\varphi|\theta)$. Then a share $\Gamma(\varphi_\theta(q)|\theta)$ of all θ-types will supply $y(\theta)$ and

the rest $y = 0$. Hence, we can write the left-hand side of (5) as

$$\begin{aligned} \frac{\partial}{\partial \tau}(1 - H^c(y;\tau))\Big|_{\tau=0} &= \frac{\partial}{\partial \tau}\int_{\theta^c(y;\tau)}^{\infty} \Gamma(\varphi_\theta^c(q;\tau)|\theta)dF(\theta)\Big|_{\tau=0} \\ = \quad & -\Gamma(\varphi_{\theta(y)}(q)|\theta(y))f(\theta(y))\,\frac{\partial \theta^c(y;\tau)}{\partial \tau}\Big|_{\tau=0} + \int_{\theta(y)}^{\infty} \gamma(\varphi_\theta(q)|\theta)\,\frac{\partial \varphi_\theta^c(q;\tau)}{\partial \tau}\Big|_{\tau=0} dF(\theta) \end{aligned} \tag{37}$$

Note that the density of y is now $h(y) = \Gamma(\varphi_{\theta(y)}(q)|\theta(y))f(\theta(y))\theta'(y)$, so the first term is the standard one derived in Appendix A. Integrating the variation defined in (10), we have

$$q(y;\tau) = q(y) + \tau(t(y) - t(0)),$$

so we can write

$$\frac{\partial \varphi_\theta^c(q;\tau)}{\partial \tau}\Big|_{\tau=0} = (t(y(\theta)) - t(0))\frac{\partial \varphi_\theta^c(q)}{\partial q(y)}.$$

Using this, the second term in (37) becomes

$$\int_{\theta(y)}^{\infty} (t(y(\theta)) - t(0))\gamma(\varphi_\theta(q)|\theta)\frac{\partial \varphi_\theta^c(q)}{\partial q(y)}dF(\theta) = \int_y^{\infty} \frac{t(z) - t(0)}{q(z)}\rho^c(z)dH(z) \tag{38}$$

where $\rho^c(y)$ is the compensated participation elasticity at y.

The right-hand side of (5) becomes

$$\int_{\theta(y)}^{\infty}\int_{-\infty}^{\varphi_\theta(q)} (\beta_{\theta,\varphi} - 1)d\Gamma(\varphi|\theta)dF(\theta) - \frac{\partial}{\partial I}\int_{\theta(y)}^{\infty}\int_0^{y(\theta;I)} t'(z)dz\Gamma(\varphi_\theta(q;I)|\theta)dF(\theta). \tag{39}$$

The first term can be rewritten as $\int_{\theta(y)}^{\infty}(\bar{\beta}_\theta - 1)\Gamma(\varphi_\theta(q)|\theta)dF(\theta)$ with

$$\bar{\beta}_\theta = \int_{-\infty}^{\varphi_\theta(q)} (\beta_{\theta,\varphi} - 1)\frac{d\Gamma(\varphi|\theta)}{\Gamma(\varphi_\theta(q)|\theta)} = \mathbb{E}\left[\beta_{\theta,\varphi}|\theta, \varphi \leq \varphi_\theta(q)\right]$$

and hence, after changing variables, as

$$\int_y^{\infty} (\bar{\beta}_z - 1)dH(z)$$

(where we slightly abused notation to write $\bar{\beta}_y = \bar{\beta}_{\theta(y)}$).

The second term in (39) equals

$$-\int_{\theta(y)}^{\infty} \frac{\partial y(\theta; I)}{\partial I} t'(y(\theta))\Gamma(\varphi_\theta(q)|\theta)dF(\theta) - \int_{\theta(y)}^{\infty} (t(y(\theta)) - t(0))\gamma(\varphi_\theta(q)|\theta)\frac{\partial \varphi_\theta(q; I)}{\partial I} dF(\theta). \tag{40}$$

The first term here is again standard and the same as in Section 5. The second term in (40) can be combined with (38) to deliver the uncompensated extensive-margin response, i.e.

$$-\int_{y}^{\infty} \frac{t(z) - t(0)}{q(z)} \rho(z) dH(z).$$

Collecting all these results and equating the left- and right-hand side yields

$$-yh(y)\frac{t'(y)}{q'(y)}\varepsilon^c(y) = \int_{y}^{\infty} (1 - \bar{\beta}_z) dH(z) - \int_{y}^{\infty} \frac{t'(z)}{q'(z)} \eta(z) dH(z) + \int_{y}^{\infty} \frac{t(z) - t(0)}{q(z)} \rho(z) dH(z)$$

and hence the condition in the main text.