Mirrlees meets Diamond-Mirrlees: Simplifying Nonlinear Income Taxation*

Florian Scheuer — University of Zurich

Iván Werning — MIT

September 2018

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 provide a simple derivation of the Mirrleesian optimal income tax formula from the Diamond-Mirrlees commodity tax formula. In the Mirrlees model, the relevant compensated cross-price 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 a direct application of the general Diamond-Mirrlees linear tax formula.

1 Introduction

The Mirrlees (1971) model is a milestone in the study of optimal nonlinear taxation of labor earnings. The Diamond and Mirrlees (1971a; 1971b) 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. We also

*For helpful comments and discussions we thank Andreas Haller, Jim Hines, Jason Huang, Louis Kaplow, Narayana Kocherlakota, Nicola Pavoni, Jim Poterba, Terhi Ravaska, Juan Rios, Casey Rothschild, Emmanuel Saez, Julian Teichgräber, Uwe Thümmel, Jean Tirole, Carlo Zanella as well as numerous conference and seminar participants.

establish a direct link between the widely used optimal tax formulas in both models. In particular, we provide a simple derivation of the nonlinear income tax formula from the linear commodity tax formula. We show that this novel approach to nonlinear taxation greatly expands the generality of the Mirrlees formula, and is useful to derive similar formulas in a variety of richer applications.

The connection between the Mirrlees and the Diamond-Mirrlees models 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 replicated 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 complication with this approach is that it requires working with a continuum of goods. In particular, in the Mirrlees model, there is a nonlinear tax on income y. Instead of treating y as the quantity for a single good, we model each y as indexing a separate good. Since any positive income y is allowed, the set of goods allowed is the positive real line.[1] We provide a formal characterization of the demand system in this infinite-dimensional commodity space and show that it satisfies standard properties familiar from the finite-goods case.

In addition to highlighting a deep connection between two canonical models in public finance, we also uncover a novel link between the associated tax formulas. Both Mirrlees (1971) and Diamond and Mirrlees (1971b) derived 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), among others. In the case of Diamond and Mirrlees (1971b), one can point to Mirrlees (1975) and especially Diamond (1975), who provided a many-person Ramsey tax formula, as well as the dynamic Ramsey models of linear labor and capital taxation that have been widely used in the macroeconomics literature (e.g. Chamley, 1986, and Judd, 1985).

Our paper provides a connection between these two literatures by showing that the Mirrlees formula can, in fact, be derived as an immediate implication of the Diamond-Mirrlees formula. In particular, we start with a version of the general Diamond-Mirrlees linear tax formula, as provided by Diamond (1975), and show that it specializes to the Mirrlees nonlinear tax formula in its integral form, as provided by Diamond (1998), Saez (2001) and others. A connection between the two formulas is natural once we have shown that Diamond-Mirrlees's framework nests Mirrlees's. However, moving from the Diamond-Mirrlees for-

[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.

mula to the Mirrlees formula is not 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]

A major benefit of demonstrating the connection between the two formulas is to offer a common economic intuition. The Diamond-Mirrlees formula (as seen through the lens of Diamond, 1975) equates two sides, each with a simple 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 formula, 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 of compensated demands in the Diamond-Mirrlees formula are 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.

The Diamond-Mirrlees formula also allows for a simple interpretation in terms of the excess burden of the tax system: the marginal deadweight loss from the tax on each good, per dollar of tax revenue, should be balanced with the redistributive benefits of taxing that good. For instance, in the absence of redistributive motives, the formula calls for equalizing the marginal deadweight loss per dollar of revenue across all goods. Our analysis shows that this intuition immediately carries over to the Mirrlees model.

Apart from providing an alternative interpretation of the Mirrlees formula in terms of the Diamond-Mirrlees formula, attacking the nonlinear tax problem this way allows us to show that the Mirrlees formula holds under weaker conditions than commonly imposed.

[2] Diamond and Mirrlees (1971b) also briefly extend their analysis to consider parametric nonlinear tax systems and derive an optimality condition. However, due to its abstract nature, they do not develop it in detail, and their particular 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.

Most importantly, 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. Our connection reveals that, in fact, the Mirrlees formula also holds when there are general equilibrium effects as in Diamond-Mirrlees, and the production efficiency theorem extends to the Mirrlees model. This is a novel substantive insight from our analysis that helps putting into perspective the recent, growing literature on Mirrleesian taxation in settings where different types of labor are imperfect substitutes (see, among others, Rothschild and Scheuer, 2013, Ales et al., 2015, Sachs et al., 2017, and Costinot and Werning, 2018).

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, stochastic productivity profile, but pays taxes based on current income. This is in line with present practice, where taxes are assessed annually, despite individuals' earnings varying significantly over their lifecycle.

In this context, 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.[3] The connection to the Diamond-Mirrlees model, however, allows us to derive a novel 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. We characterize these lifetime effects in general and show that they vanish when preferences are quasilinear. Hence, in this simple case, the formula for the annual tax in the dynamic model coincides in format with that of the static model. In other words, our connection reveals that the standard Mirrlees formula holds even in dynamic settings under natural assumptions. For example, our formulation can easily incorporate overlapping generations or rich stochastic processes driving lifecycle productivity profiles.

Our analysis highlights subtleties in applying this formula, however. It requires taking into account, for instance, that welfare weights are a function of lifetime differences in earnings, rather than current differences in annual earnings. Since inequality in lifetime earnings is smaller than inequality in annual earnings, the benefits for redistribution are smaller for a given welfare function. For this reason, even though the formula is the same, we show that

[3] Farhi and Werning (2013) characterize optimal taxes without such constraints in a life-cycle context. They then compute numerically the optimum without state-contingent or age-dependent taxes. See also Weinzierl (2011) for a quantification of the welfare gains from age-dependent taxes.

the implied optimal tax schedule involves less redistribution.

Second, we incorporate human capital investment into the lifecycle framework. Individuals choose an education level that will affect their distribution of lifetime productivities. They 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 distributions extends to this case, with the only difference that the extra term now also captures the effect of taxes on human capital. More generally, the new term 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. We show in full generality that the Diamond-Mirrlees approach provides a very straightforward way of dealing with such 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, we find that the demand system still retains an elementary structure, which is the underlying reason for why relatively simple and easily interpretable tax formulas obtain.

Of the four extensions we offer, we find the first two 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 quite 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 from the perspective of Diamond-Mirrlees. Indeed, our method is able to handle these extensions with ease while highlighting the economics in each case, summarized by the impact of different assumptions on the resulting demand system.

Related Literature. Our approach allows for a simple derivation and interpretation of the optimal nonlinear income tax formula that circumvents the complexities of the traditional mechanism design approach employed by Mirrlees (1971). An alternative approach, both in linear and nonlinear tax contexts, has been to use tax reform arguments in order to derive optimal tax formulas. For linear tax instruments, this variational approach goes back to Dixit (1975). For nonlinear taxation, Roberts (2000) and Saez (2001) have provided heuristic derivations of the Mirrlees formula[4] based on a perturbation where, starting from the optimal tax schedule, marginal tax rates are increased by a small amount in a small interval around a given income level. To a first order, this variation induces a substitution effect in that interval as well as income and welfare effects for everyone above that income. The fact that such a variation cannot improve welfare at an optimum delivers an optimal tax formula. None of these papers attempt to connect Mirrlees's nonlinear tax model, or the associated variational approaches and formulas, to the linear tax model and results in Diamond and Mirrlees (1971b) or Diamond (1975). By contrast, we show how to obtain the nonlinear income tax formula directly as a special case of the linear commodity tax formulas.

Interestingly, the linear tax formula that is our starting point also implicitly involves a variation of the tax system. As mentioned above, one side of the Diamond-Mirrlees formula features the change in compensated demand for one good when the tax rates on all goods are increased proportionally. Translated to nonlinear income taxation, this corresponds to varying the entire schedule of marginal tax rates proportionally and computing the behavioral response to that variation at a given income level. Instead, the variation in Piketty (1997), Roberts (2000) and Saez (2001) changes the marginal tax rate only locally, and then considers the effect of that local variation throughout the income distribution.

How come these two entirely different variations deliver exactly the same formula? Our analysis reveals that the answer rests in the Slutsky symmetry of the compensated demand system when interpreting each income level as a separate good. The single proportional variation underlying the Diamond-Mirrlees formula is very simple and intuitive—corresponding to a uniform expansion or contraction of the tax system. It also simplifies the computation of the relevant behavioral responses. Instead of computing the effects of a local variation in taxes on the compensated demands across all goods, and repeating this for each

[4] See also Piketty (1997) for the Rawlsian case.

possible local variation (as in Piketty, 1997, Roberts, 2000, and Saez, 2001), Slutsky symmetry allows us to reduce the problem to computing the effect of a single, common variation in taxes on the compensated demand for each given good. This is helpful especially for some of our extensions.

Golosov et al. (2014) have recently formalized the variational approach and generalized it to richer, dynamic settings. For their nonlinear tax applications, they focus on similar perturbations as in Roberts (2000) and Saez (2001) with local changes in marginal tax rates. They also consider variations within some restricted classes, for instance the set of linear taxes, and show how to obtain Ramsey-style linear tax formulas in that case. This provides a connection between linear and nonlinear taxation at a high level, in the sense that they can be based on different variational arguments. The goal of our paper is different. First, instead of showing that both formulas can be separately derived using some common underlying principle, we demonstrate that the nonlinear income tax formula is in fact a direct implication of the linear commodity tax formula.[5] Second, Golosov et al. (2014) do not aim to show that the Mirrlees (1971) model is a special case of the Diamond and Mirrlees (1971b) model.

The case with general equilibrium effects illustrates why this is useful. Under the perturbation approach with a general production function, each local change in marginal tax rates induces an infinite series of higher-order feedback effects at all income levels. To solve for optimal tax formulas, they need to be kept track of using novel techniques (Sachs et al., 2017). By contrast, Diamond and Mirrlees (1971b) showed that linear taxation in general equilibrium is straightforward, and our connection reveals that this simplicity immediately translates to the Mirrlees model. In particular, it allows us to provide conditions under which the Mirrlees formula is virtually unaffected by general equilibrium effects.

In the context of a quasilinear monopoly pricing model, Goldman et al. (1984) have provided an intuition for the optimal nonlinear pricing rule of a monopolist selling a single good by interpreting each quantity level as a separate "market," with independent demand. The standard Ramsey rule calls for a price inversely proportional to the own-price elasticity in each "market," i.e. at any given quantity level.[6] They emphasize that this connection to linear pricing fails whenever there are income effects, because in that case demands in each of these "markets" depend on inframarginal consumption and, thus, are not independent. Our approach goes beyond interpreting the optimality conditions, but actually connects the nonlinear tax model with the linear tax model itself, and does so while allowing for general income and cross-price effects, as in the general Diamond-Mirrlees demand system that we take as a starting point.

[5] The variational approach would suggest the opposite: linear tax formulas emerge from a more restricted variation than nonlinear tax formulas.

[6] See also Brown and Sibley (1986) and Tirole (2002) for textbook treatments of the relationship between second- and third-degree price discrimination.

This paper is organized as follows. Section 2 introduces both the Diamond and Mirrlees (1971b) 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 (1971b) 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 $h \in \mathcal{H}$. 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 $\mathcal{H}$. Agents face a linear budget constraint

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

with consumer prices q. Diamond and Mirrlees (1971b) consider the case where a "poll tax" is ruled out, by imposing $I = 0$. We adopt this assumption as well, but will point out below that allowing for a nonzero lump-sum tax or transfer, $I \neq 0$, would not make a difference for our analysis.

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 $h \in \mathcal{H}$.

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 it 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 (1971b), there is a finite population $\mathcal{H} = \{1, 2, \ldots, M\}$, 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

$$q \cdot x^h = \sum_{i=1}^{N} q_i x_i^h = 0, \tag{2}$$

where $q = (q_1, q_2, \ldots, q_N)$. As usual, 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, such as labor), while positive entries represent deficits or demand (i.e. buying in the market, such as consumption goods).

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 $\mathcal{H}$. The set of goods X is allowed to be infinite. Each agent h consumes a signed measure χ^h over X and is subject to a budget constraint $B(\chi^h, q) = 0$ as before, where q are consumer prices and B is a linear functional.[7] 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 (scalar) 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

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

[7] For example, if the set of goods is some interval $Z \subseteq \mathbb{R}$ and consumer prices are given by some continuous function $q : Z \to \mathbb{R}_0^+$, as will be the case below, then by the Riesz-Markov-Kakutani representation theorem, we can write the budget constraint in the familiar integral form $\int_Z q(z)\, d\chi^h(z) = 0$.

is strictly decreasing in θ (so higher θ types find it less costly to provide y). The canonical specification in Mirrlees (1971) is $U(c, y; \theta) = \tilde{U}(c, y/\theta)$ 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, in the standard Mirrlees model, 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

$$\int_\Theta c(\theta)\, dF(\theta) \leq G(H), \tag{3}$$

for some production functional G. Hence, consistent with the general technology in Diamond and Mirrlees (1971b), total output depends on the distribution of effective labor in the economy.[8] The canonical specification mentioned earlier is a special case with

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

where the second expression follows by integration by parts. An example with imperfect substitutability is the constant elasticity of substitution (CES) specification

$$G(H) = \left(\int_0^\infty y^\sigma \left(\frac{dH(y)}{dy} \right)^\sigma dy \right)^{1/\sigma},$$

where σ parametrizes the substitution elasticity across labor supply levels.

The goal is to maximize a social welfare function $W\left(\{U(c(\theta), y(\theta); \theta)\}\right)$. 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 (3). Once again, un-

[8]See Section 5.4 for a further discussion.

der 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 the consumer problem 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, so that $h = \theta$ and μ_h is the measure corresponding to the c.d.f. over skills F. The commodity space is comprised of a single consumption and a continuum of labor varieties indexed by $y \geq 0$.[9] 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

$$c \leq \int_0^\infty q(y) dH_\theta(y), \tag{4}$$

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 (Fréchet) derivatives of the production function G; in the standard Mirrlees model with linear technology $p(y) = y$.

Finally, in the canonical Mirrlees model, agents can only put full mass of unity on a single

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

value of y.[10] This is a restriction on preferences. Specifically, one may assume that agents attain utility $U(c, y; \theta)$ when they consume c and put full mass on y; they would obtain a prohibitively lower utility if they attempted to distribute mass over various points or put less than measure one. Thus, the measure corresponds to 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

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

so that the $q(y)$-schedule is effectively the retention function in the Mirrlees model.[11]

This completes the description of a particular Diamond-Mirrlees economy that nests the Mirrlees model. Under this formulation, the agents choose a measure corresponding to $H_\theta(y)$ 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 $c \leq \int_0^\infty q(y) h_\theta(y) dy$. However, in our Mirrlees formulation, we impose that H_θ has no density representation because it is a step function.

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 has a density, the economy produces a density over y in the aggregate. However, if one raises $q(y_0)$ 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 reduce $q(y_0)$ 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

In the preceding 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 to cast the Mirrlees model into the Diamond-Mirrlees framework in terms of the *marginal* price $q'(y)$ and *marginal* taxes $t'(y) = q'(y) - p'(y)$.

[10] See Section 6 for a dynamic extension where this is no longer the case, and agents can choose a non-degenerate distribution over y, which fits naturally with (4).

[11] This also reveals why allowing for a non-zero lump-sum tax or transfer $I \neq 0$ would not make a difference, since it is equivalent to a parallel shift in the function $q(y)$. In other words, when viewed through the lens of Diamond-Mirrlees, the Mirrlees model is a hedonic model (each consumer *has* to select one labor supply variety among all $y \geq 0$), which allows us to normalize $I = 0$ without loss of generality.

Integrating the budget constraint (4) by parts gives

$$c \leq \int_0^\infty q'(y)(1 - H_\theta(y))dy + \tilde{I} \tag{5}$$

where $\tilde{I} = q(0)(1 - H_\theta(0))$.[12]

Under this formulation, we reinterpret $q'(y)$ and $1 - H_\theta(y)$ as the price and quantity, respectively, for good $y > 0$. 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). Since the budget constraint is linear, standard consumer theory continues to apply. Indeed, in Appendix A, we characterize this infinite-dimensional demand system formally and show that standard relationships well-known from the case with a finite number of goods (such as Roy's identity, Shephard's lemma or Slutsky symmetry) can be naturally extended to this case.

This reformulation 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 and $\mathbb{I}$ is the indicator function. 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, one no longer expects aggregate demand for good y, defined by

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

to be necessarily discontinuous with respect to 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. We turn to this next.

4 Tax Formulas: Diamond-Mirrlees and Mirrlees

Here we briefly review the optimal tax formulas implied by both models. These formulas crystalize the main results from these theories, offer intuition and provide the starting points

[12] Hence, as is common in derivations of the Mirrlees formula, equation (5) applies to situations where bunching does not occur, by focusing on points where the consumer price schedule is differentiable (so a well-defined marginal tax rate exists). However, as we show in Appendix A, the reformulation is valid more generally whenever $q(y)$ and $H_\theta(y)$ have no common points of discontinuity.

for empirical applications. Readers familiar with these formulas can skip or quickly skim over 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 formulas, 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 various 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

$$\frac{\partial}{\partial \tau}\left(\sum_{h=1}^{M} x_i^{c,h}(q+\tau t)\right)\bigg|_{\tau=0} = \sum_{h=1}^{M} \hat{\beta}^h x_i^h. \tag{6}$$

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.[13] 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

$$\hat{\beta}^h = \beta^h - 1 + \frac{\partial}{\partial I}\left(\sum_{j=1}^{N} t_j x_j^h(q, I)\right), \tag{7}$$

Here, β^h is the marginal social benefit of increasing income for agent h. The next term, -1, captures the resource cost of providing this 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 thus revenue flows back to the government. When this last term is positive, the social cost is less than 1. Overall, the social marginal utility of income may be positive or negative. Indeed, if the poll tax I were available, then the optimality condition for I would imply that the average of the social marginal utilities of income across agents must be zero:

[13] 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.

$\sum_{h=1}^{M} \hat{\beta}^h = 0$. But formulas (6) and (7) hold even when $I = 0$ is imposed, as in Diamond and Mirrlees (1971b).

Thus, this version of the Diamond-Mirrlees optimal linear tax formula states that the discouragement (or encouragement) of the demand for 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 to whom the government wants to redistribute (i.e. 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 and the labor tax is positive.

In the special "diagonal" case where all compensated cross-price effects are zero, formula (6) simplifies to

$$\frac{t_i}{q_i} = \frac{1}{\varepsilon_{ii}^c} \frac{\sum_h \hat{\beta}_h x_i^h}{\sum_h x_i^h}, \qquad \text{where} \qquad \varepsilon_{ii}^c = \sum_h \frac{\partial x_i^{c,h}}{\partial q_i} \frac{q_i}{\sum_h x_i^h}$$

is the aggregate compensated own-price elasticity of the demand for good i. This is the heterogeneous-agent version of the "inverse elasticity rule" introduced by Ramsey (1927).

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 B 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

$$\varepsilon^u(y) = -\left.\frac{\partial y}{\partial \xi}\right|_{\xi=I=0} \frac{q'(y)}{y} \qquad \text{and} \qquad \eta(y) = -\left.\frac{\partial y}{\partial I}\right|_{\xi=I=0} q'(y), \tag{8}$$

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

$$\frac{T'(y)}{1-T'(y)}\varepsilon^c(y)yh(y) = \int_y^\infty (1-\beta_{\tilde{y}})\, dH(\tilde{y}) + \int_y^\infty \frac{T'(\tilde{y})}{1-T'(\tilde{y})}\eta(\tilde{y})dH(\tilde{y}), \tag{9}$$

at all points where no bunching takes place.[14] Here, H denotes the aggregate c.d.f. for labor supply y, h is its associated density, and β_y is the social marginal utility from consumption of individuals with labor supply y. Equation (9) must be supplemented with a boundary condition, stating that the right-hand side of (9) is equal to zero at the lower bound of the support for $H(y)$.

A version of equation (9) 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 $[y, y+dy]$ (see also Roberts, 2000, for a similar argument). Then the left-hand side of condition (9) 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 (9).[15]

One minor difference compared to Saez (2001) 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, $p(y)$,

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

[15]Golosov et al. (2014) formalize this variational approach and generalize it to richer and dynamic settings.

or both. In particular, the compensated elasticity is affected by the local curvature of q. These definitions are natural in a nonlinear taxation context and help streamline optimal tax formulas (see also Jacquet and Lehmann, 2015, and Scheuer and Werning, 2017, for the use of these elasticity concepts). Indeed, our formula (9) 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.[16]

Equation (9) can be interpreted as a first-order differential equation that implicitly characterizes the optimal tax schedule. Solving it yields the well-known ABC-formula

$$\frac{T'(y)}{1-T'(y)} = A(y)B(y)C(y) \qquad \text{with} \tag{10}$$

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

Both formulas, (9) and (10), are identical when there are no income effects, $\eta = 0$, as in the related formulas derived by Diamond (1998).[17]

5 Tax Formulas: From Diamond-Mirrlees to Mirrlees

We now show how to reach the Mirrlees formulas (9)–(10) starting from the Diamond-Mirrlees formulas (6)–(7). As a first step, it is straightforward to extend the Diamond-Mirrlees formula to an economy with infinitely many commodities, as laid out in the Mirrleesian reinterpretation put forth in Section 3.2. In Appendix A, we characterize the demand system in this economy and provide a formal proof of the Diamond-Mirrlees formula for a continuum of labor varieties $y \in \mathbb{R}_0^+$. The proportional change in all taxes underlying the left-hand side of equation (6) corresponds to changing the marginal consumer price schedule such that

$$q'(y;\tau) = q'(y) + \tau t'(y) \tag{11}$$

for all y, where $t'(y) = q'(y) - p'(y)$. Then the infinite-dimensional equivalent of the Diamond-Mirrlees formula (6) is simply

$$\left.\frac{\partial}{\partial \tau}(1 - H^c(y;\tau))\right|_{\tau=0} = \int_0^\infty (1 - H_\theta(y))\, \hat{\beta}_\theta dF(\theta) \;\; \forall y \in \mathbb{R}_0^+ \tag{12}$$

[16] Since our formulation accounts for this curvature in the elasticities, they directly correspond to the behavioral responses one would estimate, for example, based on a reform of the existing nonlinear tax schedule.

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

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

$$\hat{\beta}_\theta = \beta_\theta - 1 - \frac{\partial}{\partial I} \int_0^\infty t'(z)\,(1 - H_\theta(z; I))\,dz \tag{13}$$

is the social marginal utility of income of individual θ.[18] This is the natural continuous-goods analog of the discrete case. We next compute the left-hand side of (12) in the Mirrlees model, followed by the right-hand side.

5.1 Left-Hand Side in Diamond-Mirrlees

The left-hand side of equation (12) requires calculating

$$\frac{\partial}{\partial \tau}(1 - H^c(y;\tau))\Big|_{\tau=0}, \tag{14}$$

i.e., in words: how does the aggregate cumulative distribution of labor supply at a given y change when all marginal tax rates, throughout the income distribution, are increased proportionally? We provide a simple heuristic for how to compute this in the following; the formal derivation is contained in Appendix B. When τ is increased infinitesimally from zero, the compensated response for each agent is, by (11) and the definition of the compensated elasticity, an increase in y equal to

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

Since each agent increases y, this produces a shift in the distribution H to the right. At a particular point $(y, H(y))$, the horizontal shift equals precisely $\frac{t'(y)}{q'(y)}\varepsilon^c(y)y$. The object in (14), 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 (12) equals

$$\frac{\partial}{\partial \tau}\,(1 - H^c(y;\tau))\Big|_{\tau=0} = \frac{t'(y)}{q'(y)}\varepsilon^c(y)yh(y). \tag{15}$$

This is illustrated in Figure 1.

[18] As our proof in Appendix A demonstrates, this formula allows for individuals choosing some general c.d.f. $H_\theta(y)$ over $y \in \mathbb{R}_0^+$, consistent with the Diamond-Mirrlees model. The Mirrlees model where in fact it is a step function at a single y constitutes a special case. Moreover, note that the income effects enter with the opposite sign in (13) compared to (7), which is because good $1 - H_\theta(y)$ is a net supply, entering budget constraint (5) with the opposite sign compared to the net demands x_i^h in (2).

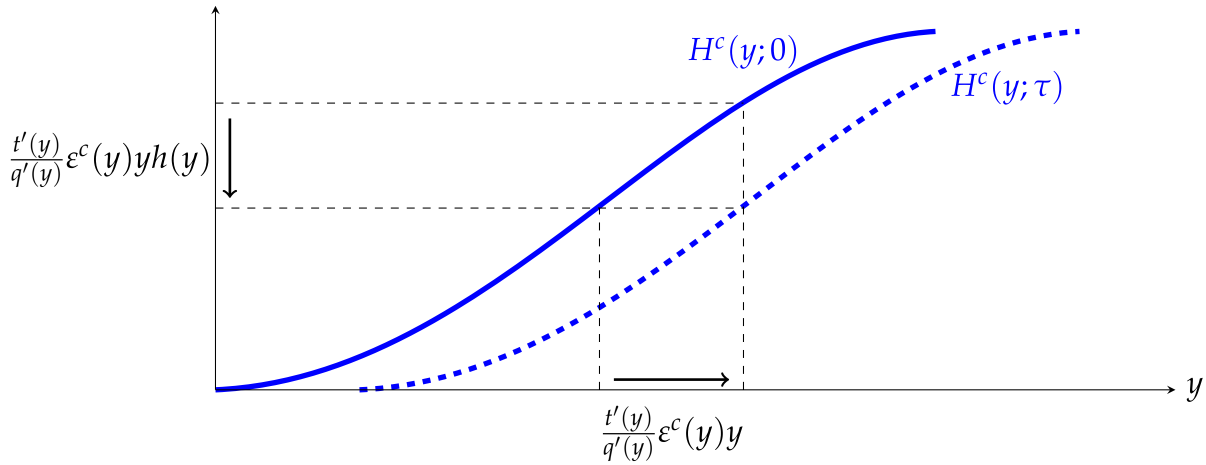

Figure 1: Shift in aggregate compensated demand $1 - H^c(y)$ at y.

Equation (15) 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 (11). 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 Ramsey's "inverse elasticity rule."

5.2 Right-Hand Side in Diamond-Mirrlees

The right-hand side of equation (12) in conjunction with (13) is

$$\begin{aligned}\int_0^\infty (1 - H_\theta(y)) \left(\beta_\theta - 1 - \frac{\partial}{\partial I}\int_0^\infty t'(z)\,(1 - H_\theta(z;I))\,dz\right) dF(\theta)& \\ = \int_{\theta(y)}^\infty \left(\beta_\theta - 1 - \frac{\partial}{\partial I}\int_0^{y(\theta;I)} t'(z)dz\right) dF(\theta)& \\ = \int_{\theta(y)}^\infty \left(\beta_\theta - 1 - t'(y(\theta))\frac{\partial y(\theta;I)}{\partial I}\right) dF(\theta),& \quad (16)\end{aligned}$$

where $\theta(y)$ denotes the inverse of $y(\theta)$.[19] Substituting $\partial y(\theta;I)/\partial I = -\eta(y)/q'(y)$ into (16) and changing variables from θ to $y = y(\theta)$ directly yields

$$-\int_y^\infty (1 - \beta_{\tilde{y}})dH(\tilde{y}) + \int_y^\infty \frac{t'(\tilde{y})}{q'(\tilde{y})}\eta(\tilde{y})dH(\tilde{y}), \quad (17)$$

[19] Recall that $y(\theta)$ is monotone increasing by the single-crossing assumption.

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

5.3 Putting it Together

Equating (15) and (17) yields

$$-\frac{t'(y)}{q'(y)}\varepsilon^c(y)yh(y) = \int_y^\infty (1-\beta_{\tilde{y}})dH(\tilde{y}) - \int_y^\infty \frac{t'(\tilde{y})}{q'(\tilde{y})}\eta(\tilde{y})dH(\tilde{y}). \tag{18}$$

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

$$\frac{t'(y)}{q'(y)} = -\frac{T'(p(y))}{1-T'(p(y))}, \tag{19}$$

which upon substitution gives precisely the Mirrlees formula (9).

It might seem surprising that applying the Diamond-Mirrlees formula (6) to the Mirrlees model immediately delivers exactly the same optimality condition as in Saez (2001) even though the underlying variations are entirely different. Recall that the left-hand side of the Diamond-Mirrlees formula measures the effect on the compensated demand for a *single* good from a proportional change in the tax rates on *all* goods. Translated to nonlinear income taxation, this corresponds to a single variation of the entire schedule of marginal tax rates (a proportional change of all marginal rates) and computing the behavioral response to that variation at a given income level. By contrast, Saez (2001) perturbs the marginal tax rate only locally, and then considers the effect of that local variation throughout the income distribution.

Figure 2 illustrates the difference between these variations. The single proportional variation in the left panel is very simple and perhaps more realistic—corresponding to a uniform expansion or contraction of the tax system—compared to the localized variations in the right panel.

The reason why both approaches lead to the same condition is the Slutsky symmetry of compensated demand, which crucially underlies the left-hand side of the Diamond-Mirrlees formula (6).[20] Instead of computing the effects of a local variation in taxes on the compensated demands across all goods, and repeating this for each possible local variation (as in Saez, 2001), Slutsky symmetry allows us to reduce the problem to computing the effect of a single, common variation in taxes on the compensated demand for each single good. Even though both ultimately coincide in the Mirrlees model because all cross-price effects vanish, the equivalence holds more generally. In particular, the variation underlying Diamond-

[20]See Appendix A for a formal proof.

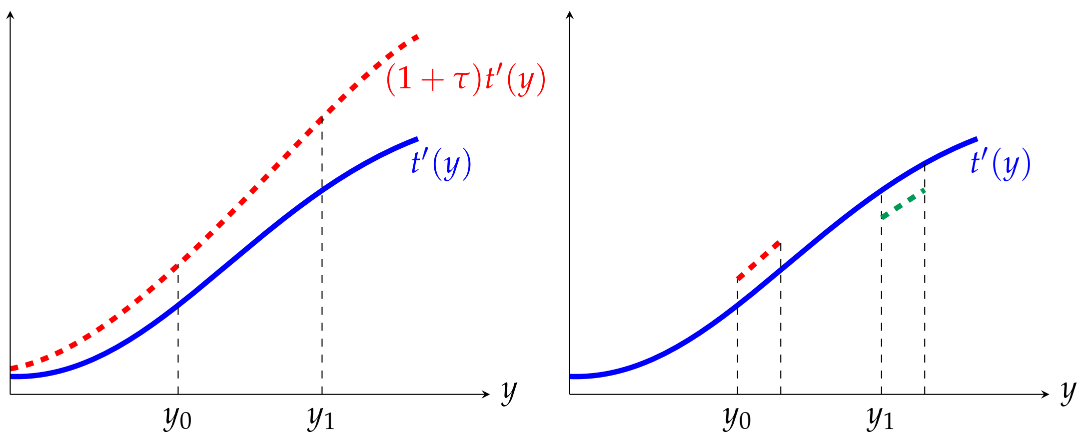

Figure 2: Single variation in Diamond and Mirrlees (1971b) (left) vs. different local variations in Saez (2001) (right).

Mirrlees, which exploits Slutsky symmetry, turns out to be useful in richer settings, such as in the dynamic extensions in Section 6.

In Appendix B, we show how to solve (18) to obtain the ABC-formula

$$-\frac{t'(y)}{q'(y)} = \frac{1}{\varepsilon^c(y)} \frac{1-H(y)}{yh(y)} \int_y^\infty (1-\beta_{\tilde{y}}) \exp\left(\int_y^{\tilde{y}} \frac{\eta(z)}{\varepsilon^c(z)} \frac{dz}{z}\right) \frac{dH(\tilde{y})}{1-H(y)}, \tag{20}$$

which upon the same substitution of the relationship (19) delivers (10). This concludes the derivation of the Mirrlees formulas (9)–(10) from the Diamond-Mirrlees formulas (6)–(7).[21]

5.4 Discussion

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 heuristic derivation (as well as the formal proof in Appendix A) makes clear, the Mirrlees formula for the optimal nonlinear income tax schedule holds for any production function $G(H)$. In other words, based on our connection, the result of Diamond and Mirrlees (1971b) that their tax formula is independent of technology turns out to immediately extend to the Mirrlees model. In particular, the standard Mirrlees formula applies even when different labor supplies are imperfect substitutes, as in the CES

[21] The boundary condition associated with the Mirrlees formula (9) also follows directly from Diamond-Mirrlees. In particular, we show in Appendix A that the first-order condition for $q(0)$ in our reformulation is $\int_0^\infty \hat{\beta}_\theta dF(\theta) = 0$. Together with (16), this implies that the right-hand side of (9) must be zero when evaluated at $y = 0$.

example in Section 2.

While it emerges as a direct side benefit of our approach, we view this neutrality result as a substantive insight in its own right. In particular, one may wonder how it relates to the active literature on optimal income taxation with general equilibrium effects, which has emphasized deviations from the standard Mirrlees formula (see e.g. Rothschild and Scheuer, 2013; 2014; Ales et al., 2015; Sachs et al., 2017). To understand the difference, consider a more general version of the Mirrlees model, with preferences $U(c_\theta, H_\theta, \theta)$, where H_θ denotes the measure that individuals choose over varieties of labor supply from some general set Y and technology is described by $G(H)$ with $H = \int H_\theta dF(\theta)$. The baseline Mirrlees model, as considered here so far, involves $Y = \mathbb{R}_0^+$ and a specification of preferences that lets individuals pick only a single element of Y (i.e. H_θ must be a unit point mass at one $y \in Y$). In other words, this model only allows choice across labor varieties y, but no choice of the intensity of supply at the chosen y. By contrast, the models cited above are in the spirit of Stiglitz (1982), which involve preferences such that individuals choose the intensity with which they supply a given labor variety y (by selecting the measure $H_\theta(y)$ attached to y).[22]

As we show in Appendix A, the Diamond-Mirrlees formula (12) and (13) holds for the general model with preferences $U(c, H_\theta, \theta)$ and technology $G(H)$. Moreover, production efficiency is always optimal by direct application of Diamond and Mirrlees (1971a). In line with Diamond and Mirrlees (1971b), this requires that we can set linear taxes on each variety $y \in Y$. In the absence of Stiglitz effects (when agents choose across different y's rather than intensities), this obviously corresponds to a single nonlinear income tax schedule across y, but otherwise it requires variety-specific (linear) taxes on intensity.[23] Hence, the difference between this clean benchmark and the recent literature is that the latter assumes the absence of sector-specific tax instruments, a crucial deviation from Diamond-Mirrlees (corresponding to a situation where the tax rates on some goods cannot be set separately). With only a single tax schedule $t(y)$ and preferences that imply Stiglitz effects, the Mirrlees formula needs to be modified to reflect general equilibrium effects on redistribution, which do depend on technology. In addition, production efficiency is not necessarily optimal with re-

[22]For example, in Stiglitz (1982), $Y = \{y_L, y_H\}$. There are two types, and low (high) types can only provide low-skill (high-skill) labor, but each type can select the quantity of labor supplied within this variety. Ales et al. (2015) consider $Y = \{y^1, y^2, ..., y^S\}$ and Sachs et al. (2017) $Y = \mathbb{R}_0^+$, each corresponding to one type $\theta \in \mathbb{R}_0^+$. Hence, no choice across elements of Y is allowed in these models. Rothschild and Scheuer (2013; 2014) consider mixed models where individuals choose both across different varieties y and their intensities.

[23]We may also consider nonlinear, variety-specific taxes by defining each quantity supplied of a given labor variety as a separate good and again applying Diamond-Mirrlees to this expanded commodity space. For example, in the Stiglitz-type model with S varieties, this corresponds to variety-specific nonlinear earnings tax schedules $t^1(y), ..., t^S(y)$ being available. Indeed, we show in Appendix B that the infinite-dimensional equivalent of the Diamond-Mirrlees formula (12) and (13) then simply holds for each sector s. This implies that the Mirrlees optimal income tax formula (18) holds for each sector-specific tax schedule $t^s(y)$ (using the sector-specific elasticities and income distributions) independent of the shape of the (multi-sector) production function $G(H)$.

stricted tax instruments (see Guesnerie, 1998, Naito, 1999, Scheuer, 2014, Jacobs, 2015, and, most recently, Costinot and Werning, 2018, Guerreiro et al., 2018, and Thuemmel, 2018).[24]

Within our class of preferences, however, there are no Stiglitz effects, so none of these considerations play a role: A single non-linear income tax schedule is sufficient, the corresponding optimal tax formula is independent of the shape of G, and production should not be distorted.

Deadweight-loss interpretation. The Diamond-Mirrlees formula has an alternative, well-known interpretation in terms of minimizing the total deadweight loss of the tax system. In particular, formula (6) in conjunction with (7) can be rewritten as

$$\frac{\partial D(t)/\partial t_i}{\sum_h x_i^h} = 1 - \frac{\sum_h \beta^h x_i^h}{\sum_h x_i^h}$$

for all goods i, where $D(t) = \sum_h D^h(t)$ and $D^h(t)$ is the (equivalent variation based) deadweight-loss of household h from the tax system t (see Appendix B for a formal derivation). The left-hand side is the marginal deadweight loss from the tax on good i relative to its marginal tax revenue. If there were no redistributive motives (the marginal social benefit of income β^h is the same across agents), the right-hand side would be independent of i. Hence, the optimal tax system would equalize the marginal deadweight loss per dollar of revenue across all goods. More generally, the formula calls for balancing the marginal deadweight loss relative to tax revenue with its redistributive benefits across goods. Goods consumed over-proportionally by those to whom the government wants to redistribute should bear a smaller marginal deadweight loss per dollar of revenue, and vice versa.

Our connection to the Mirrlees model reveals that the same interpretation immediately carries over to the Mirrleesian optimal income tax formula. Using our reinterpretation of the commodity space in the Mirrlees model, it can be written as

$$\frac{\frac{\partial}{\partial \tau} D(T' + \tau \delta_y)\Big|_{\tau=0}}{1 - H(y)} = 1 - \frac{\int_y^\infty \beta_{\tilde{y}} dH(\tilde{y})}{1 - H(y)}.$$

$D(T')$ denotes the total deadweight loss of the tax system, which is a functional that depends on the entire marginal tax schedule T', and δ_y is the Dirac function. Hence, the left-hand side is again the marginal deadweight loss of increasing the marginal tax rate at income level y, relative to its marginal tax revenue (an additional dollar from everyone above y). The right-

[24] 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.

hand side captures the redistributive benefit of this marginal tax change. Our approach makes clear that the optimal nonlinear income tax balances these two considerations across all income levels, with the same intuition as in the linear commodity tax case. In fact, if all β_y were the same, the optimum would equalize the marginal deadweight loss per dollar of revenue across all incomes, just like in the Ramsey case.

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 number 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 (6) and (7) 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.

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 B, we show that there is a direct link between the two, and we demonstrate how to rewrite formulas (9) and (10) 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.[25]

[25] 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 (2017) for the required elasticity adjustments in the context of superstar effects).

6 Four Extensions of the Mirrlees Model

In this section, we briefly consider four extensions that further 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 stochastically 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.[26] While our assumptions and results are slightly different, the main benefit of our treatment of these extensions is to demonstrate how easily these problems can be tackled based on our connection to 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.

Setup. Suppose ex ante heterogeneity is indexed by $\theta \sim F(\theta)$ as before. Each individual faces productivity shocks 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, to focus on the redistributive (rather than insurance) motives for taxation, assume that 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

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

[26] 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. (2017) and Hendren (2014) for the fourth extension.

with

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

Here, $v(y,\delta)$ is a measure of the instantaneous disutility from supplying effective labor y at a moment when productivity is δ; as usual we assume v satisfies 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 (for instance, it could be multidimensional).

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

$$-y\varepsilon^F(y)h(y)\left(\frac{t'(y)}{q'(y)} + \Lambda(y)\right) = \int_y^\infty \left(1-\bar{\beta}_{\tilde{y}}\right) dH(\tilde{y}) - \int_y^\infty \bar{\eta}(\tilde{y})\frac{t'(\tilde{y})}{q'(\tilde{y})} dH(\tilde{y}). \tag{21}$$

This is very similar to the static formula (18) except for the following differences.

First, 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).

Second, $\varepsilon^F(y) \geq 0$ 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 Frisch 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.

Third, there is an extra term on the left-hand side

$$\Lambda(y) = \int_\Theta \frac{1}{\lambda^c(\tau,\theta)} \frac{\partial \lambda^c(\tau,\theta)}{\partial \tau}\bigg|_{\tau=0} dF(\theta|y), \tag{22}$$

which captures precisely the lifetime effects on the compensated labor supply. In particular (and as explained in detail in Appendix C), $\lambda^c(\tau,\theta,;U)$ 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 (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.

Lifetime effects. To illustrate the mechanics underlying the lifetime effects Λ, consider lifetime preferences of the additively separable form

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

We show in Appendix C that, in this case,

$$\frac{1}{\lambda^c(\tau,\theta)}\frac{\partial\lambda^c(\tau,\theta)}{\partial\tau}\bigg|_{\tau=0} = \frac{\int_0^\infty t'(y)y\varepsilon^F(y)dH(y|\theta)}{\frac{u'(c(\theta))}{u''(c(\theta))} - \int_0^\infty q'(y)y\varepsilon^F(y)dH(y|\theta)}.$$

Hence, the lifetime effects depend on the entire tax schedule, Frisch elasticities throughout the income distribution, and risk aversion. Notably, under the standard conditions that $u(c)$ is concave and marginal income tax rates $T'(p(y))$ are positive (so $t'(y) \leq 0$ by (19)), we have $\Lambda(y) > 0$.

Intuitively, the marginal rate of substitution between lifetime consumption and lifetime labor supply is simply $\lambda = u'(c)$, and a proportional increase in all marginal tax rates reduces lifetime consumption and therefore increases marginal utility of consumption. Similar to income effects, this provides a force for higher marginal tax rates. On the other hand, it is straightforward to show (see Appendix C) that the Frisch elasticity, as usual in lifecycle settings, exceeds the compensated labor supply elasticity: $\varepsilon^F(y) \geq \varepsilon^c(y,\theta)$ for all θ, y. This provides a force in the opposite direction.

In the case of the quasilinear lifetime preferences with $u(c) = c$, the lifetime effects Λ vanish and the elasticities coincide. Hence, in this case, the standard formula from the static setting fully extends to the annual tax in this much richer lifecycle framework.[27]

Welfare weights. Even though the formula for the optimal annual tax in our dynamic setting coincides in structure with the formula for the static case, the lifecycle framework has important implications for the average welfare weights $\bar{\beta}_y$ at a given income y on the right-hand side of (21). 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

[27] Farhi and Werning (2013) compute these restricted taxes numerically. Assuming quasilinear and iso-elastic preferences, Golosov et al. (2014) use their general variational approach to provide a formula for the welfare effects of an age- and history-independent reform of the nonlinear labor tax schedule. Their formula features a weighted average of parameters of the age-specific labor income distributions, age-specific labor elasticities, and age-specific cross-effects on capital tax revenue (which we abstract from). Their focus is on comparing this to an age-dependent reform. Chang and Park (2017) study the interaction between nonlinear income taxes and (incomplete) private insurance, making the same assumptions on preferences. They emphasize that their optimal tax formula crucially depends on the crowing-out effects of taxes on private savings. Our formula based on Diamond and Mirrlees (1971b) holds for general preferences and only relies on the cross-sectional income distribution, highlighting the similarity to the static case.

inequality at any given point in time (which is driven by δ in addition to θ), the average welfare weights at a given income $\bar{\beta}_y$ 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. 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. In this case, $\bar{\beta}_y$ is independent of y and optimal annual taxes are zero.

Overlapping generations. It is straightforward to see that tax formula (21) not only applies to the lifecycle problem of a single cohort, but extends to a setting with multiple, overlapping generations. This is because an OLG structure can be conveniently captured by the arbitrary ex ante heterogeneity encoded in θ. In other words, different θ types could capture individuals of different generations, and the government could put different welfare weights on them. Suppose we are in a steady state where the distribution of types (including cohorts) $F(\theta)$ at any given point in time is constant. Then formula (21) for the optimal annual income tax goes through when interpreting $H(y)$ as the stationary income distribution and $\bar{\beta}_y$ as the average welfare weight at y across all generations.

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 $q(y)$.[28]

As we show in Appendix C, all the results from the basic lifecycle framework go through. In particular, the optimal tax formula (21) still applies. The effect $\Lambda(y)$ takes the same form as before (given by (22)), but now also captures the effect of taxes on individuals' human capital choices. The term Λ again vanishes if lifetime preferences take the quasilinear form $U(c, Y; \theta, e) = \tilde{U}(c - Y; \theta, e)$. More generally, the extra term can be interpreted as a "catch all" for any additional margins that affect individuals' lifetime productivity profiles and budget constraints.

[28]We abstract from exploring the optimal tax treatment of the human capital investment e by assuming that it is not taxed nor subsidized directly (see e.g. Bovenberg and Jacobs, 2005, and Stantcheva, 2017, for recent work on this issue).

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.

Returning to the static and deterministic 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 C, we show how to generalize the analysis from Section 5 to such a framework. The Mirrlees optimal tax formulas (9) and (10) 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 $\bar{\beta}_y$).[29]

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 (6)–(7).

[29] Using his 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 both an extended mechanism design approach with pooling and perturbation arguments.

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 intensive margin decisions conditional on θ.[30] 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 C that an application of the Diamond-Mirrlees formulas in this case leads to the following simple modification of formula (9):

$$\frac{T'(y)}{1 - T'(y)} \varepsilon^c(y) y h(y) = \int_y^\infty \left(1 - \bar{\beta}_{\tilde{y}} + \frac{T'(\tilde{y})}{1 - T'(\tilde{y})} \eta(\tilde{y}) - \frac{T(\tilde{y}) - T(0)}{q(\tilde{y}) - q(0)} \rho(\tilde{y}) \right) dH(\tilde{y}), \tag{23}$$

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

$$\rho(y) = \left. \frac{\partial h(y)}{\partial (q(y) - q(0))} \frac{q(y) - q(0)}{h(y)} \right|_{\{y(\theta)\}},$$

which is the percentage change in the density at y when the participation incentives measured by $q(y) - q(0)$ are 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 (6) 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 (7) leads to the additional term on the right-hand side of the optimal tax formula.[31]

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

[31] 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

A special case arises when only the extensive margin is active (see e.g. Diamond, 1980, and Choné and Laroque, 2011), in which case (23) reduces to

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

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 novel connection between the Ramsey and Mirrleesian literatures insightful and, thus, worthwhile in its own right. In addition, this line of attack on the nonlinear tax problem can easily allow for extensions and weaker conditions. We have provided four such extensions to illustrate the appeal of the Diamond-Mirrlees approach. Substantively, our connection reveals that the standard Mirrlees formula applies under natural conditions even in dynamic settings and in environments where general equilibrium forces are at play. In so doing, it allows for considerable simplicity of the derivations in these cases, relative to complementary mechanism design or perturbation methods. We conjecture that this approach could be usefully applied in other settings as well.

References

Ales, L., M. Kurnaz, and C. Sleet, "Tasks, Talents, and Taxes," *American Economic Review*, 2015, *105*, 3061–3101. 1, 5.4, 22

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

Billingsley, P., *Probability and Measure*, John Wiley Sons, 1995. A.1

Bovenberg, L. and B. Jacobs, "Redistribution and Education Subsidies are Siamese Twins," *Journal of Public Economics*, 2005, *89*, 2005–2035. 28

Brown, S. and D. Sibley, *The Theory of Public Utility Pricing*, Cambridge University Press, 1986. 6

the same formula as ours using perturbation and mechanism design approaches. Zoutman et al. (2017) and Hendren (2014) provide related formulas for the fiscal externality in the inverse optimum problem with both intensive and extensive margins.

Cartan, Henri, *Differential Calculus*, Kershaw Publishing, London, 1971. A.1

Chamley, C., "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 1986, *54*, 607–622. 1

Chang, Y. and Y. Park, "Optimal Taxation with Private Insurance," *Mimeo, University of Rochester*, 2017. 27

Choné, P. and G. Laroque, "Optimal Taxation in the Extensive Model," *Journal of Economic Theory*, 2011, *146*, 425–453. 26, 6.4, 6.4

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

Costinot, A. and I. Werning, "Robots, Trade, and Luddism," *Mimeo, MIT*, 2018. 1, 5.4

Diamond, P., "A Many-Person Ramsey Tax Rule," *Journal of Public Economics*, 1975, *4*, 335–342. 1, 1, 4.1, A.4.2

— , "Income Taxation With Fixed Hours of Work," *Journal of Public Economics*, 1980, *13*, 101–110. 1, 26, 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, 17

— **and J. Mirrlees**, "Optimal Taxation and Public Production: I–Production Efficiency," *American Economic Review*, March 1971, *61* (1), 8–27. 1, 5.4

— **and** — , "Optimal Taxation and Public Production II: Tax Rules," *American Economic Review*, June 1971, *61* (3), 261–78. 1, 2, 1, 2, 2.1, 2.1, 2.2, 4.1, 2, 5.4, 27

Dixit, A., "The Welfare Effects of Tax and Price Changes," *Journal of Public Economics*, 1975, *4*, 103–123. 1

Farhi, E. and I. Werning, "Insurance and Taxation over the Life Cycle," *Review of Economic Studies*, 2013, *810*, 596–635. 3, 27

Goldman, B., H. Leland, and D. Sibley, "Optimal Nonuniform Prices," *Review of Economic Studies*, 1984, *51*, 305–319. 1

Golosov, M., A. Tsyvinski, and N. Werquin, "A Variational Approach to the Analysis of Tax Systems," *NBER Working Paper 20780*, 2014. 1, 15, 27

Guerreiro, J., S. Rebelo, and P. Teles, "Should Robots Be Taxed?," *Mimeo, Northwestern University*, 2018. 5.4

Guesnerie, R., "Peut-on Toujours Redistribuer les Gains a la Spécialization et à l'Echange? Un Retour Pointillé 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. 26, 29, 31

Jacobs, B., "Optimal Inefficient Production," *Mimeo, Erasmus University*, 2015. 5.4

— **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, 4.2, 26, 29

— , — , 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, 26, 6.4, 31

Judd, K., "Redistributive taxation in a simple perfect foresight model," *Journal of Public Economics*, 1985, *28*, 59–83. 1

Luenberger, D., *Optimization by Vector Space Methods*, John Wiley & Sons, Inc, 1969. A.2, A.4, A.4.1, A.4.1

Mirrlees, J., "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 1971, *38*, 175–208. 1, 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, 24

Piketty, T., "La Redistribution Fiscale Face au Chômage," *Revue française d'économie*, 1997, *12*, 157–201. 1, 1, 4

Ramsey, F., "A Contribution to the Theory of Taxation," *Economic Journal*, 1927, *37*, 47–61. 4.1

Roberts, K., "A Reconsideration of the Optimal Income Tax," in P. Hammond and G. Myles, eds., *Incentives and Organization: Papers in Honour of Sir James Mirrlees*, Oxford University Press, 2000. 1, 4.2

Rothschild, C. and F. Scheuer, "Redistributive Taxation in the Roy Model," *Quarterly Journal of Economics*, 2013, *128*, 623–668. 1, 5.4, 22

_ **and** _ , "A Theory of Income Taxation under Multidimensional Skill Heterogeneity," *Working Paper*, 2014. 5.4, 22

Sachs, D., A. Tsyvinski, and N. Werquin, "Non-Linear Tax Incidence and Optimal Taxation in General Equilibrium," *Mimeo*, 2017. 1, 1, 5.4, 22

Saez, E., "Optimal Income Transfer Programs: Intensive versus Extensive Labor Supply Responses," *NBER Working Paper 7708*, 2000. 31

_ , "Using Elasticities to Derive Optimal Income Tax Rates," *Review of Economic Studies*, 2001, *68* (1), 205–29. 1, 1, 4.2, 4.2, 5.3, 5.3, 2, 26, 29

_ , "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, 26, 6.4, 31

_ , "Direct Or Indirect Tax Instruments For Redistribution: Short-Run Versus Long-Run," *Journal of Public Economics*, 2004, *88*, 503–518. 24

Scheuer, F., "Entrepreneurial Taxation with Endogenous Entry," *American Economic Journal: Economic Policy*, 2014, *6*, 126–163. 5.4

_ **and I. Werning**, "The Taxation of Superstars," *Quarterly Journal of Economics*, 2017, *132*, 211–270. 4.2, 25

Stantcheva, S., "Optimal Taxation and Human Capital Policies over the Life Cycle," *Journal of Political Economy*, 2017, *125*, 1931–1990. 28

Stiglitz, J., "Self-Selection and Pareto Efficient Taxation," *Journal of Public Economics*, 1982, *17*, 213–240. 5.4, 22, 24, B.3

Thuemmel, U., "Optimal Taxation of Robots," *Mimeo, University of Zurich*, 2018. 5.4

Tirole, Jean, *The Theory of Industrial Organization*, MIT Press, 2002. 6

Weinzierl, M., "The Surprising Power of Age-Dependent Taxes," *Review of Economic Studies*, 2011, *78*, 1490–1518. 3

Zoutman, F., B. Jacobs, and E. Jongen, "Revealed Social Preferences of Dutch Political Parties," *Journal of Public Economics*, 2017, *165*, 81–100. 26, 31

A Formal Proof of Formula (12)

We first characterize the consumer and firm problems for an infinite-dimensional goods space and translate the well-known properties from the finite case to these problems. We then use these properties to formally derive the linear tax formula (12), which is the infinite-dimensional equivalent of the Diamond-Mirrlees formula (6).

A.1 Agents

Agents are indexed by their skill type $\theta \sim F(.)$. They choose their consumption $c_\theta \in \mathbb{R}_0^+$ and measures over labor varieties $H_\theta(y)$. Let Ω denote the vector space of all finite signed measures on $\mathbb{R}_0^+$, and let Ω_s denote the convex subset of Ω that consist of those measures H_θ such that $\int_0^\infty dH_\theta(y) = s$. Preferences are represented by $u(c_\theta, H_\theta; \theta) : \mathbb{R}_0^+ \times \Omega_1 \times \mathbb{R}_0^+ \to \mathbb{R}$, which is a functional that maps consumption and measures over labor varieties into a utility level for each type θ. Therefore, each θ-type solves

$$V_\theta(q, I) \equiv \max_{c_\theta \in \mathbb{R}_0^+, H_\theta \in \Omega_1} u(c_\theta, H_\theta; \theta) \text{ s.t.}$$

$$c_\theta \leq \int_0^\infty q(y) dH_\theta(y) + I.$$

If $q(y)$ is continuous and non-decreasing, we can integrate the right-hand side by parts, so

$$\int_0^\infty q(y) dH_\theta(y) = \int_0^\infty (1 - H_\theta(y)) dq(y) + q(0)(1 - H_\theta(0)).$$

In fact, even if $q(y)$ is not continuous, this is valid whenever $q(y)$ and $H_\theta(y)$ have no common points of discontinuity (see e.g. Theorem 18.4 in Billingsley, 1995). When $q(y)$ is differentiable, this allows us to write the budget constraint as

$$c_\theta \leq \int_0^\infty q'(y) \left(1 - H_\theta(y)\right) dy + \tilde{I}$$

where $\tilde{I} = q(0)(1 - H_\theta(0)) + I$. As noted in Section 3, without loss of generality, we can drop I and express the agent's problem in terms of q' and $q(0)$ only, where q' stands short for the function $q'(y)$.

We make the following assumptions on preferences:

i. Utility is strictly increasing in consumption c_θ.

ii. Utility is decreasing in the expected labor supply implied by H_θ: if $\int_0^\infty y dH_\theta^1(y) < \int_0^\infty y dH_\theta^2(y)$ then $u(c_\theta, H_\theta^1; \theta) > u(c_\theta, H_\theta^2; \theta)$.

These assumptions imply that the budget constraint holds with equality at a solution. Moreover, we assume that u is Fréchet differentiable with respect to H_θ. In the baseline Mirrlees model, assuming that utility also depends negatively on the variance of labor supply implied by H_θ, for example, would ensure that individuals always put the full unit measure on a single labor variety y. Our proof allows for more general preferences u (and hence measures H_θ), however.

With Assumptions (i) and (ii), we have the following equivalent optimization problems for the agents:

Utility Maximization Problem (UMP)

$$V_\theta(q', q(0)) \equiv \max_{c_\theta \in \mathbb{R}_0^+, H_\theta \in \Omega_1} u(c_\theta, H_\theta; \theta) \text{ s.t.}$$

$$c_\theta = \int_0^\infty q'(y) \left(1 - H_\theta(y)\right) dy + q(0). \tag{24}$$

Expenditure Minimization Problem (EMP)

$$e_\theta(q', u) \equiv \min_{c_\theta \in \mathbb{R}_0^+, H_\theta \in \Omega_1} c_\theta - \int_0^\infty q'(y) \left(1 - H_\theta(y)\right) dy - q(0) \text{ s.t.} \tag{25}$$

$$u(c_\theta, H_\theta; \theta) = u.$$

$H_\theta(q', q(0))$ denotes the solution to (UMP) and $H_\theta^c(q', u)$ denotes the solution to (EMP). Duality of the two problems implies

$$H_\theta(q', q(0)) = H_\theta^c(q', V_\theta(q', q(0))). \tag{26}$$

For later use, it is useful to collect a number of properties (well known in the case of finitely many goods) implied by this relationship.

Shephard's Lemma. Applying the envelope theorem to (25), the Fréchet differential of e_θ with respect to q' with increment μ is

$$\left.\frac{d}{d\alpha} e_\theta(q' + \alpha\mu, u)\right|_{\alpha=0} = -\int_0^\infty \mu(y) \left(1 - H_\theta^c(y; q', u)\right) dy \tag{27}$$

for all $\mu \in C$ (where C denotes the space of continuous functions).

Roy's Identity. By duality, $V_\theta(q', e_\theta(q', u)) = u$ for all q'. Hence,

$$\left.\frac{d}{d\alpha} V_\theta(q' + \alpha\mu, q(0))\right|_{\alpha=0} + \frac{\partial V_\theta(q', q(0))}{\partial q(0)} \left.\frac{d}{d\alpha} e_\theta(q' + \alpha\mu, u)\right|_{\alpha=0} = 0 \;\; \forall \mu \in C.$$

Using (26) and Shephard's lemma (27), this implies

$$\left.\frac{d}{d\alpha} V_\theta(q' + \alpha\mu, q(0))\right|_{\alpha=0} = \frac{\partial V_\theta(q', q(0))}{\partial q(0)} \int_0^\infty (1 - H_\theta(y; q', q(0)))\mu(y) dy. \tag{28}$$

Slutsky Symmetry. By Shephard's lemma (27), recall that the Fréchet differential of the expenditure function with increment μ is

$$\delta e_\theta(q', u; \mu) \equiv \left.\frac{d}{d\alpha} e_\theta(q' + \alpha\mu, u)\right|_{\alpha=0} = -\int_0^\infty \mu(y) \left(1 - H_\theta^c(y; q', u)\right) dy.$$

Again taking the Fréchet differential of $\delta e_\theta(q', u; \mu)$ but with increment η, we obtain

$$\frac{d}{d\alpha}\delta e_\theta(q'+\alpha\eta, u; \mu)\bigg|_{\alpha=0} = \int_0^\infty \frac{d}{d\alpha} H_\theta^c(y; q'+\alpha\eta, u)\bigg|_{\alpha=0} \mu(y)dy.$$

By the symmetry of the second Fréchet derivative (see e.g. Theorem 5.1.1 in Cartan, 1971) of the expenditure function, we must therefore have

$$\int_0^\infty \frac{d}{d\alpha} H_\theta^c(y; q'+\alpha\eta, u)\bigg|_{\alpha=0} \mu(y)dy = \int_0^\infty \frac{d}{d\alpha} H_\theta^c(y; q'+\alpha\mu, u)\bigg|_{\alpha=0} \eta(y)dy \tag{29}$$

for all $\mu, \eta \in C$.

A.2 Technology and Firms

Firms are price takers and maximize profits

$$\pi = \max_{H\in\Omega_1} G(H) - \int_0^\infty p(y)dH(y).$$

$G(H) : \Omega_1 \to \mathbb{R}_0^+$ is a constant returns to scale, continuously Fréchet differentiable functional that maps aggregate measures of labor supply into total output. Similar to the reformulation of the household budget constraint, assuming $p(y)$ to be differentiable, we can rewrite profits as

$$\pi = \max_{H\in\Omega_1} G(H) - \int_0^\infty (1-H(y))\, p'(y)dy - (1-H(0))\, p(0).$$

By Theorem 1 (p. 178) in Luenberger (1969), the optimality conditions are

$$\frac{d}{d\alpha}G(H+\alpha\mu)\bigg|_{\alpha=0} + \int_0^\infty p'(y)\mu(y)dy + \mu(0)p(0) = 0 \;\; \forall \mu \in \Omega_0. \tag{30}$$

We can write (30) as

$$G'\mu + \int_0^\infty p'(y)\mu(y)dy + \mu(0)p(0) = 0 \;\; \forall \mu \in \Omega_0, \tag{31}$$

where $G'\mu$ is the Fréchet differential of G with increment μ at the optimal H. Note that the Fréchet derivative G' is, by definition, a linear operator.

Fix any $\mu \in \Omega_0$ and let μ_0 be defined such that $\mu_0(0) = 0$ and $\mu_0(y) = \mu(y)\ \forall y \neq 0$. We have

$$G'\mu_0 + \int_0^\infty p'(y)\mu_0(y)dy = 0, \tag{32}$$

which implies $G'\mu_0 = -\int_0^\infty p'(y)\mu_0(y)dy$ for all μ_0. Subtracting (32) from (31) yields

$$G'[\mu-\mu_0] + \int_0^\infty p'(y)\mu(y)dx - \int_0^\infty p'(y)\mu_0(y)dy + \mu(0)p(0) = 0.$$

Clearly, the integrals cancel because μ and μ_0 disagree only on a measure-zero set. Naturally, we assume that the marginal product of labor variety $y = 0$ is zero: $G'\mu = G'[\mu + C\delta_0]$ for any $\mu \in \Omega$ and $C \in \mathbb{R}$, where δ_0 is defined as $\delta_0(y) = 1$ if $y = 0$ and $\delta_0(y) = 0$ otherwise. Hence, $G'[\mu - \mu_0] = G'[\mu_0 + \mu(0)\delta_0] - G'[\mu_0] = 0$. We can always choose μ with $\mu(0) > 0$, so that $p(0) = 0$ and we have

$$G'\mu = -\int_0^\infty p'(y)\mu(y)dy. \tag{33}$$

A.3 Market Clearing

Labor market clearing requires

$$H(y) = \int_0^\infty H_\theta(y) dF(\theta) \;\; \forall y \in \mathbb{R}_0^+. \tag{34}$$

Goods market clearing requires

$$G(H) = \int_0^\infty c_\theta dF(\theta). \tag{35}$$

A balanced government budget requires

$$t(0) + \int_0^\infty \int_0^\infty t'(y) \left(1 - H_\theta(y)\right) dy dF(\theta) = 0. \tag{36}$$

Since G exhibits constant returns to scale, we have $\pi = 0$. As usual, this together with (34), (35), the household budget constraint (24), $q(y) - p(y) = t(y)$ for all y and $p(0) = 0$ implies that the government budget constraint (36) holds (Walras' Law).

A.4 Planning Problem

The government solves the following problem

$$\max_{q' \in C, q(0) \in \mathbb{R}} V(q', q(0)) = \int_0^\infty W(V_\theta(q', q(0))) dF(\theta)$$

subject to the market clearing conditions (34) and (35). Here, $W(.)$ is some concave, differentiable social welfare function.

Using (24), we can re-write the goods market clearing condition (34) as

$$G(H) = \int_0^\infty \int_0^\infty q'(y) \left(1 - H_\theta(y)\right) dy dF(\theta) + q(0).$$

Moreover, we can directly incorporate the labor market clearing condition (35) by writing the left-hand side as $G(H(q', q(0)))$ where $H(q', q(0)) = \int H_\theta(q', q(0)) dF(\theta)$. Therefore, the planner solves the following problem

$$\max_{q' \in C, q(0) \in \mathbb{R}} V(q', q(0)) = \int_0^\infty W(V_\theta(q', q(0))) dF(\theta) \;\; \text{s.t.}$$

$$G(H(q', q(0))) - \int_0^\infty \int_0^\infty q'(y) \left(1 - H_\theta(y; q', q(0))\right) dy dF(\theta) - q(0) = 0.$$

By Theorem 2 (p. 188) in Luenberger (1969), the solution to this problem also solves

$$\max_{q' \in C, q(0) \in \mathbb{R}} \mathcal{L} = V(q', q(0)) + \gamma \left(G(H(q', q(0))) - \int_0^\infty \int_0^\infty q'(y) \left(1 - H_\theta(y; q', q(0))\right) dy dF(\theta) - q(0) \right)$$

for some $\gamma \geq 0$.

A.4.1 Optimal q'

Applying the chain rule for Fréchet differentials (Proposition 1 (p. 176) in Luenberger, 1969), the optimality condition with respect to q' is

$$
\begin{aligned}
\frac{d}{d\alpha}V(q'+\alpha\mu,q(0))\bigg|_{\alpha=0} + \gamma G'\,\frac{d}{d\alpha}H(q'+\alpha\mu,q(0))\bigg|_{\alpha=0}& \\
-\gamma\int_0^\infty\int_0^\infty \mu(y)\left(1-H_\theta(y;q',q(0))\right)dydF(\theta)& \\
+\gamma\int_0^\infty\int_0^\infty q'(y)\,\frac{d}{d\alpha}H_\theta(y;q'+\alpha\mu,q(0))\bigg|_{\alpha=0}dydF(\theta) &= 0 \qquad (37)
\end{aligned}
$$

for all continuous functions $\mu \in C$. By the firm's optimality condition (33),

$$
G'\frac{d}{d\alpha}H(q'+\alpha\mu,q(0))\bigg|_{\alpha=0} = -\int_0^\infty p'(y)\,\frac{d}{d\alpha}H(y;q'+\alpha\mu,q(0))\bigg|_{\alpha=0}dy.
$$

Using $q'(y)-p'(y) = t'(y)$ and labor market clearing (34), optimality condition (37) therefore becomes

$$
\begin{aligned}
\frac{d}{d\alpha}V(q'+\alpha\mu,q(0))\bigg|_{\alpha=0} - \gamma\int_0^\infty\int_0^\infty \mu(y)\left(1-H_\theta(y;q',q(0))\right)dydF(\theta)& \\
+\gamma\int_0^\infty\int_0^\infty t'(y)\,\frac{d}{d\alpha}H_\theta(y;q'+\alpha\mu,q(0))\bigg|_{\alpha=0}dydF(\theta) &= 0. \qquad (38)
\end{aligned}
$$

We now re-write this first-order condition, translating the standards steps used to derive the Ramsey formula with a finite number of goods to the case we confront here, with an infinite-dimensional commodity space. Starting with the first term, we have

$$
\begin{aligned}
\frac{d}{d\alpha}V(q'+\alpha\mu,q(0))\bigg|_{\alpha=0} &= \int_0^\infty W'(V_\theta(q',q(0))\,\frac{d}{d\alpha}V_\theta(q'+\alpha\mu,q(0))\bigg|_{\alpha=0}dF(\theta) \\
&= \int_0^\infty W'(V_\theta)\frac{\partial V_\theta(q',q(0))}{\partial q(0)}\int_0^\infty(1-H_\theta(y))\mu(y)dydF(\theta) \qquad (39)
\end{aligned}
$$

where the second line uses Roy's identity (28) and we dropped the dependence of H_θ on q' and $q(0)$ to simplify notation.

To simplify the second line in (38), we use (26) to obtain

$$
\begin{aligned}
\int_0^\infty t'(y)\,\frac{d}{d\alpha}H_\theta(y;q'+\alpha\mu,q(0))\bigg|_{\alpha=0}dy &= \int_0^\infty t'(y)\,\frac{d}{d\alpha}H^c_\theta(y;q'+\alpha\mu,V_\theta)\bigg|_{\alpha=0}dy \\
&\quad+\int_0^\infty t'(y)\frac{\partial H^c_\theta}{\partial u}\,\frac{d}{d\alpha}V_\theta(q'+\alpha\mu,q(0))\bigg|_{\alpha=0}dy \qquad (40)
\end{aligned}
$$

where we again dropped arguments to simplify notation. Using Roy's identity (28) and $\frac{\partial H^c_\theta}{\partial u}\frac{\partial V_\theta(q',q(0))}{\partial q(0)} = \frac{\partial H_\theta}{\partial q(0)}$, the second line in (40) becomes

$$
\int_0^\infty t'(y)\frac{\partial H_\theta(y;q',q(0))}{\partial q(0)}dy\int_0^\infty(1-H_\theta(y;q',q(0)))\mu(y)dy.
$$

As for the first term on the right-hand side of (40), note that Slutsky symmetry (29) implies

$$\int_0^\infty t'(y)\,\frac{d}{d\alpha}H^c_\theta(y;q'+\alpha\mu,u)\bigg|_{\alpha=0} \quad = \quad \int_0^\infty \mu(y)\,\frac{d}{d\alpha}H^c_\theta(y;q'+\alpha t',u)\bigg|_{\alpha=0} dy.$$

Substituting all this in (40) yields

$$\begin{aligned}
& \int_0^\infty t'(y)\,\frac{d}{d\alpha}H_\theta(y;q'+\alpha\mu,q(0))\bigg|_{\alpha=0} dy \\
= \quad & -\int_0^\infty \mu(y)\,\frac{d}{d\alpha}\left(1-H^c_\theta(y)\right)\bigg|_{\alpha=0} dy + \int_0^\infty (1-H_\theta(y))\mu(y)dy\int_0^\infty t'(y)\frac{\partial H_\theta(y)}{\partial q(0)}dy.
\end{aligned} \tag{41}$$

Hence, using (39) and (41), we can write the optimality condition (38) as

$$\int_0^\infty\int_0^\infty \left(\phi_1(y,\theta)-\phi_2(y,\theta)\right)dF(\theta)\mu(y)dy = 0 \tag{42}$$

where

$$\phi_1(y,\theta) \equiv \left(\frac{W'(V_\theta)}{\gamma}\frac{\partial V_\theta(q',q(0))}{\partial q(0)} - 1 - \frac{\partial}{\partial q(0)}\int_0^\infty t'(z)\left(1-H_\theta(z;q',q(0))\right)dz\right)(1-H_\theta(y))$$

and

$$\phi_2(y,\theta) \equiv \frac{d}{d\alpha}\left(1-H^c_\theta(y;q'+\alpha t',u)\right)\bigg|_{\alpha=0}.$$

Condition (42) has to hold for all continuous functions $\mu \in C$. By Lemma 1 (p. 180) in Luenberger (1969), this implies that $\int_0^\infty[\phi_1(y,\theta)-\phi_2(y,\theta)]dF(\theta) = 0$ has to hold for all $y \in \mathbb{R}_0^+$. Therefore, the optimality condition for q' becomes

$$\begin{aligned}
& \frac{d}{d\alpha}\left(1-H^c(y;q'+\alpha t')\right)\bigg|_{\alpha=0} \\
= \quad & \int_0^\infty (1-H_\theta(y))\left(\frac{W'(V_\theta)}{\gamma}\frac{\partial V_\theta}{\partial q(0)} - 1 - \frac{\partial}{\partial q(0)}\int_0^\infty t'(z)\left(1-H_\theta(z)\right)dz\right)dF(\theta)
\end{aligned} \tag{43}$$

for all $y \in \mathbb{R}_0^+$.

A.4.2 Optimal Tax Formula

Defining, following Diamond (1975), the social marginal utility of income (in units of consumption) of individual θ as

$$\hat{\beta}_\theta = \beta_\theta - 1 - \frac{\partial}{\partial q(0)}\int_0^\infty t'(z)\left(1-H_\theta(z)\right)dz$$

where

$$\beta_\theta = \frac{W'(V_\theta)}{\gamma}\frac{\partial V_\theta}{\partial q(0)},$$

we can re-write (43) as

$$\frac{d}{d\alpha}\left(1 - H^c(y; q' + \alpha t')\right)\bigg|_{\alpha=0} = \int_0^\infty (1 - H_\theta(y))\,\hat{\beta}_\theta dF(\theta) \;\; \forall y \in \mathbb{R}_0^+,$$

which is equation (12).

A.4.3 Optimal $q(0)$ and Boundary Condition

The optimality condition of the Lagrangian corresponding to the planner's problem with respect to $q(0)$ is

$$\frac{\partial V(q', q(0))}{\partial q(0)} + \gamma G' \frac{\partial H(q', q(0))}{\partial q(0)} + \gamma \int_0^\infty \int_0^\infty q'(y) \frac{\partial H_\theta(y; q', q(0))}{\partial q(0)} dy dF(\theta) - \gamma = 0.$$

Using the firm's optimality condition (33), $q'(y) - p'(y) = t'(y)$ and labor market clearing (34) as in the steps leading to equation (38), this becomes

$$\frac{\partial V(q', q(0))}{\partial q(0)} + \gamma \int_0^\infty \int_0^\infty t'(y) \frac{\partial H_\theta(y; q', q(0))}{\partial q(0)} dy dF(\theta) - \gamma = 0.$$

By the definition of V, we can write this as

$$\int_0^\infty \frac{W'(V_\theta)}{\gamma} \frac{\partial V_\theta(q', q(0))}{\partial q(0)} dF(\theta) - 1 - \int_0^\infty \frac{\partial}{\partial q(0)} \int_0^\infty t'(y) H_\theta(y; q', q(0)) dy dF(\theta) = 0.$$

Finally, using the definition of $\hat{\beta}_\theta$, this simplifies to $\int_0^\infty \hat{\beta}_\theta dF(\theta) = 0$. Using the same steps as in Section 5 delivers the boundary condition corresponding to the Mirrlees formula (9).

B Further Derivations

B.1 Formal Derivation of Equation (15)

In the Mirrlees model, $H_\theta(y) = \mathbb{I}(y \geq y(\theta))$, so we can write (14) as

$$\frac{\partial}{\partial \tau} \int_0^\infty (1 - H^c_\theta(y; \tau)) dF(\theta)\bigg|_{\tau=0} = \frac{\partial}{\partial \tau} \int_{\theta^c(y;\tau)}^\infty dF(\theta)\bigg|_{\tau=0} = -\frac{\partial \theta^c(y; \tau)}{\partial \tau}\bigg|_{\tau=0} f(\theta(y)) \tag{44}$$

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

$$M(c, y; \theta) = q'(y; \tau) = q'(y) + \tau t'(y). \tag{45}$$

To compute the compensated demand, we use this equation with $c = C(v, y; \theta)$ where C 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 (45) yields

$$\frac{\partial \theta^c(y; \tau)}{\partial \tau}\bigg|_{\tau=0} = \frac{q'(y) - p'(y)}{M_\theta} = \frac{t'(y)/q'(y)}{M_\theta / M}. \tag{46}$$

Moreover, observe that the density of y is given by $h(y) = f(\theta(y))\theta'(y)$. Again differentiating (45) for $\tau = 0$ yields

$$\theta'(y) = -\frac{M_c + \frac{M_y}{M} - \frac{q''}{q'}}{M_\theta / M}. \tag{47}$$

Finally, the elasticities defined in (8) can be obtained by differentiating

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

Hence,

$$\varepsilon^u(y) = \frac{-M_c + 1/y}{M_c + \frac{M_y}{M} - \frac{q''}{q'}}, \tag{48}$$

$$\eta(y) = \frac{M_c}{M_c + \frac{M_y}{M} - \frac{q''}{q'}}, \tag{49}$$

and

$$\varepsilon^c(y) = \varepsilon^u(y) + \eta(y) = \frac{1/y}{M_c + \frac{M_y}{M} - \frac{q''}{q'}}. \tag{50}$$

Using (50) in (47) yields

$$\theta'(y) = -\frac{1}{y\varepsilon^c(y)} \frac{1}{M_\theta / M}. \tag{51}$$

Substituting all this in (44), we obtain (15).

B.2 Derivation of Equation (20)

Define

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

and write equation (18) 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 yields (20).

B.3 Multi-Sector Economies

As discussed in Section 5.4, consider an economy with S sectors and technology $G(H^1, ..., H^S)$ given by the constant-returns-to-scale functional $G : \Omega_1^S \to \mathbb{R}_0^+$. Suppose the assignment of each individual to a sector is fixed (as in Stiglitz, 1982), with sector-specific type distributions $F^s(\theta)$, and sector-specific tax schedules $t^s(y)$ are available. This implies hence consumer prices $q^s(y)$ and hence the decisions of an individual in sector s only depend on consumer prices in this sector. We can therefore write $V_\theta^s(q^{s\prime}, q^s(0))$ and $H_\theta^s(y; q^{s\prime}, q^s(0))$ for the indirect utility function and demand of type θ in sector s for good y, respectively. Hence, the planning problem is

$$\max_{\{q^{s\prime}\in C, q^s(0)\in\mathbb{R}\}_{s=1}^S} \sum_{s=1}^S \int_0^\infty W(V_\theta^s(q^{s\prime}, q^s(0)))dF^s(\theta) \text{ s.t.}$$

$$G\left(H^1(q^{1\prime}, q^1(0)), ..., H^S(q^{S\prime}, q^S(0))\right) - \sum_{s=1}^S \left[\int_0^\infty \int_0^\infty q^{s\prime}(y)\left(1 - H_\theta^s(y; q^{s\prime}, q^s(0))\right) dy dF^s(\theta) - q^s(0)\right] = 0.$$

It is clear that the problem is additively separable across sectors except through the technology G. Hence, when taking the Fréchet differential of the corresponding Lagrangian with respect to $q^{s\prime}$ with increment μ, analogous to (37), we obtain

$$\begin{aligned}
\int_0^\infty W'(V_\theta^s(q^{s\prime}, q^s(0)) \frac{d}{d\alpha} V_\theta^s(q^{s\prime} + \alpha\mu, q^s(0))\Big|_{\alpha=0} dF(\theta) \\
+\gamma G_s' \frac{d}{d\alpha} H^s(q^{s\prime} + \alpha\mu, q^s(0))\Big|_{\alpha=0} - \gamma \int_0^\infty \int_0^\infty \mu(y)\left(1 - H_\theta^s(y; q^{s\prime}, q^s(0))\right) dy dF^s(\theta) \\
+\gamma \int_0^\infty \int_0^\infty q^{s\prime}(y) \frac{d}{d\alpha} H_\theta^s(y; q^{s\prime} + \alpha\mu, q^s(0))\Big|_{\alpha=0} dy dF^s(\theta) &= 0,
\end{aligned}$$

where G_s' is the Fréchet derivative of G in dimension s. Moreover, the representative firm's profit maximization problem is

$$\pi = \max_{H^1,...,H^S \in \Omega_1} G(H^1, ..., H^S) - \sum_{s=1}^S \int_0^\infty (1 - H^s(y))\, p^{s\prime}(y) dy$$

with optimality condition

$$\frac{d}{d\alpha} G(H^1, ..., H^s + \alpha\mu, ..., H^S)\Big|_{\alpha=0} + \int_0^\infty p^{s\prime}(y)\mu(y) dy = 0 \;\; \forall s, \mu \in \Omega_0.$$

The first term can be written as $G_s'\mu$, and following the same steps as in Appendix A, we obtain

$$G_s'\mu = -\int_0^\infty p^{s\prime}(y)\mu(y) dy.$$

Substituting this in the above optimality condition for the planning problem and again proceeding as in Appendix A delivers the sector-specific optimal tax formula formula

$$\frac{\partial}{\partial\tau}(1 - H^{s,c}(y;\tau))\Big|_{\tau=0} = \int_0^\infty \left(1 - H_\theta^s(y)\right) \hat{\beta}_\theta^s dF^s(\theta) \;\; \forall y \in \mathbb{R}_0^+$$

where $H^s(y;\tau)$ is the distribution of y in sector s under price schedule $q^{s\prime}(y;\tau)$ and

$$\hat{\beta}_\theta^s = \beta_\theta^s - 1 - \frac{\partial}{\partial I} \int_0^\infty t^{s\prime}(z)\left(1 - H_\theta^s(z; I)\right) dz.$$

Hence, by the same arguments as in Appendix B, we obtain the sector-specific Mirrlees formula

$$-\frac{t^{s\prime}(y)}{q^{s\prime}(y)} \varepsilon^{s,c}(y) y h^s(y) = \int_y^\infty (1 - \beta_{\tilde{y}}^s) dH^s(\tilde{y}) - \int_y^\infty \frac{t^{s\prime}(\tilde{y})}{q^{s\prime}(\tilde{y})} \eta^s(\tilde{y}) dH^s(\tilde{y}) \;\; \forall s, y.$$

B.4 Deadweight Loss Interpretation

The deadweight loss measure, based on the equivalent variation, for agent h of the linear tax system t is

$$D^h(t) = e^h(q, V(q)) - e^h(p, V^h(q)) - \sum_i t_i x_i^{c,h}(q, V^h(q)),$$

where e^h and V^h are the expenditure and indirect utility functions of agent h, respectively, and $x_i^{c,h}$ is the Hicksian demand function for good i. Note that $e^h(q, V(q)) = e^h(p, V(p))$. Hence,

$$\frac{\partial D^h(t)}{\partial t_i} = -e_u^h \frac{\partial V^h}{\partial q_i} - x_i^h - \sum_j t_j \frac{\partial x_j^{c,h}}{\partial q_i} - \sum_j t_j \frac{\partial x_j^{c,h}}{\partial u} \frac{\partial V^h}{\partial q_i}.$$

Using Roy's identity $\partial V^h/\partial q_i = -x_i^h V_I^h$ and Slutsky symmetry, this becomes

$$\frac{\partial D^h(t)}{\partial t_i} = e_u^h V_I^h x_i^h - x_i^h - \sum_j t_j \frac{\partial x_i^{c,h}}{\partial q_j} + \sum_j t_j \frac{\partial x_j^{c,h}}{\partial u} V_I^h x_i^h.$$

Note that, by duality, $e_u^h = 1/V_I^h$ and $\partial x_j^{c,h}/\partial u \times V_I^h = \partial x_j^h/\partial I$, so this simplifies to

$$\frac{\partial D^h(t)}{\partial t_i} = -\sum_j t_j \frac{\partial x_i^{c,h}}{\partial q_j} + \sum_j t_j \frac{\partial x_j^{c,h}}{\partial I} x_i^h.$$

Summing over all agents, we obtain

$$\frac{\partial D(t)}{\partial t_i} = \sum_h \frac{\partial D^h(t)}{\partial t_i} = -\frac{\partial}{\partial \tau} \left(\sum_{h=1}^{M} x_i^{c,h}(q + \tau t) \right) \Bigg|_{\tau=0} + \sum_h x_i^h \frac{\partial}{\partial I} \left(\sum_j t_j x_j^{c,h} \right).$$

Using the Diamond-Mirrlees formula (6) and (7), we can write this as

$$\frac{\partial D(t)}{\partial t_i} = \sum_h x_i^h (1 - \beta^h),$$

which is the formula in Section 5.4.

In the Mirrlees model, the deadweight loss of the tax system can be written as

$$D(q', q(0)) = e_\theta(q', V_\theta(q', q(0))) - e_\theta(p', V_\theta(q', q(0))) + q(0) + \int_0^\infty (q'(y) - p'(y))(1 - H_\theta(y; q'))dy$$

Taking the Fréchet differential with respect to q' with increment δ_y and following the same steps as above using the analysis of the Mirrlees demand system in Appendix A delivers the equivalent, marginal deadweight loss representation of the the Mirrlees optimal tax formula.

B.5 Formulas in Terms of the Skill Distribution

Combine (44) and (46) and change variables from θ to $y(\theta)$ to write the left-hand side of (6) as

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

where we defined $\tau(\theta) = T'(p(y(\theta)))$ and $\chi(\theta) = -(M_\theta \theta/M)^{-1}$. Using this together with (17) yields

$$\frac{\tau(\theta)}{1-\tau(\theta)}\theta f(\theta)\chi(\theta) = \int_\theta^\infty (1-\beta_{\tilde{\theta}})dF(\tilde{\theta}) + \int_\theta^\infty \frac{\tau(\tilde{\theta})}{1-\tau(\tilde{\theta})}\eta(\tilde{\theta})dF(\tilde{\theta}), \tag{52}$$

where we slightly abused notation to write $\eta(\theta) = \eta(y(\theta))$. This is the equivalent of (9) written in terms of θ. Defining the left-hand side of equation (52) 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 (49), (50) and (51) (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

$$\frac{\tau(\theta)}{1-\tau(\theta)} = \frac{1}{\chi(\theta)}\frac{1-F(\theta)}{\theta f(\theta)}\int_\theta^\infty (1-\beta_{\tilde{\theta}})\exp\left(\int_\theta^{\tilde{\theta}} \frac{\eta(s)}{\varepsilon^c(s)}\frac{dy(s)}{y(s)}\right)\frac{dF(\tilde{\theta})}{1-F(\theta)}, \tag{53}$$

which is the equivalent of (10) written in terms of θ.

C Extensions

C.1 Lifecycle Framework

Derivation of formula (21). 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

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

Hence, the left-hand side of (6) simply becomes

$$\begin{aligned}
\frac{\partial}{\partial\tau}(1-H^c(y;\tau))\Big|_{\tau=0} &= \frac{\partial}{\partial\tau}\int_\Theta (1-H^c_\theta(y;\tau))dF(\theta)\Big|_{\tau=0} \\
&= \frac{\partial}{\partial\tau}\int_\Theta\int_0^\infty (1-H^c_{\delta,\theta}(y;\tau))dP(\delta|\theta)dF(\theta)\Big|_{\tau=0} \\
&= \frac{\partial}{\partial\tau}\int_\Theta\int_{\delta^c(y;\theta,\tau)}^\infty dP(\delta|\theta)dF(\theta)\Big|_{\tau=0} \\
&= -\int_\Theta p(\delta(y;\theta)|\theta)\left.\frac{\partial\delta^c(y;\theta,\tau)}{\partial\tau}\right|_{\tau=0} dF(\theta)
\end{aligned} \tag{54}$$

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

$$\varepsilon^F(y) = -\left.\frac{\partial y^F}{\partial \xi}\right|_{\xi=0}\frac{q'}{y} = \frac{\lambda}{v_{yy} - \lambda q''}\frac{q'}{y} = \frac{1/y}{v_{yy}/v_y - q''/q'}. \tag{55}$$

(Observe that this does not depend on θ and that the denominator must be non-negative by the second-order condition.) Using all this, we can write (54) 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

$$\frac{p(\delta(y;\theta)|\theta)\delta'(y;\theta)f(\theta)}{h(y)}=f(\theta|y). \tag{56}$$

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 (56).

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

$$
\begin{aligned}
& \int_\Theta\int_y^\infty\left(\beta_\theta-1+\eta(z,\theta)\frac{t'(z)}{q'(z)}\right)p(\delta(z;\theta)|\theta)\delta'(z;\theta)dzdF(\theta) \\
= & -\int_y^\infty\int_\Theta(1-\beta_\theta)\,dF(\theta|z)dH(z)+\int_y^\infty\int_\Theta\eta(z,\theta)dF(\theta|z)\frac{t'(z)}{q'(z)}dH(z) \\
= & -\int_y^\infty(1-\bar\beta_z)\,dH(z)+\int_y^\infty\bar\eta(z)\frac{t'(z)}{q'(z)}dH(z),
\end{aligned}
\tag{57}
$$

where $\bar\eta(y)$ is the average income effect and $\bar\beta_y$ the average social welfare weight at y.

Characterizing the lifecycle effect Λ. With the separable preferences assumed in the text, we have $\lambda=u'(c)$. Our goal is thus to compute the compensated effect of τ on c. To that end, consider the compensated demand system in the dynamic framework, which solves

$$\max_{c(\theta),y(\delta;\theta)}\int_0^\infty q(y(\delta;\theta);\tau)dP(\delta|\theta)-c(\theta)$$

subject to

$$u(c(\theta))-\int_0^\infty v(y(\delta;\theta),\delta)dP(\delta|\theta)=\bar U(\theta),$$

where $q'(y;\tau)=q'(y)+\tau t'(y)$.

The compensated effects of τ on the solutions $c^c(\theta, \tau)$ and $y^c(\delta; \theta, \tau)$ are therefore implicitly determined by the first-order conditions

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

and

$$u(c^c(\theta,\tau)) - \int_0^\infty v(y^c(\delta;\theta,\tau),\delta)dP(\delta|\theta) = \bar{U}(\theta)$$

for all δ, θ, τ. Differentiating these conditions with respect to τ yields (simplifying notation and changing variables from δ to y)

$$\left.\frac{\partial c^c(\theta,\tau)}{\partial \tau}\right|_{\tau=0} = \frac{\int_0^\infty t'(y)\left(\frac{v_{yy}}{v_y} - \frac{q''}{q'}\right)^{-1} dH(y|\theta)}{1 - \frac{u''(c(\theta))}{u'(c(\theta))}\int_0^\infty q'(y)\left(\frac{v_{yy}}{v_y} - \frac{q''}{q'}\right)^{-1} dH(y|\theta)}.$$

Using the definition of the Frisch elasticity in (55), we can compute

$$\begin{aligned}
\frac{1}{\lambda^c(\theta)}\left.\frac{\partial \lambda^c(\theta,\tau)}{\partial \tau}\right|_{\tau=0} &= \frac{u''(c(\theta))}{u'(c(\theta))}\left.\frac{\partial c^c(\theta,\tau)}{\partial \tau}\right|_{\tau=0} \\
&= u''(c(\theta))\frac{\int_0^\infty t'(y)y\varepsilon^F(y)dH(y|\theta)}{u'(c(\theta)) - u''(c(\theta))\int_0^\infty q'(y)y\varepsilon^F(y)dH(y|\theta)}. \qquad (58)
\end{aligned}$$

Frisch versus compensated elasticities. Consider again the compensated demand system solving

$$\max_{c(\theta),y(\delta;\theta)} \int_0^\infty (q(y(\delta;\theta)) - \xi y(\delta;\theta))dP(\delta|\theta) - c(\theta)$$

subject to

$$u(c(\theta)) - \int_0^\infty v(y(\delta;\theta),\delta)dP(\delta|\theta) = \bar{U}(\theta),$$

where ξ is the increase in the marginal tax rate underlying our definition of the elasticities (8). The compensated demands $c^c(\theta, \xi)$ and $y^c(\delta; \theta, \xi)$ solve

$$v_y(y^c(\delta;\theta,\xi),\delta) = u'(c^c(\theta,\tau))\left(q'(y^c(\delta;\theta,\xi)) - \xi\right)$$

and

$$u(c^c(\theta,\xi)) - \int_0^\infty v(y^c(\delta;\theta,\xi),\delta)dP(\delta|\theta) = \bar{U}(\theta)$$

for all δ, θ, ξ. Differentiating and tedious algebra yield

$$\begin{aligned}
\varepsilon^c(y,\theta) &= -\left.\frac{\partial y^c}{\partial \xi}\right|_{\xi=0}\frac{q'(y)}{y} \\
&= \varepsilon^F(y) + \frac{q'(y)}{y}u''(c(\theta))\frac{\int_0^\infty s\varepsilon^F(s)dH(s|\theta)}{u'(c(\theta)) - u''(c(\theta))\int_0^\infty q'(s)s\varepsilon^F(s)dH(s|\theta)}.
\end{aligned}$$

Hence, whenever $u''(c) \le 0$, we have $\varepsilon^F(y) \ge \varepsilon^c(y, \theta)$.

C.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 C.1 goes through.

C.3 General Heterogeneity

As for the left-hand side of the Diamond-Mirrlees formula (6), consider the same variation of the price schedule $q'(y)$ as in (11) 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

$$M(c,y;\theta,\phi) = q'(y;\tau) \tag{59}$$

Then we can write the left-hand side of (6) as

$$\begin{aligned} \frac{\partial}{\partial\tau}\int_\Phi \int_0^\infty (1 - H^c_{\theta,\phi}(y;\tau))dF(\theta|\phi)dP(\phi)\Big|_{\tau=0} &= \frac{\partial}{\partial\tau}\int_\Phi \int_{\theta^c(y;\phi,\tau)}^\infty dF(\theta|\phi)dP(\phi)\Big|_{\tau=0} \\ = \quad -\int_\Phi f(\theta(y;\phi)|\phi)\, \frac{\partial\theta^c(y;\phi,\tau)}{\partial\tau}\Big|_{\tau=0} dP(\phi) & \end{aligned} \tag{60}$$

where $F(\theta|\phi)$ is the c.d.f. of θ conditional on ϕ and $f(\theta|\phi)$ is the corresponding conditional density. Differentiating (59) yields

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

The cdf of income 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

$$\theta'(y;\phi) = -\frac{M_c + \frac{M_y}{M} - \frac{q''(y)}{q'(y)}}{\frac{M_\theta}{M}} = -\frac{1}{y\varepsilon^c(y,\phi)}\frac{M}{M_\theta},$$

where we used the expression, analogous to (50), for the compensated elasticity

$$\varepsilon^c(y,\phi) = \frac{1/y}{M_c(y,\theta(y;\phi),\phi) + \frac{M_y(y,\theta(y;\phi),\phi)}{M(y)} - \frac{q''(y)}{q'(y)}}.$$

Using this, (60) becomes

$$\int_{\Phi} f(\theta(y;\phi)|\phi)\frac{t'(y)}{q'(y)}\frac{M}{M_\theta}dP(\phi) \;=\; yh(y)\frac{t'(y)}{q'(y)}\int_{\Phi}\frac{f(\theta(y;\phi)|\phi)\theta'(y;\phi)}{h(y)}\varepsilon^c(y,\phi)dP(\phi)$$
$$= yh(y)\frac{t'(y)}{q'(y)}\bar{\varepsilon}^c(y), \tag{61}$$

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 C.1, the weights on the right-hand side of the first equation in (61) equal the density $p(\phi|y)$ of ϕ conditional on y).

As for the right-hand side of (6), we have

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

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

$$\int_{\Phi}\int_y^{\infty}\left(\beta_{\theta(z;\phi),\phi}-1+\eta(z,\phi)\frac{t'(z)}{q'(z)}\right)f(\theta(z;\phi)|\phi)\theta'(z;\phi)dzdP(\phi)$$
$$= -\int_y^{\infty}\int_{\Phi}\left(1-\beta_{\theta(z;\phi),\phi}\right)dP(\phi|z)dH(z)+\int_y^{\infty}\int_{\Phi}\eta(z,\phi)dP(\phi|z)\frac{t'(z)}{q'(z)}dH(z)$$
$$= -\int_y^{\infty}\left(1-\bar{\beta}_z\right)dH(z)+\int_y^{\infty}\bar{\eta}(z)\frac{t'(z)}{q'(z)}dH(z) \tag{62}$$

where $\bar{\eta}(y)$ is the average income effect and $\bar{\beta}_y$ the average social welfare weight at y. Equating (61) and (62) and following the same steps as in Appendix B delivers the results.

C.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 (6) as

$$\frac{\partial}{\partial\tau}(1-H^c(y;\tau))\bigg|_{\tau=0} = \frac{\partial}{\partial\tau}\int_{\theta^c(y;\tau)}^{\infty}\Gamma(\varphi_\theta^c(q;\tau)|\theta)dF(\theta)\bigg|_{\tau=0}$$
$$= -\Gamma(\varphi_{\theta(y)}(q)|\theta(y))f(\theta(y))\left.\frac{\partial\theta^c(y;\tau)}{\partial\tau}\right|_{\tau=0}+\int_{\theta(y)}^{\infty}\gamma(\varphi_\theta(q)|\theta)\left.\frac{\partial\varphi_\theta^c(q;\tau)}{\partial\tau}\right|_{\tau=0}dF(\theta) \tag{63}$$

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 B. Integrating the variation defined in (11), we have

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

so we can write

$$\left.\frac{\partial \varphi^c_\theta(q;\tau)}{\partial \tau}\right|_{\tau=0} = (t(y(\theta)) - t(0))\frac{\partial \varphi^c_\theta(q)}{\partial(q(y) - q(0))}.$$

Using this, the second term in (63) becomes

$$\int_{\theta(y)}^{\infty}(t(y(\theta)) - t(0))\gamma(\varphi_\theta(q)|\theta)\frac{\partial \varphi^c_\theta(q)}{\partial(q(y(\theta)) - q(0))}dF(\theta) = \int_y^\infty \frac{t(z) - t(0)}{q(z) - q(0)}\rho^c(z)dH(z) \tag{64}$$

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

The right-hand side of (6) becomes

$$\int_{\theta(y)}^{\infty}\int_{-\infty}^{\varphi_\theta(q)}(\beta_{\theta,\varphi} - 1)d\Gamma(\varphi|\theta)dF(\theta) - \frac{\partial}{\partial I}\int_{\theta(y)}^{\infty}\int_0^{y(\theta;I)} t'(z)dz\Gamma(\varphi_\theta(q;I)|\theta)dF(\theta). \tag{65}$$

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

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

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 (65) equals

$$-\int_{\theta(y)}^{\infty}\frac{\partial y(\theta;I)}{\partial I}t'(y(\theta))\Gamma(\varphi_\theta(q)|\theta)dF(\theta) - \int_{\theta(y)}^{\infty}(t(y(\theta)) - t(0))\gamma(\varphi_\theta(q)|\theta)\frac{\partial \varphi_\theta(q;I)}{\partial I}dF(\theta). \tag{66}$$

The first term here is again standard and the same as in Section 5. The second term in (66) can be combined with (64) to deliver the uncompensated extensive-margin response, i.e.

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

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

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

and hence the condition in the main text.