Advances in Dynamic Optimal Taxation

By Narayana R. Kocherlakota[1]

This paper surveys the recent literature concerning the structure of optimal taxes in dynamic economies in which agents are privately informed about skills and effort.

Keywords: optimal taxation, private information, dynamic

1. Introduction

This paper is about a now classic question in macroeconomics and public finance. A government needs to finance an exogenously given stochastic process of purchases. How do the optimal taxes behave over dates and states?

There is a large literature on this question that uses what I will term the Ramsey approach. Under this approach, the government is restricted to use *linear* taxes on *current* variables like capital and labor income. The government's main goal is then to minimize the social distortions associated with linearity.

The main weakness in the Ramsey approach is obvious: there is no explicit motivation for the restrictions that drive the analysis. Why should the government be restricted to using linear taxes? Virtually all real-world labor income tax codes display nontrivial amounts of nonlinearity. Why should the government be restricted to using functions of current variables? At least in the United States, federal taxes depend in complicated ways on the full history of assetholdings (through the use of basis calculations) and federal (social security) transfers depend in complicated ways on the history of labor incomes.

This weakness in the Ramsey approach has led to a new literature about the optimal taxation question. Under the new approach, instead of specifying an arbitrary set of tax instruments, the investigator first specifies the informational and/or enforcement frictions that limit the government's ability to extract revenue. Then, the investigator designs a tax system that implements a constrained Pareto optimal allocation given these frictions..

A key step in using this approach is knowing what frictions constrain the government. At least so far, the new literature is basically a dynamic extension of Mirrlees (1971)'s basic framework, which centers on two key insights. The first is that a - if not the - major risk in life is skill *risk*. Some people are born with the ability to produce large amounts of output with relatively little effort. Others are not. Over time, some people lose their ability to generate output (because of back pain or mental illness). Others do not.

It is straightforward to design a tax system that fully insures people against skill shocks: tax everyone at 100% and split the proceeds evenly across people, regardless of what they produced. Such a tax system works well, as long as skills and effort are fully observable. With this kind of information in hand, the government can simply order the highly-skilled people to work hard.

Mirrlees' second insight explains why this kind of system will not work well: skills and effort are often private information. In terms of the natal risks that Mirrlees himself stresses, people with high I.Q.'s can readily mimic people with low I.Q.'s. In terms of post-natal skill risks, it is easy to fake back pain or mental illness - there are few non-manipulable physical signs of these ailments.

The equal-split tax scheme by the government can still be used if skills and effort are private information. However, it is no longer desirable. If people are being taxed at 100%, all of the high-skilled people will work as if they are low-skilled. Because it can no longer directly command the high-skilled to work, the government now has to use the tax system to achieve two conflicting goals. As before, the tax system must insure people against skill shocks. But it must also provide *incentives* to motivate the skilled to produce more income than the low-skilled.

It is important to emphasize that the economic forces in this new Mirrlees literature are fundamentally different from the forces in the old Ramsey literature. Under the Ramsey approach, the government is banned from using lump-sum taxes. Its objective is to minimize the social costs associated with using linear taxes. Under the Mirrlees approach, the government is allowed to use lump-sum taxes, but chooses not to. Its objective is to find the optimal trade-off between incentives and insurance.

My goal in this essay is to provide a partial survey of the lessons of the recent dynamic Mirrlees literature. (As we shall see, "recent" really means "recent" - the first published paper in this area appeared in 2003.) The essay is divided into three sections. In the first and larger part of the essay, I focus on wealth taxation. I assume that agents are affected by idiosyncratic shocks to their skills, and that skills and effort are both private information. I impose no restriction on the time series behavior of skills or government purchases. However, I assume (as I do throughout the essay) that preferences are additively separable over time and between consumption and labor.

The main result in this first part of the paper is that it is possible to design an optimal tax system that is *linear* in *current* wealth. This tax system has two surprising features. First, the tax on wealth brought into a given period must depend on the labor income realized in that period. The optimal taxes are *regressive*, so that those who earn surprisingly high labor income pay a lower wealth tax. Those who earn surprisingly low labor incomes pay a high wealth tax. As I explain later, this regressive feature helps the tax system provide better incentives. Second, the average tax rate on wealth across individuals is always zero, regardless of the realization of government purchases. Optimal wealth taxes redistribute resources from the surprisingly low-skilled to the surprisingly high-skilled.

In the latter portion of the essay, I consider three other types of taxes. The first is estate taxes. (I turn to them next not because of their importance, but rather because they are closely linked to wealth taxes.) I discuss recent work by Farhi and Werning (2005) on this subject. They show that it is optimal for estate taxes to be negative - estates should be subsidized, not taxed. They also show that the estates of the poor should be subsidized at a higher rate than those of the rich.

I then consider inflation taxes. Here, the key source is a paper by da Costa and Werning (2005). They consider a world in which agents have fixed skills and preferences over consumption, real balances and labor. They provide conditions on these preferences such that the Friedman Rule (zero nominal interest rates) is optimal. These conditions are sufficiently weak to nest the two most popular models of money demand: cash-in-advance and shopping-time. They also show that the welfare losses from suboptimal monetary policy - that is, positive nominal interest rates - can be considerably higher in a world in which agents have heretogeneous skills.

Inflation taxes, estate taxes, and wealth taxes are all really forms of taxation of different kinds of consumption. This focus on consumption may seem surprising: after all, the heart of the Mirrlees model is the informational friction concerning an individual's willingness/ability to supply *labor*. Mirrlees' (1971) original analysis was designed to understand how an income tax schedule should be designed to resolve the tension between efficiency and equity generated by this friction. Extending his analysis to dynamic settings has proven challenging. I use the work of Battaglini and Coate (2005) discuss the properties of optimal labor income taxes in a particular setting, in which skills follow a Markov chain with a two-point support and agents are risk-neutral over consumption. The main result of their paper is that under an optimal tax system, labor income taxes are zero for anyone who has ever experienced a high-skill shock in the past.

2. Wealth Taxes

In this section, I discuss the structure of optimal wealth taxes in a wide class of economic environments. The discussion in the first five subsections closely follows Kocherlakota (2005)'s construction of an optimal tax system that is linear in wealth. The sixth subsection discusses other forms of optimal wealth taxes.

2.1 Economic Environment

In this subsection, I describe the environment. The economy lasts for T periods, where T may be infinite, and has a unit measure of agents. The economy is initially endowed with K_1^* units of the single capital good. There is a single consumption good that can be produced by capital and labor. The agents have identical preferences. A given agent has von-Neumann-Morgenstern preferences, and ranks deterministic sequences according to the function:

$$(1) \qquad \sum_{t=1}^{T} \beta^{t-1}\{u(c_t) - v(l_t)\}, 1 > \beta > 0$$

where $c_t \in R_+$ is the agent's consumption in period t, and $l_t \in R_+$ is the agent's labor in period t. I assume that $u', -u'', v'$, and v'' all exist and are positive. I also assume that the momentary utility functions u and v are bounded from above and below.

There are two kinds of shocks in the economy: public aggregate shocks and private idiosyncratic shocks. The first kind of shocks works as follows. Let Z be a finite set[2], and let μ_Z be a probability measure over the power set of Z^T that assigns positive probability to all subsets of Z^T. At the beginning of period 1, an element z^T of Z^T is drawn according to μ_Z. The random vector z^T is the sequence of public aggregate shocks; z_t is the realization of the shock in period t.

The idiosyncratic shocks work as follows. Let Θ be a Borel set in R_+, and let μ_Θ be a probability measure over the Borel subsets of Θ^T. At the beginning of period 1, an element of θ^T is drawn for each agent according to the measure μ_Θ. Conditional on z^T, the draws are independent across agents. I assume that a law of large numbers applies: conditional on any z^T, the measure of agents in the population with type θ^T in Borel set B is given by $\mu_\Theta(B)$.

Both μ_Θ and μ_Z are common knowledge.[3] Any given agent learns the realization of the public shock z_t and his own idiosyncratic shock θ_t at the beginning of period t and not before. Thus, at the beginning of period t, the agent knows his own private history $\theta^t = (\theta_1, ..., \theta_t)$ and the history of public shocks $z^t = (z_1, ..., z_t)$. This implies that his choices in period t can only be a function of this history.

What is the economic impact of these shocks? First, the shocks determine skills. In period t, an agent produces *effective labor* y_t according to the function:

$$y_t(\theta^T, z^T) = \phi_t(\theta^T, z^T) l_t(\theta^T, z^T)$$

where $\phi_t : \Theta^T \times Z^T \to (0, \infty)$ and is (θ^t, z^t)-measurable I assume that an agent's effective labor is observable at time t, but his labor input l_t is known only to him. I refer to ϕ_t as an agent's skill in history (θ^t, z^t). The idea here is that everyone shows up for eight hours per day, and their output at the end of the day is observable. However, it is hard to monitor how hard they are working and what kinds of shocks they face during the day.

The public aggregate shocks influence the aggregate production function in the following way. I define an allocation in this society to be (c, y, K) where:

(2) $K : Z^T \to R_+^{T+1}$

(3) $c : \Theta^T \times Z^T \to R_+^T$

(4) $y : \Theta^T \times Z^T \to R_+^T$

(5) K_{t+1} is z^t-measurable

(6) (c_t, y_t) is (θ^t, z^t)-measurable

Here, $y_t(\theta^T, z^T)$ $(c_t(\theta^T, z^T))$ is the amount of effective labor (consumption) assigned in period t to an agent with type θ^T, given that the public aggregate shock sequence is z^T. K_{t+1} is the per-capita amount of capital carried over from period t into period $(t+1)$.

As mentioned above, I assume that the initial endowment of capital is K_1^*. I assume that the government has exogenous per-capita purchasing needs $G_t : Z^T \to R_+$ in period t, where G_t is z^t-measurable. I define an allocation (c, y, K) to be *feasible* if for all t, z^T:

(7) $C_t(z^T) + K_{t+1}(z^T) + G_t(z^T) \leq F_t(K_t, Y_t, z^T) + (1-\delta)K_t(z^T)$

(8) $C_t(z^T) = \int_{\theta^T \in \Theta^T} c_t(\theta^T, z^T) d\mu_\Theta$

(9) $Y_t(z^T) = \int_{\theta^T \in \Theta^T} y_t(\theta^T, z^T) d\mu_\Theta$

(10) $K_1 \leq K_1^*$

Here, C_t and Y_t represent per-capita consumption and per-capita effective labor. (Note that (C_t, Y_t) are z^t-measurable.) The aggregate production function $F_t : R_+^2 \times Z^T \to R_+$ is assumed to be strictly increasing, weakly concave, homogeneous of degree one, continuously differentiable with respect to its first two arguments, and z^t-measurable with respect to its last argument.

The idiosyncratic shock θ_t is privately observable. This means that not all physically feasible allocations are actually achievable in this society. For example, consider a physically feasible allocation in which all agents always consume the same amount, but agents with high skills are required to produce a larger amount of effective labor. This allocation is inconsistent with the restriction that skills are privately observable, because agents with high skills will always mimic the low-skilled agents and will not generate a large amount of effective labor.

To find the set of achievable allocations, we exploit the Revelation Principle. Under the Revelation Principle, we can restrict attention to *direct* mechanisms in which agents make reports that lie in Θ to a central authority at each date. The Revelation Principle says that any allocation that is achievable in this setting with private information is an equilibrium allocation of some direct mechanism in which agents find it optimal to tell the truth. Note that the Revelation Principle does not say that different types of agents are necessarily treated differently.

Formally, we proceed as follows. A *reporting strategy* $\sigma : \Theta^T \times Z^T \rightarrow \Theta^T \times Z^T$, where σ_t is (θ^t, z^t)-measurable and $\sigma(\theta^T, z^T) = (\theta^{T\prime}, z^T)$ for some $\theta^{T\prime}$ (thus, the agent is required to report truthfully about the publicly observable variables). Let Σ be the set of all possible reporting strategies, and define:

(11) $$W(.; c, y) : \Sigma \rightarrow R$$

(12) $$W(\sigma; c, y) = \sum_{t=1}^{T} \beta^{t-1} \int_{Z^T} \int_{\Theta^T} \{u(c_t(\sigma)) - v(y_t(\sigma)/\phi_t)\} d\mu_\Theta d\mu_Z$$

to be the expected utility from reporting strategy σ, given an allocation (c, y). (Note that the integral over Z could also be written as a sum.) Let σ_{TT} be the truth-telling strategy $\sigma_{TT}(\theta^T, z^T) = (\theta^T, z^T)$ for all θ^T, z^T. Then, an allocation (c, y, K) is *incentive-compatible* if:

$$(13) \quad W(\sigma_{TT}; c, y) \geq W(\sigma; c, y) \text{ for all } \sigma \text{ in } \Sigma$$

An allocation which is incentive-compatible and feasible is said to be incentive-feasible.

An *optimal* allocation is an allocation (c, y, K) that solves the problem of maximizing:

$$(14) \quad \sum_{t=1}^{T} \beta^{t-1} \int_{Z^T} \int_{\Theta^T} \{u(c_t) - v(y_t/\phi_t)\} d\mu_\Theta d\mu_Z$$

subject to (c, y, K) being incentive-feasible. Under this notion of optimality, all agents are treated symmetrically from an ex-ante perspective.[4]

2.2 Aspects of the Economic Environment

The class of economic environments described above is large in the sense that skills ϕ_t may be governed by any data generation process. This is desirable, because there continues to be an ongoing empirical debate about the time series properties of wages. In particular, the class of environments is sufficiently general that it allows ϕ_t to exhibit persistence of any kind. More subtly, $Var(\phi_t|z^t)$ can be countercyclical. This kind of countercyclical variation in wage inequality has been documented by many authors (including Storesletten, Telmer and Yaron (2003)), and plays an important role in understanding the behavior of the market price of risk.

The environment is unusually general in at least two respects. It is common when analyzing models with private information to impose restrictions sufficient to guarantee that some form of Spence-Mirrlees single-crossing holds. I have not done so. The relevant results about wealth taxes are valid without single-crossing.[5] Similarly, it is common in dynamic macroeconomics to impose sufficient structure to guarantee recursivity with respect to a small number of state variables. Here, I have not imposed such structure. Recursivity plays no useful role in the characterization of optimal wealth taxes; it is easier to understand the relevant results through the lens of the sequence formulation.

The class of economic environments is special in that preferences are restricted to be additively separable between consumption and labor. With additive separability, any agent's intertemporal marginal rate of substitution for consumption is observable to the planner. Without additive separability, the intertemporal marginal rate of substitution is private information to the agent (because it depends on the agent's level of effort). This kind of private information greatly complicates the nature of optimal wealth taxes. In particular, if utility is nonseparable between consumption and labor, and the support Θ of the idiosyncratic shocks is finite, there is no optimal tax system in which taxes are differentiable in wealth (Kocherlakota (2004)).

2.3 An Intertemporal Characterization of Socially Optimal Allocations

In this subsection, I restate Proposition 1 from Kocherlakota (2005). This proposition establishes that any optimal allocation must satisfy a particular first order condition (similar to that derived in Theorem 1 of Golosov, Kocherlakota, and Tsyvinski (2003) and in Rogerson (1985)).

PROPOSITION 1: *Suppose (c^*, y^*, K^*) is an optimal allocation and that there exists $t < T$ and scalars M^+, M_+ such that $M^+ \geq c_t^*, c_{t+1}^*, K_{t+1}^* \geq M_+ > 0$ almost everywhere. Then there exists $\lambda_{t+1}^* : Z^T \rightarrow R_+$ such that:*

(15) λ_{t+1}^* *is* z^{t+1}*-measurable*

(16) $\lambda^*_{t+1} = \beta[E\{u'(c^*_{t+1})^{-1}|\theta^t, z^{t+1}\}]^{-1}/u'(c^*_t)$ *a.e.*

(17) $E\{\lambda^*_{t+1}(1-\delta+F^*_{K,t+1})|z^t\} = 1$ *a.e.*

where $F^*_{K,t+1}(z^T) = F_{K,t+1}(K^*_{t+1}(z^T), Y^*_{t+1}(z^T), z^T)$ *for all* z^T.

PROOF. See Kocherlakota (2005).

In this proposition, λ^*_{t+1} is the shadow price of consumption in public history z^{t+1} relative to consumption in public history z^t. (If one wrote out the full planner's problem, λ^*_{t+1} would be the ratio of the multiplier on the resource constraint in z^{t+1} to the multiplier on the resource constraint in z^t.) This shadow price does not depend on idiosyncratic shocks.

Hence, the content of this proposition is twofold. First, it establishes that:

$$\beta\{E(u'(c^*_{t+1})^{-1}|\theta^t, z^{t+1})\}^{-1}/u'(c^*_t)$$

is independent of θ^t. This result is obviously true without private information, because in that case the optimal c^*_t is independent of θ^t. In the presence of private information, it is generally optimal to allow c^*_t to depend on θ^t in order to require high-skilled agents to produce more effective labor. Proposition 1 establishes that even in that case, the *harmonic* mean of $\beta u'(c^*_{t+1})/u'(c^*_t)$, conditional on θ^t and z^{t+1}, is independent of θ^t. The second part of the proposition is more obvious - it simply states that the shadow price λ^*_{t+1} can be used to determine the optimal level of capital accumulation between period t and period $(t+1)$.

A complete proof of this proposition appears in Kocherlakota (2005). Here is a sketch of the proof, for the case in which Θ is a finite set. Suppose (c^*, y^*, K^*) is an interior optimum, and fix a positive probability history $(\overline{\theta}^t, \overline{z}^{t+1})$. Consider a perturbation similar to that used by Rogerson (1985), and define a new consumption allocation c' to be the same as c^* except that:

$$u(c'_t(\overline{\theta}^t, \overline{z}^t)) = u(c^*_t(\overline{\theta}^t, \overline{z}^t)) - \varepsilon$$
$$u(c'_{t+1}(\overline{\theta}^t, \theta, \overline{z}^{t+1})) = u(c^*_{t+1}(\overline{\theta}^t, \theta, \overline{z}^{t+1})) + \beta^{-1}\varepsilon \text{ for all } \theta \text{ in } \Theta$$

where ε is small and positive. For the agents with skill history $\overline{\theta}^t$, this change is designed to reduce momentary utility in period t by ε, and increase momentary utility in period $(t+1)$ (given any continuation skill history) by $\beta^{-1}\varepsilon$.

The key to the proof is that, by construction, any sequence of reports generates the same (ex-ante) utility under (c', y^*) as under (c^*, y^*). Hence, the ranking of reporting strategies must be the same under (c', y^*) as under (c^*, y^*). It follows that, if (c^*, y^*) is incentive-compatible, the new plan (c', y^*) must also be incentive-compatible.

Similarly, the new consumption plan (c', y^*) does not change the planner's objective. It follows that the new timing of consumption payments cannot result in extra resources for the planner or (c^*, y^*, K^*) is not optimal. There are $\mu(\overline{\theta}^t)$ agents who have shock history $\overline{\theta}^t$. Hence, for small ε, the new plan frees up $\mu(\overline{\theta}^t)\varepsilon/u'(c^*_t(\overline{\theta}^t, \overline{z}^t))$ units of consumption in period t. Similarly, it costs $\beta^{-1}\varepsilon\sum_{\theta\in\Theta}\mu(\overline{\theta}^t, \theta)/u'(c^*_{t+1}(\overline{\theta}^t, \theta, \overline{z}^{t+1}))$ in period $(t+1)$. By saving the extra period t consumption into period $(t+1)$, the planner has extra resources after paying the costs unless:

$$(18) \quad \mu(\overline{\theta}^t)/u'(c^*_t(\overline{\theta}^t, \overline{z}^t)) \leq \beta^{-1}\lambda_{t+1}(\overline{z}^{t+1})\sum_{\theta\in\Theta}\mu(\overline{\theta}^t, \theta)/u'(c^*_{t+1}(\overline{\theta}^t, \theta, \overline{z}^{t+1}))$$

where, as in Proposition 1, $\lambda_{t+1}(\overline{z}^{t+1})$ is the shadow price of consumption in history $\overline{z}^{t+1}$ relative to consumption in history $\overline{z}^t$. Making the same argument with ε small and negative implies the reverse inequality. In other words, at an optimum, the two sides of (18) are equated, which implies Proposition 1.

It is important to note that even if θ^T is public information (so that there is no incentive problem), Proposition 1 is still valid. In this case, full insurance is possible and $u'(c_t^*)$ is deterministic for all t. Proposition 1 immediately implies the standard first order condition:

$$u'(c_t^*) = \beta E\{(1 - \delta + F_{K,t+1}^*)u'(c_{t+1}^*)|z^t\}$$

Thus, the incentive problem does not create the restriction in Proposition 1. Rather, the incentive problem determines the variance of the marginal utility process that gets plugged into the formula in Proposition 1.

Proposition 1 immediately implies the following corollary.

COROLLARY 1: *Suppose* $\Pr(Var(u'(c_{t+1}^*)|z^{t+1}, \theta^t) > 0) > 0$. *Then:*

$$\beta E\{u'(c_{t+1}^*)(1 - \delta + F_{K,t+1})|z^t, \theta^t\} > u'(c_t^*) \tag{19}$$

PROOF. Jensen's inequality.

The hypothesis in this corollary is that the incentive problem is such that it is optimal for $u'(c_{t+1}^*)$ to be a nondegenerate function of θ_{t+1}. It will not be satisfied, for example, if, conditional on θ^t, ϕ_{t+s} is independent of θ_{t+1} for all $s \geq 1$. In this case, the agent knows his future path of skills with certainty and the inequality in Corollary 1 becomes an equality. More generally, though, it is optimal to have a wedge between the intertemporal marginal rate of substitution and the intertemporal marginal rate of transformation: an individual's marginal expected utility from selling capital tomorrow exceeds his marginal disutility from buying capital today.

2.4 The Distinction between Wedges and Taxes

Agents decide how much capital to bring into period $(t+1)$ using information available up through period t. Hence, it seems logical that the period $(t+1)$ capital tax itself should be a function only of information up through period t. Under this assumption, the optimal capital tax would be positive. To see this, note that if agents can buy and sell capital in a competitive market subject to a tax, they face the following first order condition:

$$(20)\quad \beta E\{u'(c_{t+1})(1-\delta+F_{K,t+1})(1-\tau_{t+1}^{k})|\theta^t, z^t\} = u'(c_t).$$

where τ_{t+1}^{k} is the marginal tax rate on capital. If τ_{t+1}^{k} is (θ^t, z^t)-measurable, then it must be larger than 0 if the equilibrium allocation is to be optimal.

I show in this subsection that, at least for tax systems that are differentiable in wealth, the above natural logic starts from a wrong premise: even though the capital accumulation decision is made in period t, the tax on capital accumulation must depend on period $(t+1)$ information. To see this, consider the following example (which is similar to ones described in Albanesi and Sleet (2005), Golosov and Tsyvinski (2005), and Kocherlakota (2005)). The example has two periods, with no shocks in period 1, and two equally likely possible realizations of skills in period 2.

Specifically, let $u(c) = \ln(c)$, $v(l) = l^2/2$, and $\beta = 1$. Suppose too that $T = 2, \Theta = \{0,1\}$, $Z = \{1\}$ (so that there are not aggregate shocks), $F(K,Y) = rK + wY$, and $\delta = 1$. As well, suppose $\phi_1(\theta) = 1$, $\phi_2(\theta, z) = \theta$, $v(l) = l^2/2$, and $\Pr(\theta_2 = 1) = 1/2$. Set $G = 0$. Let (c_{2i}, y_{2i}) denote consumption and effective labor when $\theta = i$. Then, we can re-write the planner's problem as:

$$\max_{c_1, c_{2h}, c_{2l}, y_1, y_{2h}, K_2} \ln(c_1) - y_1^2/2 + \ln(c_{21})/2 + \ln(c_{20})/2 - y_{20}^2/4$$

$s.t.\ c_1 + K_2 = rK_1 + wy_1$

$c_{21}/2 + c_{20}/2 = rK_2 + wy_{21}/2$

$\ln(c_{21}) - y_{21}^2/2 \geq \ln(c_{20})$

$c_{21}, c_{20}, y_{21}, K_2, y_1 \geq 0$

(In this statement of the problem, I have set $y_{20} = 0$, as would be true in a social optimum.) The solution to this problem must satisfy the following first order conditions:

$c_1^* + K_2^* = rK_1 + wy_1^*$

$c_{21}^*/2 + c_{20}^*/2 = rK_2^* + wy_{21}^*/2$

$\ln(c_{21}^*) - y_{21}^{*2}/2 = \ln(c_{20}^*)$

$1/c_1^* = r/[0.5c_{21}^* + 0.5c_{20}^*]$

$w/c_{21}^* = y_{21}^*$

$y_1^* = w/c_1^*$

Note that in the social optimum, the highly-skilled agent is indifferent between acting high-skilled and acting low-skilled.

Now consider a tax system τ such that if an agent has capital holdings k_2 at the beginning of period 2, and effective labor y_2 in period 2, then in period 2, he pays taxes $\tau^k(k_2) + \tau^y(y_2)$, where τ^k is differentiable with respect to k_2. The additive separability guarantees that the marginal tax rate on capital is known at the time the agent invests. Given this tax system, an equilibrium in this economy is a specification of $(\widehat{c}_1, \widehat{c}_{21}, \widehat{c}_{20}, y_1, \widehat{y}_{21}, \widehat{k}_2)$

such that it solves:

$$\max_{c_1,y_1,c_{2h},c_{2l},y_{2h},k_2} \ln(c_1) - y_1^2/2 + \ln(c_{21})/2 + \ln(c_{20})/2 - y_{21}^2/4$$

$$s.t.\ c_1 + k_2 = rk_1 + wy_1$$

$$c_{21} = rk_2 + wy_{21} - \tau^k(k_2) - \tau^y(y_{21})$$

$$c_{20} = rk_2 - \tau^k(k_2) - \tau^y(0)$$

$$k_2, c_{21}, c_{20}, y_{21}, y_1 \geq 0$$

and markets clear:

$$\begin{aligned} \widehat{c}_1 + \widehat{k}_2 &= rk_1 + wy_1 \\ \widehat{c}_{21}/2 + \widehat{c}_{20}/2 &= r\widehat{k}_2 + w\widehat{y}_{21}/2 \end{aligned}$$

Note that in equilibrium, $\tau^k(\widehat{k}_2)+0.5\tau^y(\widehat{y}_{21})+0.5\tau^y(0) = 0$, which is the government's budget constraint.

Consider a socially optimal allocation in which $y_{21}^* > 0$. I claim that there is no tax system such that this socially optimal allocation is an equilibrium. Suppose the claim is false. In the socially optimal allocation, a high-skilled agent is indifferent between working y_{21}^* and working 0. Hence, in equilibrium, the agent is indifferent between setting $(y_{21}, k_2) = (y_{21}^*, K_2^*)$ and setting $(y_{21}, k_2) = (0, K_2^*)$. Now suppose the agent chooses $k_2 = K_2^* + \varepsilon$ and sets $y_{21}^* = 0$. His utility from this budget-feasible plan is:

$$\ln(c_1^* - \varepsilon) + \ln(c_{2l}^* + r\varepsilon - \tau^k(K_2^* + \varepsilon))$$

To be in equilibrium, the socially optimal allocation must satisfy the agent's intertemporal first order condition :

$$1/c_1^* = r(1 - \tau^{k\prime}(K_2^*))[0.5/c_{21}^* + 0.5/c_{20}^*]$$

which implies that:

$$1/c_1^* < r(1 - \tau^{k\prime}(K_2^*))/c_{20}^*$$

which means that:

$$\begin{aligned} & \ln(c_1^* - \varepsilon) + \ln(c_{20}^* + r\varepsilon - \tau^k(K_2^* + \varepsilon)) \\ > \; & \ln(c_1^*) + \ln(c_{20}^*) \\ = \; & \ln(c_1^*) + 0.5 \ln(c_{21}^*) + 0.5 \ln(c_{20}^*) - y_{21}^{*2}/2 \end{aligned}$$

Given this tax system, the agent can improve upon the socially optimal allocation by saving a little more and then not working.

Intuitively, we have set the capital tax rate to guarantee that the agent does not save too much or too little - assuming that he tells the truth about his type. The optimal allocation pushes the agent to be indifferent between telling the truth or lying. If he saves a little bit more, and wealth effects are nonzero, then he will prefer to pretend to be disabled when he is actually abled. The *joint* or *double* deviation of saving too much and then shirking beats saving the right amount and telling the truth about one's type.

What this means is that the wedge derived in the previous subsection does not immediately translate into a conclusion about taxes. We have to find a different way to make a connection between the wedge and tax rates.

2.5 A General Optimal Linear Tax System on Wealth

In this subsection, I describe how to design an optimal tax system that is linear in current wealth. I assume throughout that agents hold a positive amount of capital, so that they are interior in their capital-holdings. I begin by reconsidering the two-period example and designing an optimal system in that context. I then discuss how to design a similar system for the general model described earlier.

2.5.1 Optimal System: 2 Period Example

Recall the two-period example discussed in the prior subsection. Let $(c_1^*, c_{21}^*, c_{20}^*, K_2^*, y_{21}^*)$ be a socially optimal allocation. Define a period 2 tax code by:

$$\begin{aligned}\tau(k_2, y_2) &= \tau_1 k_2 + \alpha_1 \text{ if } y_2 > 0 \\ &= \tau_0 k_2 + \alpha_0 \text{ if } y_2 = 0\end{aligned}$$

where α_i and τ_i are set so as to satisfy:

$$\beta(1-\tau_i)u'(c_{2i}^*)r = u'(c_1^*) \text{ for } i = 0, 1$$

$$wy_{2i} + k_2^*(1-\tau_i)r - \alpha_i = c_{2i}^* \text{ for } i = 0, 1$$

Given this tax system, the individual's choice problem is:

$$\max_{c_1, y_1, c_{2h}, c_{2l}, y_{2h}, k_2} \ln(c_1) - y_1^2/2 + \ln(c_{2h})/2 + \ln(c_{2l})/2 - y_{2h}^2/4$$

$$s.t.\ c_1 + k_2 = rk_1 + wy_1$$

$$c_{21} = r(1-\tau_1)k_2 + wy_{21} - \alpha_1 \text{ if } y_{21} > 0$$

$$c_{21} = r(1-\tau_0)k_2 - \alpha_0 \text{ if } y_{21} = 0$$

$$c_{20} = r(1 - \tau_0)k_2 - \alpha_0$$

$$k_2, c_{20}, c_{21}, y_{21}, y_1 \geq 0$$

I claim that the socially optimal allocation solves this choice problem.

Why is this true? Suppose that the agent works $y_{21} > 0$ in period 2 when abled. Then, his solution for his other choice variables is characterized by:

$$\begin{aligned} 1/c_1 &= r[0.5(1 - \tau_1)/c_{21} + 0.5(1 - \tau_0)/c_{20}] \\ w/c_1 &= y_1 \\ c_1 + k_2 &= rk_1 + wy_1 \\ c_{21} &= r(1 - \tau_1)k_2 + wy_{21} - \alpha_1 \\ c_{20} &= r(1 - \tau_0)k_2 - \alpha_0 \end{aligned}$$

The starred allocation satisfies these first order conditions.

What if the agent works 0 in period 2 when abled? Then, his first order conditions become:

$$1/c_1 = r(1 - \tau_{k0})/c_{20}$$

$$w/c_1 = y_1$$

$$c_1 + k_2 = rk_1 + wy_1$$

$$c_{21} = r(1 - \tau_0)k_2 - \alpha_0$$

$$c_{20} = r(1 - \tau_0)k_2 - \alpha_0$$

Setting $(c_1, y_1, k_2, c_{21}, c_{20})$ equal to $(c_1^*, y_1^*, k_2^*, c_{20}^*, c_{20}^*)$ satisfies these first order conditions.

Hence, the agent is indifferent between working y_{21}^* in period 2 (when able) and not working in period 2.

Thus, we can implement the optimal allocation using a tax schedule that is linear in capital income and nonlinear in labor income. In this implementation, even though agents make their decisions about capital-holdings before period 2, the tax rate on capital must depend on the amount of labor income earned in period 2. The tax rate on capital brought into period 2 is necessarily stochastic from the point of view of period 1.

2.5.2 Optimal System: The General Case

The key to the optimal tax system in the two-period example is that wealth tax rates were set so as to equate agents' ex-post after-tax marginal rates of substitution with the societal marginal rate of transformation. The general optimal tax system works in the same way. Let (c^*, y^*, K^*) be a socially optimal allocation and let $\lambda_{t+1}^*(z^T)$ be the shadow price of consumption in public history z^{t+1} relative to consumption in public history z^t. Define the tax *rate* τ_{t+1} on wealth brought into period $(t+1)$ by:

$$(21) \quad (1 - \tau_{t+1}(\theta^T, z^T)) = \frac{\beta^{-1} u'(c_t^*(\theta^T, z^T))\lambda_{t+1}^*(z^T)}{u'(c_{t+1}^*(\theta^T, z^T))}$$

This tax rate equates the agent's after-tax ex-post marginal rate of substitution with the societal marginal rate of transformation (that is, λ).

Kocherlakota (2005) formally proves that, in conjunction with a well-designed labor income tax code, this tax system implements (c^*, y^*, K^*) as an equilibrium. The basic intuition for the proof is as follows. The agent's choice problem is to choose how much to work and how much to save. It is easy to design the labor taxes so that no agent ever chooses an effective labor sequence that is not associated with some type under the socially optimal allocation. (For example, such a choice could result in high fines.) Hence, the agent's problem is basically the same as jointly choosing whether to mimic some other type of person and which capital accumulation strategy to use.

Suppose that the agent chooses his effective labor as if his true type is $\theta^{T\prime}$. He receives a consumption sequence $c^*(\theta^{T\prime}, z^T)$ and generates an effective labor sequence $y^*(\theta^{T\prime}, z^T)$. The agent has the ability to change this consumption stream by buying more or less capital at any date t. But the tax system deters him from doing so; his first-order condition for savings is given by:

$$
\begin{aligned}
& u'(c_t^*(\theta^{T\prime}, z^T)) - \beta E\{(1 - \tau_{t+1}(\theta^{T\prime}, z^T))u'(c_{t+1}^*(\theta^{T\prime}, z^T)(1 - \delta + F_{K,t+1}^*)|z^t, \theta^T = \theta^{T\prime}\} \\
= \; & u'(c_t^*(\theta^{T\prime}, z^T)) - u'(c_t^*(\theta^{T\prime}, z^T))E\{\lambda_{t+1}^*(1 - \delta + F_{K,t+1}^*)|z^t\} \\
= \; & 0
\end{aligned}
$$

This is the key step in the proof: regardless of what type the agent pretends to be, he has no incentive to alter the resultant consumption profile by buying or selling capital. In other words, adding the ability to buy and sell capital, subject to the tax rate τ, does not change the agent's reporting problem. Since (c^*, y^*) is incentive-compatible, the agent finds it optimal not to mimic anyone else.

It is important to underscore the role of wealth taxes in this setting. It does not matter who owns the capital stock in this economy. Hence, the ability on the part of agents to trade capital is not socially useful; it only serves to undercut the ability of government to achieve desirable outcomes in this setting. Thus, one optimal tax system is for the government to take over the capital stock, and then tax any capital holdings by private agents at 100%. The tax system described above is much less draconian. Private citizens are allowed to hold capital. Instead of banning trade in capital, the linear tax system is carefully designed so that the agents choose not to deviate from their appropriate consumption profiles.

2.5.3 Properties of the Optimal System

What is the expected wealth tax rate in period $(t+1)$, conditional on (θ^t, z^{t+1})? To answer this question, we compute:

$$\begin{aligned} & E\{(1-\tau^*_{t+1}(\theta^T, z^T))|\theta^t, z^{t+1}\} \\ = \; & E\{\beta^{-1}\lambda^*_{t+1}u'(c^*_{t+1})^{-1}u'(c^*_t)|\theta^t, z^{t+1}\} \\ = \; & \beta^{-1}\lambda^*_{t+1}u'(c^*_t)E\{u'(c^*_{t+1})^{-1}|\theta^t, z^{t+1}\} \text{ by } (\theta^t, z^{t+1})\text{-measurability of } \lambda^*_{t+1}u'(c^*_t) \\ = \; & 1 \end{aligned}$$

where the last step follows from Proposition 1. Thus, the expected wealth tax rate is zero.

This result has a second, slightly more subtle, implication: wealth taxes are purely redistributional because the government raises no net revenue from them in any public history z^{t+1}. Suppose k^* is an equilibrium process of capital-holdings given that wealth taxes as a function of (θ^T, z^T) equal τ^{**}. Then, we can calculate the total revenue from wealth taxes in each public history:

$$\begin{aligned} & \int_{\theta^T \in \Theta^T} \tau^*_{t+1}(\theta^T, z^T)k^*_{t+1}(\theta^T, z^T)(1-\delta+MPK^*_{t+1}(z^T))d\mu_\Theta \\ = \; & (1-\delta+MPK^*_{t+1}(z^T))E(\tau^{**}_{t+1}k^*_{t+1}|z^{t+1}) \\ = \; & (1-\delta+MPK^*_{t+1}(z^T))E(E(\tau^{**}_{t+1}|\theta^t, z^{t+1})k^*_{t+1}|z^{t+1}) \\ = \; & 0 \end{aligned}$$

The key step in this calculation is the penultimate one, in which I exploit the Law of Iterated Expectations and the fact that k^*_{t+1} is (θ^t, z^t)-measurable.

Who pays the higher tax? This is also easy to see. Conditional on (θ^t, z^{t+1}), the variance in the wealth tax rate derives from the dependence of $u'(c^*_{t+1})^{-1}$ on θ_{t+1}. The after-tax rate $(1 - \tau^{**}_{t+1})$ is surprisingly high for agents with a surprisingly high $1/u'(c^*_{t+1})$ - that is, a high c^*_{t+1}. Intuitively, the high wealth tax rate on the unskilled is needed to deter agents from doing a joint deviation of saving too much and then working too little when skilled in the following period. In this way, the "regressive" nature of the optimal system - surprisingly poor agents face high wealth taxes - enables the government to provide better social insurance.

The expected optimal tax rate is zero. Nonetheless, for the tax to be optimal, there must be a wedge between the individuals' intertemporal marginal rates of substitution and the social marginal rate of transformation. The optimal tax system generates this wedge in a subtle fashion. The wealth tax rate is high exactly when agents have low consumption. This negative correlation between wealth tax rates and consumption means that the after-tax rate of return on capital is riskier because of taxes. This extra risk is what creates the wedge.

2.6 Other Forms of Optimal Wealth Taxes

It is well-known (Chari and Kehoe (1999)) that even if taxes are restricted to be linear, there may be many optimal tax systems. This indeterminacy is even more true if taxes are not restricted to be linear, but may be any function of current and past actions. In this section, I discuss some other types of optimal wealth taxes.

2.6.1 Non-Zero Expected Wealth Taxes

Albanesi and Sleet (2005) consider an economy in which there are no aggregate shocks and idiosyncratic skill shocks are independently and identically distributed over time. In this setting, they prove that it is possible to construct an optimal tax system in which taxes are restricted to be a function only of current wealth and of current labor income. They show that this optimal tax system is not necessarily linear in wealth, and the expected marginal tax rate on wealth need not be zero.

Albanesi and Sleet's tax system is shaped by two key forces. The first is that when skill shocks are i.i.d. over time, the socially optimal allocation has a simple intertemporal structure. In period t, each agent announces his skill realization. The planner responds to this realization by giving the agent some current consumption and promising the agent a certain amount of continuation utility. Then, in period $(t+1)$, agents are given consumption and future promised utility as a function only of their current skill shock and their promised utility from period t. In this way, because of the i.i.d. shock structure, the socially optimal allocation only depends on an agent's past reports through that agent's continuation utility (as originally described in Green (1987)).

This Markov structure is crucial to Albanesi and Sleet's implementation. In a market economy, there is typically a one-to-one mapping between wealth and continuation utility. Since optimal allocations depend on the past only through continuation utility, optimal tax systems need only condition on current wealth and current income, rather than conditioning on the full history of skill shocks.

There is a second, more subtle, feature to Albanesi and Sleet's (2005) analysis. In a centralized allocation, agents are simply given future promised utility and current consumption as a function of their reports. They cannot re-adjust these after the planner's allocation. In contrast, in a market economy, agents have the ability to trade between current consumption and future wealth (that is, promised utility). Albanesi and Sleet (2005) show that, at least when preferences are additively separable between consumption and labor, the optimal allocation is structured so that agents do not want to deviate in such intertemporal trades.

Why are optimal expected marginal taxes non-zero in Albanesi and Sleet's (2005) system? In the linear tax system described earlier, the government keeps track of the entire history of agents' skills. The government gets no information about the past from wealth, and so taxing wealth is only necessary to guard against socially suboptimal deviations. In Albanesi and Sleet's system, the government keeps track of the past through wealth. This mnemonic function played by wealth implies that expected marginal tax rates are no longer necessarily zero.

Albanesi and Sleet restrict their attention to situations in which skill shocks are i.i.d. over time. If skill shocks are not i.i.d., optimal allocations are no longer Markov in promised utility. Hence, it will no longer be possible to construct an optimal tax system in which taxes depend only on current income and current wealth. In this sense, their analysis is pretty special.

However, there is an important general lesson to be learned from their paper. An agent's financial position provides information about his past skill shock realizations.[6] Albanesi and Sleet's work teaches us that the properties of optimal wealth tax rates will depend crucially on how the tax system uses this information about the past encoded in wealth.

2.6.2 Asset Testing

Golosov and Tsyvinski (2004) suggest an alternative solution to the optimal tax problem which centers on *asset-testing*. To understand their solution, again consider the two-period example described in section 2.4, in which $T = 2$, $u(c) = \ln(c)$, $v(l) = l^2/2$, and $\beta = 1$. Suppose too that $\Theta = \{0, 1\}$, $Z = \{1\}$ (so that there are no aggregate shocks), $F(K, Y) = rK + wY$, and $\delta = 1$. As well, suppose $\phi_1(\theta) = 1$, $\phi_2(\theta, z) = \theta$, $v(l) = l^2/2$, and $\Pr(\theta_2 = 1) = 1/2$. Set $G = 0$. Let (c_{2i}, y_{2i}) denote consumption and effective labor when

$\theta = i$.

Suppose $(c_1^*, k_2^*, c_{2i}^*, y_{2i}^*)_{i=1,2}$ is a socially optimal allocation. Consider a tax system τ such that if an agent has capital holdings k_2 at the beginning of period 2 and generates effective labor y_2 in period 2, then he pays period 2 taxes $\tau(k_2, y_2)$, where:

$$\begin{aligned} \tau(k_2, y_2) &= \alpha_0 \text{ if } y_2 = 0 \text{ AND } k_2 \leq k_2^* \\ &= \alpha_1 \text{ otherwise} \end{aligned}$$

where:

$$\alpha_i = wy_{2i}^* + k_2^* r - c_{2i}^* \text{ for } i = 0, 1$$

Under this tax system, the agent receives a disability insurance payment from the government only if his assets are sufficiently low. In this sense, the disability insurance system exhibits asset-testing.

Under this tax system, the agent's decision problem is:

$$\max_{c_1, y_1, c_{21}, c_{20}, y_{21}, k_2} \ln(c_1) - y_1^2/2 + \ln(c_{21})/2 + \ln(c_{20})/2 - y_{21}^2/4$$

$$s.t. \; c_1 + k_2 = rk_1 + wy_1$$

$$c_{21} = rk_2 + wy_{21} - \tau(k, y_{21})$$

$$c_{20} = rk_2 - \tau(k, 0)$$

$$k_2, c_{21}, c_{20}, y_{21}, y_1 \geq 0$$

The tax on labor income is lump-sum if $y_{21} > 0$; hence, the agent finds it optimal to set y_{21} equal to y_{21}^* or 0. If the agent chooses $k_2 = k_2^* - \Delta$, $\Delta > 0$, and sets $y_{21} = y_{21}^*$, then the agent's marginal benefit from saving more is:

$$0.5\beta r/(c_{21}^* - r\Delta) + 0.5\beta r/(c_{20}^* - r\Delta) - 1/(c_1^* + \Delta)$$

which is positive from Corollary 1. So the agent wants to save more. This argument applies even more strongly if $y_{21} = 0$, because the agent's second period consumption is even lower in that case.

Hence, the agent chooses $k_2 \geq k_2^*$. Suppose then that the agent chooses $(c_1', y_1', k_2', (c_{2i}', y_{2i}')_{i=1,2})$, where $k_2' > k_2^*$. Under this choice, the agent makes a positive tax payment to the government regardless of his skill realization in period 2. It follows that the primed allocation is incentive-feasible and uses fewer resources than the original allocation; it must provide less utility to the agent.

Thus, this tax system with asset-testing implements the socially optimal allocation. Golosov and Tsyvinski (2004) show how a similar system can be used in multi-period settings, given that skill shocks follow a two-state Markov chain in which the low realization is an absorbing state. In their system, the asset cutoffs are a function of age only. My own guess is that asset-testing can be used much more generally, as long as the asset cutoffs are allowed to depend on the full history of past income realizations.

2.6.3 Non-Differentiable Taxes

In all of the systems described so far, taxes are nonseparable in wealth and income. But this nonseparability is not an essential feature of tax systems, if one allows the tax system to have kinks as a function of wealth. To see this, return to the two-period example described in the subsection 2.4. In this setting, suppose $(c_1^*, k_2^*, c_{2i}^*, y_{2i}^*)_{i=1,2}$ is a socially optimal allocation. Consider a tax system τ such that if an agent has capital holdings k_2 at the beginning of period 2, and effective labor y_2 in period 2, then in period 2, he pays taxes

$\tau^k(k_2) + \tau^y(y_2)$. Define wealth taxes by:

$$\begin{aligned}\tau^k(k) &= rk \text{ if } k > k_2^* \\ &= rk_2^* \text{ if } k \leq k_2^*\end{aligned}$$

and define labor taxes by:

$$\begin{aligned}\tau^y(y) &= \alpha_1 \text{ if } y > 0 \\ &= \alpha_0 \text{ if } y = 0\end{aligned}$$

where:

$$\alpha_i = wy_{2i} - c_{2i}^* \text{ for } i = 0, 1$$

I claim that given this tax system, agents find it individually optimal to choose the socially allocation. Given this tax system, a typical agent solves the decision problem:

$$\max_{c_1, y_1, c_{21}, c_{20}, y_{21}, k_2} \ln(c_1) - y_1^2/2 + \ln(c_{21})/2 + \ln(c_{20})/2 - y_{21}^2/4$$

$$s.t.\ c_1 + k_2 = rk_1 + wy_1$$

$$c_{21} = rk_2 + wy_{21} - \tau^k(k_2) - \tau^y(y_{21})$$

$$c_{20} = rk_2 - \tau^k(k_2) - \tau^y(0)$$

$$k_2, c_{21}, c_{20}, y_{21}, y_1 \geq 0$$

Clearly, agents choose $k_2 \leq k_2^*$; choosing to save more than that is tantamount to discarding wealth. Suppose an agent chooses $k_2 = k_2^* - \Delta$, where $\Delta > 0$. The agent's marginal benefit from saving more is:

$$
\begin{aligned}
& \beta r[0.5/(c_{21}^{*}-r\Delta)+0.5/(c_{20}^{*}-r\Delta)]-1/(c_{1}^{*}+\Delta) \\
> \; & \beta r[0.5/c_{21}^{*}+0.5/c_{20}^{*}]-1/c_{1}^{*} \\
> \; & 0
\end{aligned}
$$

Hence, the agent finds it optimal to choose $k_2 = k_2^*$. Given this choice of capital, it is easy to prove that the agent's incentive constraint implies that it is optimal to choose $y_{21} = y_{21}^*$.

Thus, this tax system, which is non-differentiable in wealth, implements the optimal allocation. Unlike the linear tax system described earlier or Albanesi and Sleet's (2005) optimal system, the tax system is separable between wealth and labor. While the above analysis takes place in a two-period context, the result is much more general: if taxes are allowed to be non-differentiable, then the tax rate on wealth need not depend on the ex-post realization of labor income.

2.6.4 General Lessons?

There are many forms of optimal wealth taxes. This finding is, as I indicated earlier, nothing new in optimal tax theory. Even under the Ramsey approach, in which taxes are restricted to be linear, there is a great deal of indeterminacy. The situation is similar to that in portfolio theory. There, if some assets have perfectly correlated payoffs, then the optimal portfolio is indeterminate. Similarly, if the government has a lot of flexibility in terms of tax instruments, there are many ways to achieve the socially optimal outcome. Indeed, in some ways, we have greatly reduced the potential for indeterminacy by assuming that the government is solely responsible for social insurance. This indeterminacy only becomes more pronounced once one takes into account the possibility that the private sector might also provide social insurance.

I do believe that there are some reasons why the linear tax system is preferable to other forms of taxation. The asset cutoff system or the kinks system are not robust to small errors in the taxation authority's ability to measure a household's wealth. (The Albanesi-Sleet (2005) system does not have this problem, because taxes are differentiable with respect to wealth. However, it is very much tailored to the unrealistic case in which skill shocks are i.i.d. over time.) The linear system that I described relies on careful measurement of labor incomes, but is more robust to problems in the measurement of wealth.

It would be desirable to provide necessary conditions that all optimal tax systems must satisfy. However, the above discussion indicates that this task is not an easy one. The necessary conditions end up depending precisely on auxiliary functional form restrictions. For example, if one restricts attention to systems that are differentiable in wealth, then the discussion in section 2.4 shows that the marginal tax rate on wealth must depend on realized income in an optimal system. The discussion in the previous section shows that this implication disappears once one allows the tax system to be non-differentiable.

This kind of indeterminacy is taken by some economists to mean that optimal tax theory has no useful implications. This is not true. For example, at a minimum, any optimal tax system must be consistent with the Euler equation described in Proposition 1. In ongoing research, Luigi Pistaferri and I are examining to what extent this Euler equation holds in United States household consumption data and using our results to suggest ways in which the current wealth tax system might be improved.

3. The Mirrlees Approach and Other Kinds of Taxes

I now turn to the results of using the Mirrlees approach to study three other types of taxes: estate taxes, inflation, and labor income taxes.

3.1 Optimal Estate Taxation

In a recent paper, Farhi and Werning (2005) consider the problem of optimal estate taxation. They study a multi-period setting with i.i.d. skill shocks. However, it is easier to understand the economics of their analysis in a simpler model. Suppose there are two periods, and a unit measure of families. Each family consists of a parent and a child. The parent is born at the beginning of period 1 and dies at the end of that period; the child is born at the beginning of period 2 and dies at the end of that period. The parent has utility function given by:

$$u(c_1) - v(l_1) + \beta V_2, \beta > 0$$

where $V_2 = u(c_2) - v(l_2)$ is the utility of the child. Here, $c_t(l_t)$ is the consumption (labor) of the family member alive in period t. I assume that $u', v', -u'', v'' > 0$.

Skills are private information and a child's skill is the same as that of its parent. In this world, the parent and the child share the same ranking over consumption/labor profiles. Hence, we can write the incentive-compatibility constraint of the high-skilled as:

$$\begin{aligned} & u(c_{1H}) - v(y_{1H}/\theta_H) + \beta u(c_{2H}) - \beta v(y_{2H}/\theta_H) \\ \geq \; & u(c_{1L}) - v(y_{1L}/\theta_H) + \beta u(c_{2L}) - \beta v(y_{2L}/\theta_H) \end{aligned}$$

Here as before, y represents effective labor. (I have dropped the incentive-compatibility constraint of the low-skilled households, because it will not bind.) I assume that there are no aggregate shocks (Z is a singleton) and that $F(K, Y) = rK + wY$.

Suppose there is a planner who puts weight 1 on all parents and puts weight β' on their children. Then, we can write the planner's ex-ante objective as:

$$\mu_H[u(c_{1H}) - v(y_{1H}/\theta_H) + (\beta + \beta')\{u(c_{2H}) - v(y_{2H}/\theta_H)\}]$$

$$+\mu_L[u(c_{1L}) - v(y_{1L}/\theta_L) + (\beta + \beta')\{u(c_{2L}) - v(y_{2L}/\theta_L)\}]$$

The key element of this objective is that, as long as $\beta' > 0$, the planner puts more weight on children than do their parents.

3.1.1 Characterizing the Optimal Allocation

Using methods akin to those in Kocherlakota (2005), it can be shown that if (c^*, y^*) is socially optimal, then:

(22) $$r[\beta + \beta' \frac{u'(c^*_{1i})}{\mu_H u'(c^*_{1H}) + \mu_L u'(c^*_{1L})}] \frac{u'(c^*_{2i})}{u'(c^*_{1i})} = 1, i = h, l$$

The restriction (22) has two important implications. The first is that:

(23) $$r\beta \frac{u'(c^*_{2i})}{u'(c^*_{1i})} < 1$$

In this setting, with a dynastic transition from period 1 to period 2, it is socially optimal for the social rate of return to be lower than the parents' rate of time preference. The reason for this result is simple. Society puts more weight on children than parents do. Hence, parents need to be subsidized to provide a socially optimal amount of consumption to their children.

The second implication is more subtle. Suppose first that there is no private information. Then, it is efficient to insure households against their skill shocks by equating c^*_{1H} and c^*_{1L}. The optimality condition (22) implies that:

(24) $$r[\beta + \beta'] \frac{u'(c^*_{2i})}{u'(c^*_{1i})} = 1 \text{ for } i = h, l$$

Without private information, it is socially optimal for the planner to equate his version of the family's intertemporal marginal rate of substitution - namely, $[\beta+\beta']\frac{u'(c_{2i}^*)}{u'(c_{1i}^*)}$ – to the marginal rate of transformation $1/r$. With private information, (22) implies that:

$$(25) \quad r[\beta+\beta']\frac{u'(c_{2H}^*)}{u'(c_{1H}^*)} > 1 > r[\beta+\beta']\frac{u'(c_{2L}^*)}{u'(c_{1L}^*)}$$

The planner's version of the high-skilled (low-skilled) family's intertemporal marginal rate of substitution is larger (smaller) than the marginal rate of transformation.

The surprising feature of this optimal allocation is that the high-skilled family is distorted relative to the full information case. This result violates the general rule in private information problems that in an optimal allocation, there is no distortion at the top; the low-skilled family's incentive constraint does not bind and so the high-skilled's allocation should not be distorted. There is a simple argument for why this rule breaks down here. Suppose that:

$$(26) \quad r[\beta+\beta']\frac{u'(c_{2H}^*)}{u'(c_{1H}^*)} = 1$$

The planner contemplates the following perturbation: increase c_{1H}^* by ε and decrease c_{2H}^* by $r\varepsilon$. This perturbation lowers the planner's objective. However, it makes the high-skilled household better off, which frees up the incentive constraint. The planner can then do a second perturbation (transferring from the high-skilled household to the low-skilled household) that reduces ex-ante risk. The loss to the planner from the first perturbation is only second-order, because the planner is at a tangency. The gain to the planner from the second perturbation is first-order, because the incentive constraint was binding.[7]

3.1.2 Taxes

As we have seen, we need to be cautious in drawing conclusions about taxes from these kinds of characterizations of the socially optimal allocations. Nonetheless, any optimal system that is differentiable with respect to bequests must have two attributes.

First, all households should receive marginal estate subsidies. The planner puts more weight on children than the child's parents do. Hence, the tax system needs to encourage the parent to transfer more to its child.

Second, the size of this marginal subsidy should vary across households. In particular, the estate of a parent who is highly-skilled should be subsidized at a lower marginal rate than that of a parent who is low-skilled. This prescription may sound egalitarian. However, it emerges purely from incentive constraints. In particular, if the planner for some reason put *less* weight on children than parents did, then the result would be exactly reversed: The optimal subsidy system would subsidize high-skilled parents at a higher rate than low-skilled parents.

Like Farhi and Werning (2005), I have couched the above discussion in terms of estate taxes. In fact, these tax results apply to *all* transfers from parents to children. Thus, the results provide indirect support for subsidies to child education and other kinds of inter-vivo transfers. Again, the marginal subsidy rate should be larger for low-skilled parents than high-skilled parents.[8]

3.2 Inflation

In a recent working paper, da Costa and Werning (2005) study the behavior of optimal inflation in an economy in which agents receive utility from consumption, leisure, and real balances. They assume that agents have fixed skills that are private information and that the planner maximizes a utilitarian objective. They also assume that the planner cannot observe an agent's holdings of any assets including money.

In this setting, a positive nominal interest rate is a way to tax real balances relative to other forms of consumption. da Costa and Werning assume that the agents' utility function is such that high-skilled agents demand no more real balances than does a low-skilled agent (given that both types of agent have the same allocation of effective labor and consumption). Under this assumption, a positive nominal interest rate is a tax on a good that the low-skilled likes better than the high-skilled. Such a tax tightens the incentive constraint by making it more attractive for the high-skilled to mimic the low-skilled, given a certain level of reservation utility for the low-skilled. Positive nominal interest rates are suboptimal. In other words, it is optimal for the government to follow the Friedman Rule.

Putting real balances in the utility function is a reduced-form way to capture money demand. Da Costa and Werning consider two deeper models of money demand. One is a cash-credit model in which agents are required to buy some consumption goods with cash. The other is a shopping-time model in which agents can economize on the time spent on shopping using real-balances. They show that in both models, the high-skilled demand no more real balances than the low-skilled. Thus, in both models, the Friedman Rule is optimal.

There is a large literature describing optimal monetary policy using a Ramsey approach (see Chari and Kehoe (1999) for a discussion). The Friedman Rule also emerges as a prime policy recommendation from this literature. What do we learn from the Mirrlees approach over the Ramsey approach? Researchers using the Ramsey approach have emphasized that the Friedman Rule may be suboptimal if the income elasticity of money demand is sufficiently large. This elasticity, which is somewhat delicate to estimate precisely, is irrelevant under the Mirrlees approach.

The second lesson is probably more important. da Costa and Werning (2005) assess the welfare losses associated with using a suboptimal monetary policy. They show that these losses are considerably higher in a Mirrleesian economy, in which agents are heterogeneous in skills, than in a representative agent economy with the same aggregate money demand function. Positive nominal interest rates make it more costly for the government to provide incentives to the high-skilled. The representative agent economy ignores this effect completely.

3.3 Labor Income Taxes

As of this writing, Battaglini and Coate (2005) have written the only paper that attempts to address the question of the dynamic behavior of optimal labor income taxes when skill shocks are only imperfectly persistent. They consider a setting in which agents are risk-neutral over consumption and face skill shocks that follow a positively autocorrelated Markov chain with two possible realizations. This environment is highly specialized. However, it has the benefit that Battaglini and Coate can solve explicitly for optimal allocations.

Up until now, we have focused on ex-ante Pareto optima. These allocations are not of interest to Battaglini and Coate. With risk-neutral consumers, there is no trade-off between insurance and incentives. Hence, the ex-ante Pareto optimal tax systems would be lump-sum.

Instead, Battaglini and Coate study allocations that are Pareto optimal, conditional on the realization of period 1 skills. Of course, one of these allocations corresponds to the ex-ante Pareto optimum and is again uninteresting to Battaglini and Coate. Battaglini and Coate focus on the set of interim Pareto optima such that which the planner puts sufficiently high weight on the initially low-skilled agents that the incentive constraint of the initially highly-skilled agents binds.

In this world, if skills are fixed over time, then the labor supply of the highly-skilled is not distorted (no-distortion-at-the-top principle) and the labor supply of the low-skilled is distorted downwards. Battaglini and Coate's main result is for the case in which skills are persistent but not perfectly so. Let the support of the skill shocks be $\{\theta_H, \theta_L\}$ where $\theta_H > \theta_L$ and the agents' disutility of labor be given by the function v. Consider an agent who has skill history $\overline{\theta}^t$ such that $\overline{\theta}_s = \theta_H$ for some $s \leq t$. Then in any optimal allocation (c, y):

$$(27) \quad v'(y_t(\overline{\theta}^t)/\overline{\theta}_t) = \overline{\theta}_t$$

After an agent is highly-skilled, it is never optimal to distort his labor supply, even if he later becomes low-skilled.

The basic idea of the proof of this result is simple. Assume that only high-skilled agents can misrepresent themselves. (Battaglini and Coate prove that the incentive constraints of the low-skilled never bind, so this assumption is without loss of generality). Suppose:

$$(28) \quad v'(y_t(\overline{\theta}^t)/\overline{\theta}_t) < \overline{\theta}_t$$

so that the agent's labor supply is distorted downwards at time t. Let $s = \max\{r | \overline{\theta}_r = \theta_H\}$. Then, consider a perturbation to the allocation (c, y) in which we simultaneously increase $y_t(\overline{\theta}^t)$ and increase $c_s(\overline{\theta}^s)$ so as to keep the agent's lifetime utility unchanged in history θ^s. Because the agent's labor supply was distorted downward, this perturbation provides more resources for the planner. After period s, the perturbation increases the loss from pretending to be low-skilled. In period s and earlier periods, because the shocks are Markov, the incentive to mis-report has been left unchanged Hence, this perturbation is incentive-compatible and provides the same welfare to agents in the economy at lower cost.

The crux of this perturbation is that we have to roll the gains of increasing labor supply backwards in time to a high type. If we cannot do that, the perturbation will not work. Hence, in general, the labor supply of any agent who has always been low-skilled is distorted downwards. (Battaglini and Coate prove that the size of this distortion is declining to zero over time.)

Battaglini and Coate show that their sharp characterization of the Pareto optimal allocations does not survive if preferences are allowed to exhibit risk aversion. (Fortunately, at least for finite horizons, the optimal allocations are continuous with respect to risk aversion. Thus, for sufficiently low levels of risk aversion, the distortions are small for agents who have ever been highly-skilled.) The structure of optimal labor income taxes in a setting with persistent skill shocks and risk aversion remains an important open question.[9]

4. Where Do We Go From Here?

In this last section, I discuss some possible future directions for the new dynamic optimal taxation literature.

4.1 Hidden Asset Income

In Section 2, we discussed a number of optimal tax systems. All of these tax systems assume that the government is responsible for social insurance. However, Golosov and Tsyvinski (2005) point out that it is possible to implement Pareto optimal allocations without any government intervention at all (except through the enforcement of long-term contracts). In these implementations, a large number of private insurers compete before period 1 to sign up agents to contracts that last for the duration of the economy. These contracts specify consumption and labor supply for all agents, as a function of reports that they send to their insurer over their lifetimes.

Golosov and Tsyvinski then change the environment by positing that the agents are able to engage in secret trades of a risk-free asset without being detected by the taxation authority or other insurers. These secret trades allow agents to engage in wealth tax arbitrages, and so all agents must face uniform wealth taxes. (The taxes need not be zero; agents are able to secretly borrow and lend with one another, but not secretly accumulate capital.)

Golosov and Tsyvinski show that it is generally optimal for these wealth taxes to be non-zero. They also show that if skills are i.i.d. over time, then the wealth taxes are optimally positive. Intuitively, in the i.i.d. case, agents want to deviate from socially optimal allocations by saving secretly and then shirking. By taxing the return on wealth, this deviation becomes less attractive. Golosov and Tsyvinski argue that only a government has the ability to affect equilibrium prices in this fashion. They conclude that hidden intertemporal side-trades provide a novel rationalization for government intervention in social insurance.

An important difficulty with Golosov and Tsyvinski's analysis is that borrowing and lending require *enforcement*: somebody must compel the borrower to repay their obligation. What kinds of courts can enforce the *secret* intertemporal side-trades in Golosov and Tsyvinski's setup? Golosov and Tsyvinski make reference to informal enforcement mechanisms, such as those that exist within families. However, it is unnatural to model such informal side-contracts through anonymous trades of a risk-free asset. Side-trades within families or other informal coalitions are generally neither anonymous nor non-contingent. It is important for future work to verify whether Golosov and Tsyinvski's conclusions about government intervention and wealth taxes survive a more careful modeling of how exactly secret intertemporal side-trades are being enforced.

As I have done above, Golosov and Tsyvinski assume that capital holdings are fully observable. It would be useful to understand the structure of optimal taxes given weaker assumptions about capital holdings. For example, it may be necessary for the government to pay an auditing cost to ascertain an individual's holdings of physical capital. Similarly, it may be costly for the government to distinguish whether a given flow of income is derived from physical capital or labor.

4.2 Limited Commitment

In deriving the optimal tax systems in sections 2 and 3, we implicitly assumed that the government could fully commit to a tax system at date 0. In the real world, this kind of commitment is difficult to achieve. What happens to the optimal tax systems when the government is allowed to renege on its past promises?

The nature of the time-consistency problem in the Mirrleesian setting is quite different from the usual time inconsistencies that occur in a world with linear taxes. With linear taxes, the government is always tempted to tax current period capital income at a high rate, because such a tax is virtually lump-sum. In the Mirrleesian setup, lump-sum taxes are already available to the government. This temptation does not exist.

Instead, the relevant temptation is that the government may want to exploit information that it has learned about agents in earlier periods. The simplest example of this phenomenon is when skills can take on two possible values and are fixed over time. In this case, the optimal tax system with commitment is constant over time and involves distortion of the labor supply of the low-skilled. However, at the beginning of period 2, the government knows who is low-skilled. The government then wants to change the future tax code so that it becomes lump-sum.

There have been several attempts to analyze the optimal (equilibrium?) choices of taxes once the government is no longer able to commit (see Berliant and Ledyard (2004), Sleet and Yeltekin (2005), and Roberts (1984)). However, these analyses have focused on rather simple examples (with either i.i.d. shocks or two periods). Much more needs to be done.

Bisin and Rampini (2005)'s recent article is especially provocative. In Golosov and Tsyvinski (2005) and the prior related literature, hidden side-trades are a nuisance that prevent the government and society from achieving desirable outcomes. Bisin and Rampini (2005) consider a series of examples in which the government is not able to commit. They show that in these settings, it may be *optimal* for society to allow agents to engage in hidden side-trades. The hidden side-trades prevent the government from fully exploiting its ex-post informational advantage.

This line of argument needs to be explored much more fully. Our standard rationalization for markets is the First Welfare Theorem. But it merely says that markets are just as good as benevolent governments. Bisin and Rampini provide a novel rationalization for markets: they are in fact necessary in order to discipline the ex-post opportunism of benevolent governments.

4.3 What Are the Right Frictions?

I have often heard the criticism that the Mirrleesian approach is inferior to the Ramsey approach because the optimal Mirrleesian tax systems do not resemble those used in the real world as well as the optimal Ramsey systems do. In particular, it is often said that the Mirrlees approach leads to tax systems that are overly complicated relative to those used in practice, while the Ramsey approach does not.

This criticism strikes me as misguided in two respects. The first problem with it is that the optimal tax systems that emerge from a Mirrleesian analysis are not at all complicated compared to the hundreds of pages necessary to describe actual tax codes. Both the Mirrleesian and Ramsey tax systems are unrealistically simple once one takes this perspective.

But there is a second more fundamental problem with this criticism. Regardless of which approach one uses, optimal taxation is a normative, not a positive, exercise. One does not judge a normative analysis by its prescriptions. The whole point of such analyses is that the government is using a flawed policy and the analyst is attempting to discover better policies. These better policies may differ greatly from conventional practice. Thus, many economists have argued for many years in favor of rarely-seen policies like setting farm subsidies or trade barriers to zero.

This does not mean that normative analyses are conducted independently of empirical considerations. One judges a normative analysis by the realism of its assumptions. In this regard, an attractive feature of the Mirrlees approach is that it is more disciplined than the Ramsey approach, Under the Ramsey approach, the set of tax instruments is entirely at the disposal of the modeler. Under the Mirrlees approach, the optimal tax instruments are an endogenous response to exogenously specified private information frictions.

Of course, if this discipline is to be useful, we need to have a way to figure out which frictions are really operational in the world. Right now, the literature has been proceeding by individual researchers thinking up new frictions (or really old ones!) and plugging them into the existing framework. We need to have a much more systematic way to proceed.

In principle, this is not an impossible task. Informational frictions are a specification of a particular type of technology. For example, when we say "effort is hidden", we are really saying that it is infinitely costly for society to monitor effort. The desired approach would be to devise optimal tax systems for different specifications of the costs of monitoring different activities and/or individual attributes.

To be able to implement this approach, we need to accomplish two goals. One is to extend our modes of technical analysis to allow for costs of monitoring other than zero or infinity. The other is to figure out empirical methods of measuring these costs. Neither of these goals will be easy to achieve. But both are essential if the dynamic Mirrlees approach is to survive as more than a collection of theoretical results.

University of Minnesota, Federal Reserve Bank of Minneapolis, and NBER

5. Notes

1. The initial draft of this essay was completed while I was a professor at Stanford University. I thank Ellen McGrattan and Torsten Persson for their comments. I acknowledge the support of NSF SES-0350833. The views expressed herein are mine and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

2. I believe that the analysis could be extended to allow Z to be an infinite set, but I have not done so.

3. Following Wilson (1987), Bergemann and Morris (2005) are rightly critical of making these kinds of strong common knowledge assumptions in mechanism design. They urge the use of what they call *robust* mechanism design, which allows agents to hold differing and unknown beliefs. It would be useful to use robust mechanism design to analyze optimal taxes.

4. The notion of social optimality (ex-ante symmetric Pareto optima) may seem restrictive. However, the results about wealth taxes are readily generalized to other notions of optimality: ex-ante asymmetric Pareto optima, or interim Pareto optima (conditional on period 1 realizations of θ).

5. In deriving his results about the structure of the optimal contract in a setting with repeated moral hazard, Rogerson (1985) does not require single-crossing. In what follows, the results about wealth taxes rely on a line of attack similar to that adopted by Rogerson.

6. Kapicka (2005) studies a dynamic Mirrlees tax system in which the government can only condition taxes on current incomes. This restriction seems unduly severe given the potential mnemonic role of financial wealth.

7. More generally, as long as the planner's objective deviates from those of the individual decision-maker (because of externalities), then the no-distortion-at-the-top principle

will break down. See Amador, Angeletos, and Werning (2005).

8. Farhi and Werning (2005) focus on a multi-period environment in which skill shocks are i.i.d over time. I have described how their estate taxes work in a two-period setting with fixed skills. However, it is possible to show that these restrictions are not essential: analogs of the results can be derived for any skill shock process and for any horizon length. Then, the optimal policy is one of subsidizing all estates, with a greater subsidy for low-wealth parents.

9. Zhang (2005) provides a preliminary analysis based on a novel continuous-time solution method.

REFERENCES

[1] ALBANESI, S., and SLEET, C. (2005): "Dynamic Optimal Taxation with Private Information," forthcoming, *Review of Economic Studies.*

[2] AMADOR, M., ANGELETOS, G.-M., AND WERNING, I. (2005): "Commitment versus Flexibility," forthcoming, *Econometrica.*

[3] BATTAGLINI, M., and COATE, S. (2005): "Pareto Efficient Income Taxation with Stochastic Abilities," Princeton University working paper.

[4] BERGEMANN, D., and MORRIS, S. (2005): "Robust Mechanism Design," *Econometrica* 73, 1771-1813.

[5] BERLIANT, M. and LEDYARD, J. (2004): "Dynamic Optimal Nonlinear Income Taxation with No Commitment," Washington University working paper.

[6] BISIN, A., and RAMPINI, A. (2005): "Markets as Beneficial Constraints on Governments, " forthcoming, *Journal of Public Economics.*

[7] CHARI, V. V., and KEHOE, P. (1999): "Optimal Fiscal and Monetary Policy," in *Handbook of Macroeconomics,* ed. Taylor, J., and Woodford, M., (New York, NY: Elsevier).

[8] DA COSTA, C., and WERNING, I. (2005): "On the Optimality of the Friedman Rule with Heterogeneous Agents and Nonlinear Income Taxation," MIT working paper.

[9] FARHI, E., and WERNING, I. (2005): "Inequality, Social Discounting, and Estate Taxation," MIT working paper.

[10] GOLOSOV, M., KOCHERLAKOTA, N., and TSYVINSKI, A. (2003): "Optimal Indirect and Capital Taxation," *Review of Economic Studies* 70, 569-587.

[11] GOLOSOV, M., and TSYVINSKI, A. (2004): "Designing Optimal Disability Insurance," MIT working paper.

[12] GOLOSOV, M., and TSYVINSKI, A. (2005): "Optimal Taxation with Endogenous Insurance Markets," MIT working paper.

[13] GREEN, E. (1987): "Lending and the Smoothing of Uninsurable Income," in Prescott, E., and Wallace, N., eds., *Contractual Arrangements in Intertemporal Trade*, Minneapolis, MN: University of Minnesota Press.

[14] KAPICKA, M. (2005): "Optimal Taxation with Hidden Human Capital Accumulation," University of California-Santa Barbara working paper.

[15] KOCHERLAKOTA, N. (2004): "Wedges and Taxes," *American Economic Review* 94, 109-113.

[16] KOCHERLAKOTA, N. (2005): "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation," *Econometrica,*.73, 1587-1621.

[17] MIRRLEES, J. (1971): "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies,* 38, 175-208.

[18] ROBERTS, K. (1984): "The Theoretical Limits of Redistribution," *Review of Economic Studies* 51, 177-95.

[19] ROGERSON, W. (1985): "Repeated Moral Hazard," *Econometrica,* 53, 69-76.

[20] SLEET, C. and YELTEKIN, S. (2004): "Credible Social Insurance, " Carnegie-Mellon working paper.

[21] WILSON, R. (1987): "Game-Theoretic Analyses of Trading Processes," in *Advances in Economic Theory: Fifth World Congress,* ed. T. Bewley, Cambridge, UK: Cambridge University Press, Chap. 2, 33-70.

[22] ZHANG, Y. (2005): "Dynamic Contracts, Persistent Shocks, and Optimal Taxation," Federal Reserve Bank of Minneapolis working paper.