ZERO EXPECTED WEALTH TAXES: A MIRRLEES APPROACH TO DYNAMIC OPTIMAL TAXATION

By Narayana R. Kocherlakota[1]

In this paper, I consider a dynamic economy in which a government needs to finance a stochastic process of purchases. The agents in the economy are privately informed about their skills, which evolve stochastically over time; I impose no restriction on the stochastic evolution of skills. I construct a tax system that implements a symmetric constrained Pareto optimal allocation. The tax system is constrained to be linear in an agent's wealth, but can be arbitrarily nonlinear in his current and past labor incomes. I find that wealth taxes in a given period depend on the individual's labor income in that period and previous ones. However, in any period, the expectation of an agent's wealth tax rate in the following period is zero. As well, the government never collects any net revenue from wealth taxes.

Keywords: Optimal taxation, capital, private information.

People are born with different abilities to turn effort into output—that is, with different *skills*. These skills evolve *stochastically*. Talented people may awake one day with crippling back pain or chronic fatigue syndrome that renders them low skilled. Unemployed people may suddenly find a good job opportunity. Some people learn faster or forget slower than others. It is plausible to think of skills at birth as private information (as in Mirrlees (1971)), but these kinds of *changes* in skills are also difficult for outside observers to verify directly.

Motivated by these observations, I ask the following question. Suppose people are privately informed about their skills and those skills evolve stochastically over time. What are the properties of optimal wealth taxes, given that the government wants to use the tax system to fund its purchases and to insure people against skill shocks?

To answer this question, I restrict attention to preferences that are additively separable between consumption and leisure. I place *no* restriction on the process that generates skills or on the process that generates aggregate shocks. I prove that a symmetric constrained Pareto optimal allocation of resources

can be implemented using a tax system that is *linear* in current wealth. The wealth tax rate in any period is a function of current and past labor incomes; however, the expected period $(t+1)$ wealth tax rate for a person, conditional on his period t information, is zero. I then show that total wealth tax collections are zero in every date and state, regardless of the realization of the aggregate shocks.

In the aggregate, wealth taxes are zero, but it is crucial that wealth taxes not be zero for every person. I find that in the optimal system, a given agent's period $(t+1)$ wealth tax is a strictly decreasing function of consumption in period $(t+1)$. Thus, people who are surprisingly low skilled in period $(t+1)$ pay a wealth tax; people who are surprisingly highly skilled receive a wealth subsidy. The intuition behind these nonzero taxes is simple. The optimal tax system is designed to trade off social insurance of skill shocks against incentives to provide effort. The "regressive" wealth taxes deter agents from bringing a large amount of wealth into period $(t+1)$ and then reducing work effort. This dynamic deterrence allows the society to have a tax system that provides better social insurance for a given level of output.

Given my analysis, actual tax systems are suboptimal in two ways. First, the average tax rate on net capital income is around 50% across OECD countries (Carey and Rabesona (2002)). Second, in real-world tax codes, there is little explicit dependence of wealth taxes on labor income shocks. It is known from prior work that the first kind of suboptimality can lead to large welfare losses (Lucas (2003)). I show through a numerical example that the lack of labor income dependence can potentially lead to even greater welfare losses.

My paper is a dynamic extension of Mirrlees' (1971) analysis. The question of optimal wealth taxes in dynamic versions of the Mirrlees model has also been taken up in recent papers by Albanesi and Sleet (2005) and Golosov and Tsyvinski (2004). The limitation of these papers is that they restrict attention to particular shock processes for individual skills (and, less importantly, assume that there are no aggregate shocks). Golosov and Tsyvinski (2004) focus on disability insurance and so they assume that skills have an absorbing state. Albanesi and Sleet (2005) assume individual skills are independently and identically distributed over time.[2]

Neither of these assumptions is a particularly good approximation to what we know about individual skills from the empirical literature on individual wages. This literature documents that individuals experience large and persistent shocks to their wages (and presumably to their skills as well) throughout their lives (see, among others, Storesletten, Telmer, and Yaron (2001) and Meghir and Pistaferri (2004)). Thus, my primary contribution to the existing literature on dynamic optimal taxation is that my results are applicable

---

[2]Golosov and Tsyvinski construct a system in which agents face a age-dependent asset test to receive welfare benefits. In Albanesi and Sleet's system, the optimal taxes are a function only of current wealth and current labor income (although this function is allowed to be arbitrarily nonlinear).

for *empirically relevant* specifications of the stochastic process that generates skills. I document in the context of a numerical example how persistence in individual-level skill shocks and the presence of aggregate shocks affect optimal wealth taxes.

Golosov, Kocherlakota, and Tsyvinski (2003) (GKT) provide a partial intertemporal characterization of socially optimal allocations in a class of environments similar to mine. I generalize their characterization to allow for aggregate shocks. This generalized characterization plays a key role in my proof that optimal wealth taxes should be zero on average. However, my main contribution over their work is that, unlike them, I explicitly consider a decentralized tax system. In particular, GKT find that it is optimal for a wedge to exist between the intertemporal marginal rate of transformation and individuals' intertemporal marginal rates of substitution. I show through an example that this optimal *wedge* does not translate directly into a positive *tax* on capital in a decentralized tax system.[3] Instead, the optimal tax system creates the wedge in a subtle fashion: Individuals face idiosyncratic wealth tax risk that serves to deter saving.

Chari, Christiano, and Kehoe (CCK) (1994) ask the question, "What is the structure of optimal capital income taxes when government purchases are stochastic?" However, they use what might be termed the Ramsey approach: they assume that there is no heterogeneity across individuals and assume that the government can only use linear taxes on labor and capital income. Assuming preferences are additively separable between consumption and leisure (as I do in this paper), they find that the optimal tax rate on capital income is zero in any period $t > 2$ if preferences are homothetic in consumption. In contrast, I find that it is optimal for aggregate capital tax collections to be zero in every period, including periods 1 and 2.

I discuss these differences in results in the body of the paper. However, the difference is really hardly surprising, because entirely different forces are at work in the two kinds of analyses. In the Ramsey approach used by CCK, the goal is to minimize the deadweight loss associated with the distortions generated by the linearity of taxes. (Sufficiently nonlinear taxes are nondistorting, because they are lump-sum.) Under the Mirrlees approach that I use, the goal is to design taxes so as to minimize the deadweight loss associated with providing good incentives.

1. ENVIRONMENT

In this section, I describe the environment. The description is similar to that in GKT, except that I allow for the possibility of publicly observable aggregate shocks.

---

[3]The example closely follows similar ones in Albanesi and Sleet (2005) and Golosov and Tsyvinski (2004).

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

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

where $c_t \in \mathbb{R}_+$ is the agent's consumption in period t and $l_t \in \mathbb{R}_+$ 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 work as follows. Let Z be a finite set 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 $\mathbb{R}_+$ 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 $\mu_\Theta(B)$.

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 $\theta^t = (\theta_1, \ldots, \theta_t)$ and the history of public shocks $z^t = (z_1, \ldots, z_t)$. 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

$$(2) \qquad y_t(\theta^T, z^T) = \phi_t(\theta^T, z^T) l_t(\theta^T, z^T),$$

$$(3) \qquad \phi_t : \Theta^T \times Z^T \to (0, \infty),$$

$$(4) \qquad \phi_t \text{ is } (\theta^t, z^t)\text{-measurable.}$$

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

$$K : Z^T \to \mathbb{R}_+^{T+1},$$

$$c : \Theta^T \times Z^T \to \mathbb{R}_+^T,$$

$$y : \Theta^T \times Z^T \to \mathbb{R}_+^T,$$

K_{t+1} is z^t-measurable,

(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 $G_t : Z^T \to \mathbb{R}_+$ 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,

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

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

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

$$K_1 \le K_1^*.$$

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 $F_t : \mathbb{R}_+^2 \times Z^T \to \mathbb{R}_+$ is assumed to be strictly increasing, weakly concave, homogeneous of degree 1, continuously differentiable with respect to its first two arguments, and z^t-measurable with respect to its last argument.

Both ϕ_t and F_t are allowed to depend on the history of shocks in potentially complicated nonlinear ways. In particular, in keeping with recent empirical descriptions of idiosyncratic shocks to wages (Storesletten, Telmer, and Yaron (2001)), the variance of ϕ_t, conditional on θ^t and z^{t+1}, may well be a nondegenerate function of z^{t+1}.

Because θ_t is only privately observable, allocations must respect incentive-compatibility conditions. (The following definitions correspond closely to those in GKT.) A *reporting strategy* $\sigma : \Theta^T \times Z^T \to \Theta^T \times Z^T$, 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

$$W(\cdot; c, y) : \Sigma \to \mathbb{R},$$

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

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 $\sigma_{\mathrm{TT}}(\theta^T, z^T) = (\theta^T, z^T)$ for all θ^T, z^T. Then an allocation (c, y, K) is *incentive-compatible* if

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

An allocation that 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

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

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

2. AN INTERTEMPORAL CHARACTERIZATION OF OPTIMAL CONSUMPTION ALLOCATIONS

In this section, I provide a partial characterization of optimal allocations that is valid for *any* specification of the exogenous elements of the model $(\phi, F, \mu_\Theta, \mu_Z, u, v, \beta, Z, \Theta)$. The main contribution is that I extend GKT's intertemporal characterization into this setting with aggregate shocks.

The key proposition is the following. It establishes that any optimal allocation must satisfy a particular first order condition (similar to that derived in Theorem 1 of GKT (2003) and in Rogerson (1985)).

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

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

(6) $\lambda_{t+1}^* = \beta \dfrac{[E(u'(c_{t+1}^*)^{-1}|\theta^t, z^{t+1})]^{-1}}{u'(c_t^*)}$ *a.e.,*

(7) $E\{\lambda_{t+1}^*(1 - \delta + F_{K,t+1}^*)|z^t\} = 1$ *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}.

The proof is given in Appendix A.

The content of this proposition is twofold. First, it establishes that

$$\beta \frac{\{E(u'(c_{t+1}^{*})^{-1}|\theta^{t}, z^{t+1})\}^{-1}}{u'(c_{t}^{*})}$$

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} so as to require highly 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}.

Second, the theorem establishes that this conditional harmonic mean is equal to the social discount factor (λ) between period t and period $(t+1)$. The social discount factor can then be used to determine the optimal level of capital accumulation between period t and period $(t+1)$.

Why does the relationship involve harmonic means, as opposed to arithmetic means? Assume Θ is finite and think about the marginal benefit to the planner of getting ε extra units of per capita consumption in history z^{t}. At first glance, one might think that the marginal benefit is proportional to the arithmetic mean of marginal utilities:

$$\varepsilon \sum_{\theta^{t} \in \Theta^{t}} \mu_{\Theta}(\theta^{t}) u'(c_{t}(\theta^{t}, z^{t})).$$

(For the purposes of this intuitive argument, I write c_{t} as a function of (θ^{t}, z^{t}), not (θ^{T}, z^{T}). This is without loss of generality, because c_{t} is (θ^{t}, z^{t})-measurable.) However, this implicitly assumes that each agent is receiving ε units of consumption regardless of history, which will typically violate incentive constraints.

Instead, the extra resources should be split so that each agent θ^{t} receives $\eta(\theta^{t})$, where $\sum_{\theta^{t} \in \Theta^{t}} \eta(\theta^{t}) \mu_{\Theta}(\theta^{t}) = \varepsilon$ and, for all $\theta^{t}, \theta^{t\prime}$,

$$u(c_{t}(\theta^{t}, z^{t}) + \eta(\theta^{t})) - u(c_{t}(\theta^{t\prime}, z^{t}) + \eta(\theta^{t\prime})) = 0$$

or, using a first order approximation,

$$u'(c_{t}(\theta^{t}, z^{t}))\eta(\theta^{t}) = u'(c_{t}(\theta^{t\prime}, z^{t}))\eta(\theta^{t\prime}) = B$$

for some B. We can solve for B using

$$\varepsilon = \sum_{\theta^{t} \in \Theta^{t}} \frac{B\mu_{\Theta}(\theta^{t})}{u'(c_{t}(\theta^{t}, z^{t}))},$$

so that the marginal gain to the planner is given by

$$\sum_{\theta^t \in \Theta^t} \mu_\Theta(\theta^t) u'(c_t(\theta^t, z^t)) \eta(\theta^t) = B$$

$$= \varepsilon \left[\sum_{\theta^t \in \Theta^t} \frac{\mu_\Theta(\theta^t)}{u'(c_t(\theta^t, z^t))} \right]^{-1}.$$

The shadow value of resources in a history z^t is given by the harmonic mean of marginal utilities, not the arithmetic mean.[4]

Proposition 1 immediately implies that there is an intertemporal wedge of the sort established by GKT. By using Jensen's inequality, we get

$$(8) \qquad \beta E\{u'(c^*_{t+1})(1 - \delta + F_{K,t+1}) | z^t, \theta^t\} > u'(c^*_t)$$

with positive probability if $\mathrm{Var}(u'(c^*_{t+1}) | z^{t+1}, \theta^t)$ has a positive probability of being larger than zero. Thus, we get 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.

3. TAXES AND WEDGES

At this stage, we have provided an intertemporal characterization of the Pareto optimal quantities in this dynamic Mirrlees world. What does this result say about taxes?

A Problem with the Natural Tax/Wedge Connection

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 linear tax, they face the first-order condition

$$(9) \qquad \beta E\{u'(c_{t+1})(1 - \delta + F_{K,t+1})(1 - \tau^k_{t+1}) | \theta^t, z^t\} = u'(c_t).$$

---

[4]Note that the proposition reduces to Theorem 1 of GKT if Z is a singleton (so there are no aggregate shocks). The proof of Proposition 1 also resembles the proof of Theorem 1 in GKT. Both proofs work by first establishing that the optimal allocation must satisfy a particular resource minimization problem, but the nature of the minimization problem is different. The GKT proof constructs the constraint set in the resource minimization problem by keeping the utility from consumption along all realizations of θ^T the same as in a putative optimum. In my proof, I construct the constraint set by keeping the utility differential between any two paths the same.

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 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) and Golosov and Tsyvinski (2004)).[5] Let $u(c) = \ln(c)$, $v(l) = l^2/2$, and $\beta = 1$. Suppose too that $T = 2$, $\Theta = \{0, 1\}$, $Z = \{1\}$, $F(K, Y) = rK + wY$, 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 $G = 0$. Then, we can rewrite the planner's problem as

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

$$\text{s.t.} \quad c_1 + K_2 = rK_1 + wy_1,$$

$$\frac{c_{2h}}{2} + \frac{c_{2l}}{2} = rK_2 + \frac{wy_{2h}}{2},$$

$$\ln(c_{2h}) - \frac{y_{2h}^2}{2} \geq \ln(c_{2l}),$$

$$c_{2h}, c_{2l}, y_{2h}, K_2, y_1 \geq 0.$$

(In this statement of the problem, I have set $y_{2l} = 0$, as would be true in a social optimum.) The solution to this problem must satisfy the first-order conditions

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

$$\frac{c_{2h}^*}{2} + \frac{c_{2l}^*}{2} = rK_2^* + \frac{wy_{2h}^*}{2},$$

$$\ln(c_{2h}^*) - \frac{y_{2h}^{*2}}{2} = \ln(c_{2l}^*),$$

$$\frac{1}{c_1^*} = \frac{r}{0.5c_{2h}^* + 0.5c_{2l}^*},$$

$$\frac{w}{c_{2h}^*} = y_{2h}^*,$$

$$y_1^* = \frac{w}{c_1^*}.$$

The obvious way to implement this allocation is as follows. Suppose that there is a single firm that owns the technology. The firm rents capital and labor

---

[5]See Chiappori et al. (1994) for a similar example.

in each period to produce output. In period 1, agents decide how much to work and how much capital to accumulate, given a linear tax on capital income. In period 2, the agents decide how much to work. If they generate zero income, they get a handout α_{2l}. If they earn positive income, they get a handout α_{2h} (which may be negative). So, the proceeds from the linear tax on capital income are being used to fund the subsidy to the disabled/unemployed agents in period 2.

More formally, define a tax mechanism in this world by $(\tau_k, \alpha_{2h}, \alpha_{2l})$. Then an equilibrium in this economy is a specification of $(c_1, c_{2h}, c_{2l}, y_1, y_{2h}, k_2)$ such that it solves

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

$$\text{s.t.} \quad c_1 + k_2 = rk_1 + wy_1,$$

$$c_{2h} = r(1 - \tau_k)k_2 + wy_{2h} + \alpha_{2h} \quad \text{if} \quad y_{2h} > 0,$$

$$c_{2h} = r(1 - \tau_k)k_2 + \alpha_{2l} \quad \text{if} \quad y_{2h} = 0,$$

$$c_{2l} = r(1 - \tau_k)k_2 + \alpha_{2l},$$

$$k_2, c_{2h}, c_{2l}, y_{2h}, y_1 \geq 0,$$

and markets clear

$$c_1 + k_2 = rk_1 + wy_1,$$

$$\frac{c_{2h}}{2} + \frac{c_{2l}}{2} = rk_2 + \frac{wy_{2h}}{2}.$$

Note that in equilibrium, $r\tau_k k_2 = \alpha_{2h}/2 + \alpha_{2l}/2$, which is the government's budget constraint.

Assume that the tax mechanism is such that the equilibrium value of $y_{2h} > 0$. Then the first-order conditions to the agent's problem are

$$\frac{1}{c_1} = r(1 - \tau_k)\left[\frac{0.5}{c_{2h}} + \frac{0.5}{c_{2l}}\right],$$

$$y_1 = \frac{w}{c_1},$$

$$\frac{w}{c_{2h}} = y_{2h},$$

$$\ln(c_{2h}) - \frac{y_{2h}^2}{2} \geq \ln(c_{2l}),$$

$$c_{2h} = r(1 - \tau_k)k_2 + wy_{2h} + \alpha_{2h},$$

$$c_{2l} = r(1 - \tau_k)k_2 + \alpha_{2l}.$$

How do we pick the tax mechanism so as to make the solution to these first-order conditions coincide with the equilibrium allocation? We set

$$(10) \qquad (1 - \tau_k) = \frac{[0.5c_h^* + 0.5c_l^*]^{-1}}{0.5/c_{2h}^* + 0.5/c_{2l}^*},$$

$$(11) \qquad \alpha_{2h} = c_{2h}^* - r(1 - \tau_k)K_2^* - wy_{2h}^*,$$

$$(12) \qquad \alpha_{2l} = c_{2l}^* - r(1 - \tau_k)K_2^*.$$

Then the equilibrium first-order conditions line up exactly with the social optimality first-order conditions. Note that the capital tax is positive.

However, there is a problem with this analysis. Under this tax mechanism, the optimal allocation satisfies the agent's first-order conditions. Nonetheless, the agent can do better than choose the optimal allocation. Why is this? Note first that

$$\frac{1}{c_1^*} < \frac{r(1 - \tau_k)}{c_{2l}^*},$$

because

$$\frac{1}{c_1^*} = r(1 - \tau_k)\left[\frac{0.5}{c_{2h}^*} + \frac{0.5}{c_{2l}^*}\right],$$

$$c_{2h}^* > c_{2l}^*.$$

Now suppose the agent saves $k_2^* + \varepsilon$ and set $y_{2h}^* = 0$. His utility from this budget-feasible plan is

$$\ln(c_1^* - \varepsilon) + \ln(c_{2l}^* + r(1 - \tau_k)\varepsilon)$$

as opposed to

$$\ln(c_1^*) + \frac{\ln(c_{2h}^*)}{2} - \frac{y_{2h}^{*2}}{2} + \frac{\ln(c_{2l}^*)}{2} = \ln(c_1^*) + \ln(c_{2l}^*).$$

Because $1/c_1^* < r(1 - \tau_k)/c_{2l}^*$, then the agent is better off using the new plan.

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 nontrivial, then he will prefer to pretend to be disabled when he is actually abled. Saving too much and shirking beats saving the right amount and telling the truth about one's type.

What this means is that the wedge 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.

Fixing the Problem

The above problem resulted from the fact that even though the agent was happy with saving k_2^* when he told the truth, he wanted to save a different amount when he lied. How do we fix this problem? One way is to tailor the tax rates on saving to the agent's announcements.

In particular, define a new tax mechanism $(\tau_{kh}, \tau_{kl}, \alpha_{2h}, \alpha_{2l})$. This mechanism works like this. If the agent produces 0 effective labor in period 2, then he receives a handout α_{2l} and his savings tax rate is τ_{kl}. If the agent produces a positive amount of effective labor in period 2, he receives a handout α_{2h} and his savings tax rate is τ_{kh}. His problem becomes

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

$$\text{s.t.} \quad c_1 + k_2 = rk_1 + wy_1,$$

$$c_{2h} = r(1 - \tau_{kh})k_2 + wy_{2h} + \alpha_{2h} \quad \text{if} \quad y_{2h} > 0,$$

$$c_{2h} = r(1 - \tau_{kl})k_2 + \alpha_{2l} \quad \text{if} \quad y_{2h} = 0,$$

$$c_{2l} = r(1 - \tau_{kl})k_2 + \alpha_{2l},$$

$$k_2, c_{2h}, c_{2l}, y_{2h}, y_1 \geq 0.$$

Define $(\tau_{kl}, \tau_{kh}, \alpha_{2l}, \alpha_{2h})$ so that

$$\frac{(1 - \tau_{kl})r}{c_{2l}^*} = \frac{1}{c_1^*},$$

$$\frac{(1 - \tau_{kh})r}{c_{2h}^*} = \frac{1}{c_1^*},$$

$$\alpha_{2i} = c_{2i}^* - r(1 - \tau_{ki})k_2^*, \quad i = h, l.$$

Then I claim that under this tax mechanism, the equilibrium allocation coincides with the optimal allocation.

Why? Suppose that the agent works $y_{2h}^* > 0$ in period 2 when abled. Then his solution for his other choice variables is

$$\frac{1}{c_1} = r\left[\frac{0.5(1 - \tau_{kh})}{c_{2h}} + \frac{0.5(1 - \tau_{kl})}{c_{2l}}\right],$$

$$\frac{w}{c_1} = y_1,$$

$$c_1 + k_2 = rk_1 + wy_1,$$

$$c_{2h} = r(1 - \tau_{kh})k_2 + wy_{2h} + \alpha_{2h},$$

$$c_{2l} = r(1 - \tau_{kl})k_2 + \alpha_{2l}.$$

The starred allocation satisfies these first-order conditions.

What if the agent works $y_{2h} = 0$ in period 2 when abled? Then his first-order conditions become

$$\frac{1}{c_1} = \frac{r(1 - \tau_{kl})}{c_{2l}},$$

$$\frac{w}{c_1} = y_1,$$

$$c_1 + k_2 = rk_1 + wy_1,$$

$$c_{2h} = r(1 - \tau_{kl})k_2 + \alpha_{2l},$$

$$c_{2l} = r(1 - \tau_{kl})k_2 + \alpha_{2l}.$$

Setting $(c_1, y_1, k_2, c_{2h}, c_{2l})$ equal to $(c_1^*, y_1^*, k_2^*, c_{2l}^*, c_{2l}^*)$ satisfies these first-order conditions. Hence, the agent is indifferent between working y_{2h}^* 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. Note that $\tau_{kl} > \tau_{kh}$; people who do not work get hit with a higher savings tax rate than those who work.

I want to emphasize that it is still optimal to have a wedge between the intertemporal marginal rate of substitution and the intertemporal marginal rate of transformation. However, the only way to decentralize this wedge using linear taxes on savings is to use state-contingent tax rates.

4. A GENERAL IMPLEMENTATION

In this section, I use the above two-period analysis to build a general implementation of an optimal allocation. In the implementation, taxes on wealth are restricted to be linear.

Taxes in a Sequential Markets Economy

I begin by describing a notion of equilibrium, given arbitrary *nonlinear* taxes on labor income and *linear* taxes on wealth. In the economy, there is a single representative firm that owns the technology of production, and rents capital and hires effective labor in each period. The firm takes period t capital rents r_t and period t wages w_t as given.

The agents in the economy all begin life with K_1^* units of capital. They trade capital, labor, and consumption in a sequence of competitive markets. They face a labor tax schedule $\psi : \mathbb{R}_+^T \times Z^T \to \mathbb{R}^T$, where ψ_t is (y^t, z^t)-measurable. According to this schedule, an agent who has an effective labor sequence $\{y_t\}_{t=1}^T$ pays labor taxes $\psi_t((y_t)_{t=1}^T, z^T)$ in period t.

The agents also face a linear tax on their wealth. The tax rate, though, may depend on their effective labor history. Thus, let $\tau : \mathbb{R}_+^T \times Z^T \to \mathbb{R}^T$, where τ_t is (y^t, z^t)-measurable. Then an agent with wealth W_t at the beginning of period t and with effective labor $(y_t)_{t=1}^T$ pays wealth taxes $\tau_t((y_t)_{t=1}^T, z^T)W_t$.

Formally, the typical agent takes as given a tax system (ψ, τ) and prices (r, w). He then has a choice problem of the form

$$\max_{c,y,k} \sum_{t=1}^T \beta^{t-1} \int_{z^T \in Z^T} \int_{\theta^T \in \Theta^T} \left\{ u(c_t(\theta^T, z^T)) - v\left(\frac{y_t(\theta^T, z^T)}{\phi_t(\theta^T, z^T)} \right) \right\} d\mu_\Theta \, d\mu_Z$$

$$\text{s.t.} \quad c_t(\theta^T, z^T) + k_{t+1}(\theta^T, z^T)$$
$$\leq \left(1 - \tau_t\big(y(\theta^T, z^T), z^T\big)\right)(1 - \delta + r_t(z^T))k_t(\theta^T, z^T)$$
$$+ w_t(z^T)y_t(\theta^T, z^T) - \psi_t\big(y(\theta^T, z^T), z^T\big) \quad \text{for all } (\theta^T, z^T),$$

(c_t, k_{t+1}, y_t) is (θ^t, z^t)-measurable and nonnegative,

$k_1 \leq K_1^*$.

Given a tax system (ψ, τ), an equilibrium in this economy is a specification of (c, y, k) and (r, w) such that (c, y, k) solves the agent's problem, given ψ, τ, r, and w, such that $r_t(z^T) = F_{kt}(K_t(z^T), Y_t(z^T), z^T)$ and $w_t(z^T) = F_{Yt}(K_t(z^T), Y_t(z^T), z^T)$, and such that markets clear for all t and z^T:

$$\int_{\theta^T \in \Theta^T} c_t(\theta^T, z^T) \, d\mu_\Theta + G_t(z^T) + K_{t+1}(z^T)$$
$$= F_t\big(K_t(z^T), Y_t(z^T), z^T\big) + (1 - \delta)K_t(z^T),$$

$$K_t(z^T) = \int_{\theta^T \in \Theta^T} k_t(\theta^T, z^T) \, d\mu_\Theta,$$

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

Note that in this definition of equilibrium, the government's budget is balanced in every period:

$$G_t(z^T) = \int_{\theta^T \in \Theta^T} \tau_t\big(y(\theta^T, z^T), z^T\big)\big\{(1 - \delta + r_t(z^T))k_t(\theta^T, z^T)\big\} \, d\mu_\Theta$$

$$+ \int_{\theta^T \in \Theta^T} \psi_t\big(y(\theta^T, z^T), z^T\big) \, d\mu_\Theta.$$

Elements of the Implementation

In this subsection, I describe how to set ψ and τ such that the resulting equilibrium allocation is optimal. I make one key assumption: in the optimal allo-

cation, current consumption depends on current and past skills only through current and past effective labor. This assumption allows me to implement an optimal allocation using a tax schedule that is written in terms of effective labor, not in terms of θ^t.

Formally, suppose (c^*, y^*, K^*) is a socially optimal allocation. Let DOM_t be a subset of $\mathbb{R}_+^T \times Z^T$ defined as follows: (y^T, z^T) is in DOM_t if and only if there exists θ^T in Θ^T, $(y_s')_{s=t+1}^T$ in $\mathbb{R}_+^{T-t}$, and $(z_s')_{s=t+1}^T$ in Z^{T-t} such that

$$\left(y^t, (y_s')_{s=t+1}^T, z^T\right) = \left(y^*(\theta^T, z^T), z^T\right),$$

where $z^T = (z^t, (z_s')_{s=t+1}^T)$ for some $(z_s')_{s=t+1}^T$ in Z^{T-t}. In words, (y^T, z^T) is in DOM_t if in the socially optimal allocation, there exists some type in Θ^T that receives the effective labor history y^t when the public history is z^t. Note that $DOM_t \subseteq DOM_{t-1}$ for all t. I then make the following assumption.

ASSUMPTION 1: *There exists a sequence of functions* $\widehat{c}^* = (\widehat{c}_t^*)_{t=1}^T$, *where* $\widehat{c}_t^*: DOM_t \to \mathbb{R}_+$, $\widehat{c}_t^*$ *is* (y^t, z^t)-*measurable*, *and*

$$\widehat{c}_t^*\left(y^*(\theta^T, z^T), z^T\right) = c_t^*(\theta^T, z^T)$$

for all (θ^T, z^T).

In a one-period setting, this assumption is a trivial consequence of incentive compatibility. However, in a dynamic setting, an agent may receive information about both his current skills and his future skills. The planner induces truthful revelation about current skills by having consumption and effective labor covary positively across agents as in the static case. In contrast, the planner induces truthful revelation about *future* skills by offering agents more momentary utility today in exchange for more effective labor in the future. It is optimal for this increase in current momentary utility to come from both consumption and leisure. Thus, information about future skills induces a negative covariance between current consumption and current effective labor. It is then no longer obvious that current consumption depends on current and past skills only through current and past effective labor.

It is easy to prove that Assumption 1 is satisfied in an environment like that of Albanesi and Sleet (2005), in which skill shocks are i.i.d., because the agent receives no information about his future skills. If skill shocks are mean-reverting, it is intuitively plausible that the effect of information about future skills is outweighed by the effect of information about current skills. This intuition would suggest that Assumption 1 continues to be valid when skills are mean-reverting. (This intuition is further supported by the fact that when skills are fixed over time, it is optimal for consumption and effective labor to covary positively in period 1.)

However, suppose skills actually are an explosive process in the sense that a high current level of skills implies that the future growth rate of skills is high.

In this case, the above intuition suggests that Assumption 1 may be violated. I present an example of this kind in Appendix B.[6]

Given Assumption 1, I next construct the details of the tax system that implements the optimal allocation (c^*, y^*, K^*). I begin with the optimal tax on wealth. Given the optimal allocation (c^*, y^*, K^*), we know from Proposition 1 that there exists $\lambda_{t+1}^* : Z^T \to \mathbb{R}_+$ such that λ_{t+1}^* is z^{t+1}-measurable and

$$(13) \qquad \lambda_{t+1}^* = \beta \frac{[E(u'(c_{t+1}^*)^{-1} | \theta^t, z^{t+1})]^{-1}}{u'(c_t^*)}.$$

Define $\tau_{t+1}^* : \mathbb{R}_+^T \times Z^T \to \mathbb{R}$ by

$$(14) \qquad \tau_{t+1}^*(y^T, z^T)$$
$$= \begin{cases} 1 - \dfrac{\lambda_{t+1}^*(z^T) u'(\widehat{c}_t^*(y^T, z^T))}{\beta u'(\widehat{c}_{t+1}^*(y^T, z^T))}, & \forall (y^T, z^T) \text{ in } DOM_{t+1}, \\ 1, & \forall (y^T, z^T) \text{ not in } DOM_{t+1}. \end{cases}$$

Verbally, τ_{t+1}^* equates the ex post individual marginal rate of substitution with the ex post societal marginal rate of transformation if (y^T, z^T) is in DOM_{t+1}. If (y^T, z^T) is not in DOM_{t+1}, the agent loses all of his wealth. Note that τ_{t+1}^* is (y^{t+1}, z^{t+1}) measurable.

Next, I describe the labor tax code. First, define

$$(15) \qquad MPK_t^*(z^t) \equiv F_K\big(K_t^*(z^{t-1}), Y_t^*(z^t), z^t\big),$$

$$(16) \qquad MPL_t^*(z^t) \equiv F_Y\big(K_t^*(z^{t-1}), Y_t^*(z^t), z^t\big).$$

Then let $(\psi^{**}, \widehat{k}^*) : DOM_T \to \mathbb{R}^T \times \mathbb{R}_+^T$ be defined so that

$$(17) \qquad \widehat{c}_t^*(y^T, z^T) + \widehat{k}_{t+1}^*(y^T, z^T)$$
$$= \big(1 - \tau_{t+1}^*(y^T, z^T)\big)\big(1 - \delta + MPK_t^*(z^T)\big)\widehat{k}_t^*(y^T, z^T)$$
$$+ MPL_t^*(z^T) y_t - \psi_t^{**}(y^T, z^T),$$

$$(18) \qquad \int_{\theta^T \in \Theta^T} \widehat{k}_{t+1}^*\big(y^*(\theta^T, z^T), z^T\big) d\mu_\Theta = K_t^*(z^T),$$

$$(19) \qquad \widehat{k}_1^* = K_1^*$$

for all t and for all (y^T, z^T) in DOM_T. Here ψ^{**} describes the labor taxes, given that the agent chooses an effective labor sequence in DOM_T and $\widehat{k}^*$ describes the agent's capital holdings so as to satisfy the flow budget constraint.

---

[6]I used an insight provided by Ivan Werning to construct this example.

However, this does not quite complete the specification of the labor taxes; we need to describe labor taxes when the agent chooses an effective labor sequence not in DOM_T. First, note that for all t, there exists an extension $\psi^{**}_{\text{ext},t}: DOM_t \to \mathbb{R}$ such that $\psi^{**}_{\text{ext},t}(y^T, z^T) = \psi^{**}(y^T, z^T)$ for all (y^T, z^T) in DOM_T and such that $\psi^{**}_{\text{ext},t}$ is (y^t, z^t)-measurable. Then define $\psi^*: \mathbb{R}_+^T \times Z^T \to \mathbb{R}^T$ by

$$(20) \qquad \psi^*_t(y^T, z^T) = \begin{cases} \psi^{**}_{\text{ext},t}(y^T, z^T), & \text{if } (y^T, z^T) \text{ is in } DOM_t, \\ 2y_t w_t(z^T), & \text{for any } (y^T, z^T) \text{ not in } DOM_t. \end{cases}$$

This function ψ^* describes an optimal labor tax code. As is made clear in the proof of Proposition 2, the tax when (y^T, z^T) is outside of DOM_t is sufficient to make only effective labor strategies that lie in DOM_T budget-feasible.

A Second Welfare Theorem

I now claim that (c^*, y^*, k^*) is an equilibrium given (ψ^*, τ^*), where $k^*(\theta^T, z^T) = \widehat{k}^*(y^*(\theta^T, z^T), z^T)$. As usual, we use social shadow values to construct equilibrium prices. Let $r^*_t = MPK^*_t$ and let $w^*_t = MPL^*_t$. Clearly, given these prices, the firm's first-order conditions are satisfied. The optimal allocation satisfies market clearing. Hence, we need only verify that given prices (r^*, w^*), and tax system (ψ^*, τ^*), the allocation (c^*, y^*, k^*) is individually optimal for an agent in the economy. To prove this claim, we need the following proposition.

PROPOSITION 2: *Given prices (r^*, w^*) and tax system (ψ^*, τ^*), and given that the typical agent chooses a budget-feasible y', his optimal choices of (c, k) are $c'_t(\theta^T, z^T) = \widehat{c}^*_t(y'(\theta^T, z^T), z^T)$ and $k'_t(\theta^T, z^T) = \widehat{k}^*_t(y'(\theta^T, z^T), z^T)$.*

PROOF: First, note that any budget-feasible y' satisfies the property that $(y'(\theta^T, z^T), z^T)$ is in DOM_T for all (θ^T, z^T). To see this, suppose y' is a strategy such that $(y'(\theta^T, z^T), z^T)$ is not in DOM_T for some (θ^T, z^T). Note that $DOM_T = \bigcap_{t=1}^T DOM_t$, even if T is infinite. Hence, there always exists some finite t such $(y'(\theta^T, z^T), z^T)$ is not in DOM_t. In this period t, the agent loses all of his accumulated wealth and owes twice his current labor income in taxes. This is impossible for him to afford, so such a y' is not budget-feasible.

Define $\tau': \Theta^T \times Z^T \to \mathbb{R}$ by $\tau'(\theta^T, z^T) = \tau^*(y'(\theta^T, z^T), z^T)$ and $\psi'(\theta^T, z^T) = \psi^*(y'(\theta^T, z^T), z^T)$. Then, given that the agent chooses effective labor strategy y', his intertemporal consumption problem becomes

$$\max_{c,k} \sum_{t=1}^T \beta^{t-1} \int_{\theta^T \in \Theta^T} \int_{z^T \in Z^T} u(c_t) \, d\mu_\Theta \, d\mu_Z$$

$$\text{s.t.} \quad c_t + k_{t+1} = (1 - \tau'_t)(1 - \delta + r_t)k_t + w_t y'_t - \psi'_t,$$

c_t, k_{t+1} are (θ^t, z^t)-measurable and nonnegative,

$k_1 \le K_1^*.$

The first-order conditions to this problem are

$$\beta E\{(1 - \tau_{t+1}')u'(c_{t+1})(1 - \delta + r_{t+1})|\theta^t, z^t\} = u'(c_t),$$
$$c_t + k_{t+1} = (1 - \tau_t')(1 - \delta + r_t)k_t + w_t y_t' - \psi_t'.$$

If T is finite, the first-order conditions are obviously necessary and sufficient (once y' is fixed). If T is infinite, the necessity and sufficiency is implied by the assumption that u is bounded.

My claim is that (c', k') satisfy these first-order conditions. Clearly, from the definition of ψ^* and $\widehat{k}^*$, they satisfy the flow budget constraints. What about the Euler equations? We know that for all (y^T, z^T) in DOM_T,

$$\left(1 - \tau_{t+1}^*(y^T, z^T)\right)\lambda_{t+1}^*(z^T)^{-1} = \beta^{-1}\frac{u'(\widehat{c}_t^*(y^T, z^T))}{u'(\widehat{c}_{t+1}^*(y^T, z^T))}$$

and so for all (θ^T, z^T),

$$\left(1 - \tau_{t+1}'(\theta^T, z^T)\right)\lambda_{t+1}^*(z^T)^{-1} = \left(1 - \tau_{t+1}^*\big(y'(\theta^T, z^T), z^T\big)\right)\lambda_{t+1}^*(z^T)^{-1}$$
$$= \beta^{-1}\frac{u'(\widehat{c}_t^*(y'(\theta^T, z^T), z^T))}{u'(\widehat{c}_{t+1}^*(y'(\theta^T, z^T), z^T))}$$
$$= \beta^{-1}\frac{u'(c_t'(\theta^T, z^T))}{u'(c_{t+1}'(\theta^T, z^T))}.$$

Hence,

$$\beta E\{(1 - \tau_{t+1}')u'(c_{t+1}')(1 - \delta + r_{t+1})|\theta^t, z^t\} - u'(c_t')$$
$$= \left[E\{\lambda_{t+1}^*(1 - \delta + r_{t+1})|\theta^t, z^t\} - 1\right]u'(c_t')$$
$$= 0.$$

This proves the proposition.   Q.E.D.

Proposition 2 considers an agent who chooses an arbitrary effective labor strategy y' that is budget-feasible. Because $(y'(\theta^T, z^T), z^T) \in DOM_T$ for all (θ^T, z^T), there exists a reporting strategy $\sigma': \Theta^T \times Z^T \to \Theta^T \times Z^T$ that satisfies

$$y^*(\sigma'(\theta^T, z^T)) = y'(\theta^T, z^T) \quad \text{for all } (\theta^T, z^T).$$

The content of Proposition 2 is that if an agent chooses y', it is optimal for him to choose an asset allocation plan that gives him consumption $c'(\theta^T, z^T)$, where for all (θ^T, z^T),

$$
\begin{aligned}
c'(\theta^T, z^T) &= \widehat{c}^*\big(y'(\theta^T, z^T), z^T\big) \\
&= \widehat{c}^*\big(y^*(\sigma'(\theta^T, z^T)), z^T\big) \\
&= c^*(\sigma'(\theta^T, z^T)).
\end{aligned}
$$

We can now use Proposition 2 to show that given prices and taxes, a consumer's optimal choice from his budget set is (c^*, y^*, k^*), where $k^*(\theta^T, z^T) = \widehat{k}^*(y^*(\theta^T, z^T), z^T)$. To complete the argument, we need only show that the optimal effective labor strategy is y^*. We know from Proposition 2 that an agent who chooses y' and then chooses an optimal consumption-savings strategy, receives utility $W(\sigma'; c^*, y^*)$, where σ' is defined as above. However, this utility is no larger than $W(\sigma_{TT}; c^*, y^*)$, which can be achieved by choosing y^* and then saving optimally. The agent is weakly better off choosing y^*.

Thus, we have successfully implemented the optimal allocation as an equilibrium allocation using the tax mechanism (ψ^*, τ^*). In the implementation, agents can only trade capital and consumption. However, it is straightforward to extend the analysis to allow agents to trade z^{t+1}-contingent claims that are available in zero net supply. Indeed, the structure of the optimal taxes τ^* is left unaltered by adding these financial asset markets.

5. IMPLICATIONS FOR OPTIMAL TAXES

In this subsection, I discuss some properties of the optimal tax system derived in Section 4.

Zero Taxes

It is easy to prove that in the above implementation, the expected wealth tax rate in period $(t + 1)$, conditional on (θ^t, z^{t+1}), is zero. Define

$$
\tau^{**}_{t+1}(\theta^T, z^T) = \tau^*_{t+1}\big(y^*(\theta^T, z^T), z^T\big) \quad \text{for all } (\theta^T, z^T).
$$

By construction,

$$
(1 - \tau^{**}_{t+1}) = \beta^{-1}\lambda^*_{t+1} u'(c^*_{t+1})^{-1} u'(c^*_t),
$$

so that the after-tax ex post marginal rate of substitution is set equal to the

social discount factor. Then

$$
\begin{aligned}
E\{(1 - \tau_{t+1}^{**})|\theta^t, z^{t+1}\} &= E\{\beta^{-1}\lambda_{t+1}^* u'(c_{t+1}^*)^{-1} u'(c_t^*)|\theta^t, z^{t+1}\} \\
&= \beta^{-1}\lambda_{t+1}^* u'(c_t^*) E\{u'(c_{t+1}^*)^{-1}|\theta^t, z^{t+1}\} \\
&\qquad \text{by } (\theta^t, z^{t+1})\text{-measurability of } \lambda_{t+1}^* u'(c_t^*) \\
&= 1,
\end{aligned}
$$

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

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 that consists of saving too much and then working too little when skilled in the following period.

This result implies immediately that any given individual's expected wealth tax rate is zero. However, there is a second, slightly more subtle, implication: under any optimal system, wealth taxes are purely redistributional because the government raises no net revenue from them in any public history z^{t+1}. This result may seem surprising at first because many individual capital holdings processes k^* are consistent with optimality. Nonetheless, 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:

$$
\begin{aligned}
\int_{\theta^T \in \Theta^T} &\tau_{t+1}^{**}(\theta^T, z^T) k_{t+1}^*(\theta^T, z^T)\big(1 - \delta + MPK_{t+1}^*(z^T)\big)\, d\mu_\Theta \\
&= \big(1 - \delta + MPK_{t+1}^*(z^T)\big) E(\tau_{t+1}^{**} k_{t+1}^*|z^{t+1}) \\
&= \big(1 - \delta + MPK_{t+1}^*(z^T)\big) E\big(E(\tau_{t+1}^{**}|\theta^t, z^{t+1}) k_{t+1}^*|z^{t+1}\big) \\
&= 0.
\end{aligned}
$$

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.

The Key Assumption

The above zero mean wealth tax result is a direct consequence of Propositions 1 and 2. These propositions are, respectively, about the socially and individually optimal intertemporal allocation of consumption. Both of them rely on one crucial feature of the environment: individuals' intertemporal marginal

rates of substitution for consumption (IMRS) are publicly observable.[7] This aspect of the setup means that private information does not directly impact on the intertemporal first order conditions for consumption.

We can use this point to better understand how the zero mean tax result does and does not generalize. Propositions 1 and 2 are *not* valid if preferences are not additively separable between consumption and labor[8] (because then unobservable labor affects the consumption IMRS). They are also not valid if there are unobservable shocks to agents' marginal utility of consumption or hidden endowment shocks. On the other hand, Propositions 1 and 2 do remain valid if agents can engage in multiple tasks or if they can invest in multiple forms of capital.

More subtly, the propositions remain valid even if skills are endogenous. Suppose agents can engage in human capital accumulation, but it is not possible for outsiders to monitor their investment. As long as this hidden investment is only in the form of time, then Propositions 1 and 2, and the zero mean wealth tax result go through. However, if the unmonitorable investment includes consumption goods, the agents' marginal utilities of consumption are no longer public information, and neither of Propositions 1 and 2 is true.

Comparison with Ramsey Taxation

There is a large literature on optimal capital and labor income taxation that uses what might be termed the Ramsey approach. Under this approach, it is assumed that all agents are identical. The government is restricted to use linear taxes on capital and labor income, but these taxes are allowed to depend in arbitrary ways on the history of aggregate shocks. The government is not allowed to impose a tax on capital income in period 1, because such a tax is lump-sum. As well, the government and agents can trade in a complete set of date- and state-contingent claims to consumption.

Suppose the agents' utility from consumption is given by $u(c) = c^{1-\gamma}/(1-\gamma)$ or $u(c) = \ln(c)$. Then their overall utility function is homothetic over consumption goods at different dates and separable between consumption and labor. Chari, Christiano, and Kehoe (1994) show that under this assumption,[9] the optimal capital tax is positive in period 2 and is zero in any period $t > 2$. As they explain, the zero tax result is an application of the uniform commodity taxation theorem, where the various commodities are consumption at different dates and states.

---

[7]GKT (2003) also emphasize this aspect of the environment.

[8]Kocherlakota (2004) shows that if Θ is finite, and preferences are not additively separable between consumption and labor, the optimal wealth taxes are necessarily nondifferentiable as a function of wealth.

[9]CCK generalize Chamley's (1986) result for deterministic economies by allowing for shocks to government purchases. Atkinson and Sandmo (1980) consider the Ramsey problem in an overlapping generations context.

How do the Ramsey capital taxes derived by Chari, Christiano, and Kehoe compare to the optimal tax system described in Section 4? The systems agree in the sense that in any period $t > 2$, they imply that aggregate capital tax collections should be zero. However, they differ in two important respects. First, in period 2, the Ramsey taxes are nonzero. The government is using these period 2 taxes as an imperfect substitute for the instrument that it would like to use: the lump-sum period 1 taxes on capital income. In the tax system derived in Section 4, the optimal capital tax collections are zero in every date and state, including period 2. Lump-sum taxes are always available—there is no need to use the period 2 taxes to substitute for them.

Of course, this difference between the Mirrlees and Ramsey systems vanishes after one period. The more long-lasting difference is that the Mirrlees system requires the use of nonzero capital taxes, while the Ramsey system does not. The nonzero capital taxes play an essential role in inducing optimal savings and effort given the private information friction. As we shall see in Section 6, the welfare losses associated with having uniform capital taxes across agents can be substantial in the setting with private information.

Remarks

This subsection points out several other important features of the optimal tax system described in Section 4.

REMARK 1: The labor income taxes ψ^* are indeterminate. There is a large set of labor income tax schedules and individual capital holdings $(\psi^*, \widehat{k}^*)$ that can be used as part of a tax mechanism that supports a given optimal allocation (c^*, y^*, K^*). Loosely speaking, these various optimal tax systems differ in terms of the timing of tax collections. For example, suppose $T = 2$, but people earn labor income only in period 1 (which implies in turn that optimal capital taxes are zero for everyone). Suppose one optimal tax system is to tax agents with high income \$10,000 in period 1 and not tax agents with low income. Then we can construct another optimal tax system by taxing high-income agents \$1,000 in period 1, and \$9,000$(1 + r)$ in period 2, while transferring \$9,000 in period 1 to low-income agents and then taxing them \$9,000$(1 + r)$ in period 2. This tax system is also optimal, because the present value of the tax burden for each possible report is kept the same. However, individual capital holdings in equilibrium change (high-income agents hold less capital under the second system, while low-income agents hold more).

REMARK 2: In the above class of optimal mechanisms, the government's budget is balanced in every period. However, using the reasoning in the above paragraph, it is possible to construct optimal tax structures with alternative streams of government debt: there is simply no notion of an optimal debt structure in this world. This is a consequence of the richness of the tax structure: as

Bassetto and Kocherlakota (2004) emphasize, when taxes can depend on past incomes, debt is irrelevant.

REMARK 3: The tax system is linear in wealth, but it is not arbitrage-free. Consider an agent who faces no future skill risk: in the optimal tax system, he faces no wealth taxes. Other agents face wealth tax risk that is correlated with their equilibrium consumptions. They would like to shield their wealth from taxes by making an off-book loan to the "no-risk" agent that allows him to do all of the capital accumulation in the economy.[10]

REMARK 4: An agent's expected wealth tax rate is zero, but there is a wedge between the social intertemporal marginal rate of transformation and the agent's marginal rate of substitution. This wedge is created by tax risk. The agent faces high wealth taxes when his consumption is low and low wealth taxes when his consumption is high. The optimal tax system inflates the agent's "consumption beta" of investing in physical capital and thereby deters saving.

REMARK 5: I focus on symmetric Pareto optima, in which the planner treats all agents the same. It is possible, though, to extend the results to the case of asymmetric Pareto optima, in which agents are given different initial lifetime utility promises. In particular, given an asymmetric Pareto optimal allocation, it can be implemented using a tax system that is linear in wealth and in which the expected wealth tax rate is zero for all agents. This optimal tax system is a nonlinear function of initial lifetime utility promises and current and past labor incomes.

6. A NUMERICAL EXAMPLE AND ITS LESSONS

Section 5 focused on one particular aspect of the optimal tax system described in Section 4: the average wealth taxes are zero in every date and state. In this section, I use a numerical example to explore other aspects of the optimal tax system constructed in Section 4.

Before proceeding to the example, however, a word of warning is in order. It would be desirable to use an infinite horizon example as in Albanesi and Sleet (2005). However, to answer the questions of interest, the example would have to include aggregate shocks, persistent hidden state variables, and probably should allow for endogenous physical state variables. At a conceptual

---

[10]Like I do, Golosov and Tsyvinski (2005) consider an optimal tax problem in a dynamic Mirrlees economy. However, they assume that the government is restricted to using arbitrage-free taxes on wealth. (They motivate this restriction, like Guesnerie (1995), by assuming that agents can freely engage in intertemporal side trades. I suspect that they could also have motivated it by assuming that agents could form coalitions, as in Hammond (1987).) The optimal tax is typically nonzero in their setting. Its sign depends on details of the data generation process for skills.

level, it is known how to attack problems of this kind, thanks to the work of Fernandes and Phelan (2000) and Doepke and Townsend (2005). Practically, it is still impossible to implement their procedures in an example that includes the elements of interest.

For these reasons, I focus on a two-period example. Obviously, this analysis can only be viewed as suggestive of what would happen in a full-blown infinite horizon application. Nonetheless, I believe that the example contains some useful lessons beyond what we learn from the infinite horizon analysis of Albanesi and Sleet (2005) (who require skill shocks to be i.i.d. over time).

Details of the Example

The example works as follows. I assume that $T = 2$, $u(c) = \ln(c)$, $v(y) = y^2/2$, and $\beta = 1$. In terms of the shocks, I assume that

$$Z = \{0, 1\},$$

$$\Theta = \{0.8, 1.2\},$$

$$\Pr(\theta_1 = 0.8) = 0.5,$$

$$\Pr(\theta_2 = \theta | \theta_1 = \theta) = \pi \quad \text{for} \quad \theta = 0.8, 1.2,$$

$$\Pr(z_1 = 1) = 0.5,$$

$$\Pr(z_2 = z | z_1 = z) = 0.9 \quad \text{for} \quad z = 0, 1.$$

I leave π as a free parameter that determines the persistence of the idiosyncratic shocks. Given the Markov structure, the autocorrelation of the idiosyncratic shocks equals $2\pi - 1$.

The production technology works as follows. I set $K_1 = 1$. The production technology in both periods has the same form,

$$F_t(K, Y) = K + Y,$$

and capital fully depreciates ($\delta = 1$). Hence, the aggregate shocks do not affect the aggregate production function.

Finally, I set government purchases to be

$$G_t(z_t) = \begin{cases} g_H, & \text{if } z_t = 0, \\ 0, & \text{if } z_t = 1. \end{cases}$$

Again, I treat g_H as a free parameter to be varied. Thus, this example can be used to analyze the impact of g_H and π on endogenous variables like capital tax rates.

I solve the planner's problem using a direct numerical procedure. I form a relaxed program by discarding the incentive constraints for the low types at

every node. I then write down the first order conditions of the relaxed problem and use the nonlinear equation solver in GAUSS (eqsolve) to find the solution to these first-order conditions. (There are 58 equations in 58 unknowns, which is easily manageable.) I conclude by checking to make sure that the incentive constraints of the low types are in fact satisfied by the solution to the relaxed problem.

Impact of Idiosyncratic Shock Persistence

I begin by varying the idiosyncratic shock parameter π. I assume that $g_H = 0$, so that there are no aggregate shocks in the economy. In this two-period setting, optimal capital tax rates depend on skill realizations in periods 1 and 2. I plot these in Figure 1 as a function of the autocorrelation of the individual-
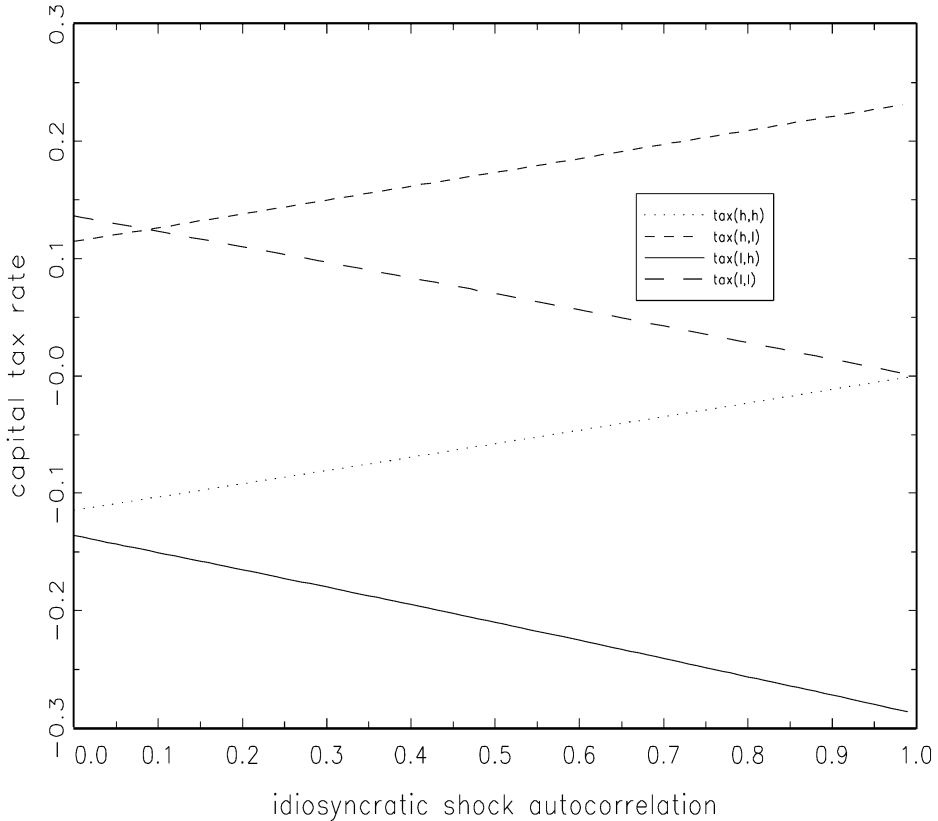

FIGURE 1.—Impact of the autocorrelation of individual-specific skill shocks on optimal capital tax rates.

specific skill shocks. The major lesson of this figure is that the optimal capital tax rates depend crucially on the persistence of the skill shock. In particular, as the shocks become more persistent, the capital tax rates associated with high probability transitions (low to low and high to high) fall in absolute value. In contrast, the capital tax rates associated with low probability transitions rise in absolute value. These findings are not entirely unexpected given the zero mean tax rate result derived in Section 5.

In terms of magnitudes, it is worth emphasizing that these tax rates are being applied to the entire stock of capital, not just to flow income (which in this model economy is actually zero). A capital tax rate of 5% is equivalent to a capital income tax rate of 100%, if the rate of return on capital is 5%.

Impact of Shocks to Government Purchases

Much of the Ramsey literature on optimal capital taxation is concerned with how shocks to government purchases should affect the mix between capital taxes and labor taxes. In this subsection, I explore how government purchases affect capital tax rates in this two-period economy. I set $\pi = 0.75$ (which is equivalent to an autocorrelation of 0.5 in skill shocks), and then solve the planner's problem for $g_H = 0$ (the no aggregate shocks case) and $g_H = 0.3$. This shock to government purchases is large: in the computed optimal allocations, the ratio of government purchases to output is about 1/3 when $g_H = 0.3$ and is exactly zero when $g_H = 0$. This kind of jump is about the same order of magnitude as that seen in the United States at the onset of World War II.

The results of this exercise are depicted in Table I. The zero mean wealth tax result says that for any level of government purchases, the cross-sectional average of the capital tax rate is zero; Table I is consistent with this result. More generally, Table I documents that period 2 government purchases are the key determinant of the capital tax rates faced by any given individual. The capital tax rates are highly similar in rows 1, 2, and 3 in Table I, because g_2 is equal to 0 in all three rows. The capital tax rates are highly similar in rows 4 and 5 in Table I, because g_2 is equal to 0.3 in both rows. However, the cross-sectional variance of capital tax rates if $g_2 = 0.3$ is quite different than if $g_2 = 0$. When $g_2 = 0.3$, the cross-sectional standard deviation of capital tax rates is

TABLE I

AGGREGATE SHOCKS AND CAPITAL TAX RATES

	τ_{hh}	τ_{hl}	τ_{lh}	τ_{ll}
$g_H = 0$ (no agg. shocks)	−0.058	0.173	−0.210	0.070
$(g_1, g_2) = (0, 0)$	−0.057	0.171	−0.208	0.069
$(g_1, g_2) = (0.3, 0)$	−0.057	0.170	−0.221	0.074
$(g_1, g_2) = (0, 0.3)$	−0.072	0.216	−0.254	0.085
$(g_1, g_2) = (0.3, 0.3)$	−0.072	0.216	−0.269	0.090

around 0.125 for agents who were highly skilled in period 1. When $g_2 = 0.3$, the cross-sectional standard deviation of capital tax rates is only 0.09 for those same agents. As we shall see, this difference of 3.5% between the two standard deviations is comparable to the average capital tax rate in the OECD.

Intuitively, when purchases are high in period 2, it is socially optimal for effective labor to be high. It becomes even more socially beneficial to deter individuals from pursuing the strategy of saving and shirking, and therefore socially optimal to impose relatively high capital taxes on those with low labor income. Thus, the cross-sectional mean of capital tax rates is always zero, but the cross-sectional variance of capital tax rates covaries positively with aggregate government purchases.

Positivity or Labor Income Independence: What Is Worse?

This is a normative paper. It describes the properties of a particular optimal tax system, given a plausible specification of informational and technological restrictions. In the real world, there is no reason to expect politically determined tax systems to be the same as the socially optimal one described in this paper.

Nonetheless, it is useful to consider how real-world tax systems can be improved. Tax systems in developed countries have two striking differences from the prescriptions of this paper. First, they have a higher mean capital tax. They are substantially different from zero—the average tax rate on net capital income is about 50% (Carey and Rabesona (2002)). This tax on the *flow* of capital income is equivalent to a tax of about 2.5% on the capital *stock,* assuming that the net return on capital is roughly 5%.

The second important prescription of this paper is that capital taxes should be systematically higher on individuals with unexpectedly low labor income. There is some element of this in the United States welfare system, because welfare payments are asset-tested (see Golosov and Tsyvinski (2004) for a full discussion). Nonetheless, I am not aware of income tax codes that have explicit "regressive" wealth taxes of the form described in this paper.

Thus, relative to the optimal tax system described in Section 4, real-world tax systems have two apparent deficiencies: the tax on capital is too *high* and too *uniform* (across people). As mentioned in the introduction, the prior Ramsey literature has emphasized the welfare losses associated with the first problem. How large are the welfare losses associated with the second problem?

To answer this question, I compare agents' welfare under three regimes in the two-period model with $g_H = 0$. The first is the optimal tax system described in Section 4. The second is a similar tax system, except that it is only optimal subject to the restriction that the average capital tax rate in period 2 is 2.5% (the average tax rate on the capital stock in OECD countries). In this system, the capital tax rates are allowed to depend in an arbitrary way on an agent's history of labor incomes. The third is the optimal tax system given that the

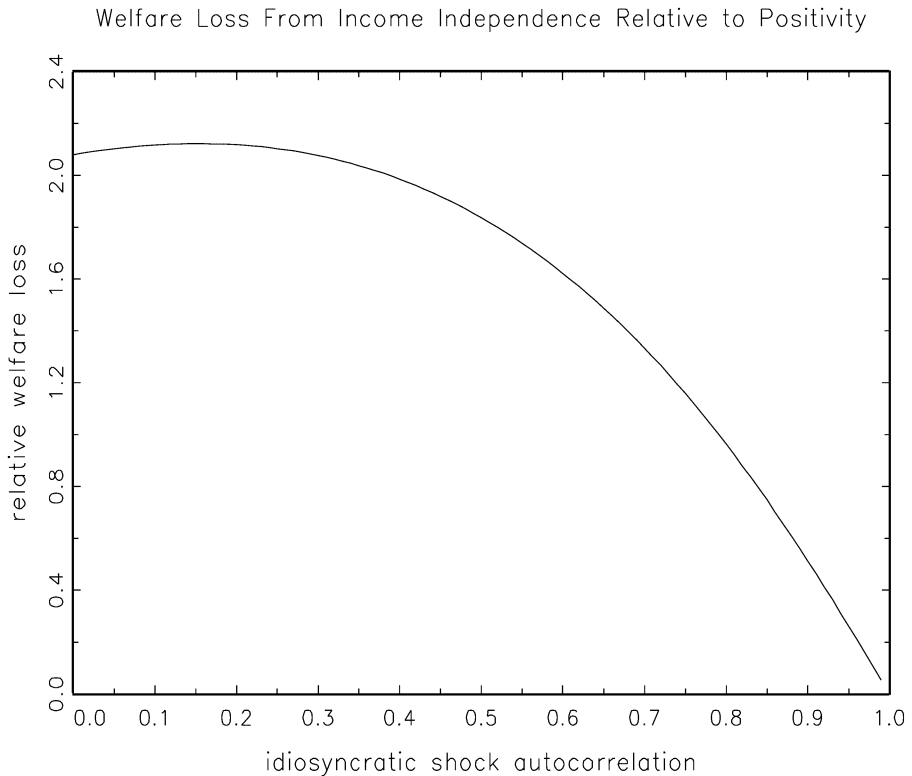Welfare Loss From Income Independence Relative to Positivity

FIGURE 2.—The ratio of two welfare losses. The numerator is the welfare loss from imposing a zero capital tax rate (independent of second-period labor income) as opposed to the optimal tax system. The denominator is the welfare loss from requiring the average capital tax rate to be 0.025 as opposed to the optimal tax system. Both welfare losses are expressed in terms of the percentage of consumption agents in the better regime are willing to give up to avoid being in the worse regime.

government cannot impose taxes on capital accumulation. In this last tax system, the average tax rate is zero, but there is no dependence of capital taxes on individual labor income.

In Figure 2, I graph the ratio of the welfare difference between regimes 1 and 3 to the welfare difference between regimes 1 and 2. (I measure the welfare difference as the percentage loss of consumption an agent in regime 1 would be willing to incur to avoid being in the other regime.) Figure 2 has two basic lessons. First, the welfare loss associated with income independence can be much larger than the welfare loss associated with positive capital taxes. The second lesson is that income independence is much less costly if skill shocks are highly persistent. This latter result is hardly surprising. If skills are fixed over time, it is optimal for capital taxes to be zero for all agents. Hence, the

welfare loss of imposing zero taxes is relatively small when the autocorrelation of the shocks is near 1.

The conclusions depicted in Figure 2 are tentative. The welfare losses associated with either restriction on the tax system are likely to be quite different in model economies with more periods (or with more carefully calibrated individual skills processes). Nonetheless, I believe that this picture indicates that having capital taxes depend on labor income can be a first order consideration in designing good social insurance systems.

Lessons from the Simulations

There are three conclusions to be drawn from this numerical exercise. First, the cross-sectional standard deviation of optimal capital tax rates is a decreasing function of the persistence of the underlying skill shocks. Second, the cross-sectional standard deviation of optimal capital tax rates is an increasing function of government purchases. Third, the welfare loss from imposing uniform capital tax rates can be much larger than that from imposing positive capital tax rates. In future work, it would be useful to see the extent to which these lessons can be generalized to more elaborate environments.

7. CONCLUSION

In this paper, I describe a general implementation for the ex ante symmetric Pareto optima in a dynamic Mirrlees economy. The implementation relies on a tax system that is nonlinear in labor income and linear in wealth. As in GKT (2003), it is socially optimal to have social shadow rates of return be higher than individual shadow rates of return. However, it is not possible to implement the optimum by equating this wedge with a tax on wealth. Instead, the tax on wealth accumulated from period t to period $(t+1)$ must be designed to equate the ex post individual after-tax rate of return with the social shadow rate of return. The resulting average wealth tax rate is zero and the government never collects any net tax revenue from wealth taxes.

The government is potentially free to use tax schedules that are nonlinear functions of current and past labor incomes and wealths. This flexibility means that there are many optimal tax systems. Not all of them have the feature that expected marginal tax rates on wealth are zero. In particular, for the case in which skill shocks are i.i.d. over time, Albanesi and Sleet (2005) consider a class of tax systems that are arbitrary nonlinear functions of *current* labor income and *current* wealth. They show that there is a tax system of this kind that implements a constrained Pareto optimal allocation. In this optimal system, current wealth is used as a summary statistic that encodes the relevant information from the past history of labor incomes. Because wealth is playing this record-keeping role, Albanesi and Sleet (2005) prove that the expected marginal tax rate on wealth need not be zero in their optimal system.

In my paper, the government is treated as the sole provider of insurance against skill shocks. It is clear that the results about wealth taxation are sensitive to this assumption. Suppose instead that agents are ex ante identical and can sign long-term contracts with insurance entities (as, for example, in Atkeson and Lucas (1995)). Then the social insurance can be handled by the private sector. There is still a need for taxation—to fund government expenditures—but these taxes optimally take the form of lump-sum levies. There is no need for either labor income taxes or capital income taxes.[11]

Nonetheless, it remains true that much social insurance in highly developed economies is done by the government. I view the analysis in this paper as taking this fact as given and then providing a partial characterization of the nature of optimal dynamic taxation. Understanding why the government plays such a large role in social insurance—using efficiency or other considerations—is an important goal for future research.

Dept. of Economics, Stanford University, Stanford, CA 94305, U.S.A.; Federal Reserve Bank of Minneapolis, and National Bureau of Economic Research; nkocher@stanford.edu.

APPENDIX A: PROOF OF PROPOSITION 1

The proof has two parts. In the first part, I establish that (c^*, y^*, K^*) solves a particular resource minimization problem. In the second part, I derive the first order conditions to that minimization problem.

Part 1

Note first that we can use Lemma 1 of GKT (2003) to show that any optimal allocation satisfies all feasibility constraints with equality.

Next, define $\mu(z_{t+1}|\bar{z}^t)$ to be the conditional probability of z_{t+1}, given $\bar{z}^t$. Consider the following minimization problem *MIN* (I abuse notation slightly by writing $c_t^*(\theta^T, \bar{z}^t)$ to refer to $c_t^*(\theta^T, \bar{z}^t, z_{t+1}, z_{t+2}, \ldots, z_T)$):

$$(21) \qquad \min_{c_t, c_{t+1}, K_{t+1}, \zeta} \int_{\theta^T \in \Theta^T} c_t(\theta^T)\, d\mu_\Theta + K_{t+1}$$

$$(22) \qquad \text{s.t.} \quad u(c_t(\theta^T)) = u(c_t^*(\theta^T, \bar{z}^t)) + \beta \sum_{z_{t+1} \in Z} \zeta(\theta^T, z_{t+1}) \mu(z_{t+1}|\bar{z}^t)$$

$$\text{for almost all } \theta^T \text{ in } \Theta^T,$$

$$(23) \qquad u(c_{t+1}(\theta^T, z_{t+1})) = u\big(c_{t+1}^*(\theta^T, \bar{z}^t, z_{t+1})\big) - \zeta(\theta^T, z_{t+1})$$

$$\text{for all } z_{t+1} \text{ in } Z \text{ and almost all } \theta^T \text{ in } \Theta^T,$$

---

[11]Golosov and Tsyvinski (2005) provide a formal justification of this basic intuition.

$$(24) \qquad \int_{\theta^T \in \Theta^T} c_{t+1}(\theta^T) \, d\mu_\Theta - F_{t+1}\big(K_{t+1}, Y_{t+1}^*(\overline{z}^t), \overline{z}^t\big) - (1-\delta)K_{t+1}$$

$$= -K_{t+2}^*(\overline{z}^t, z_{t+1}) - G_{t+1}(\overline{z}^t, z_{t+1}) \quad \text{for all } z_{t+1} \text{ in } Z,$$

$$(25) \qquad c_t : \Theta^T \to \mathbb{R}_+, \quad c_t \theta^t\text{-measurable},$$

$$(26) \qquad c_{t+1} : \Theta^T \times Z \to \mathbb{R}_+, \quad c_{t+1} \theta^{t+1}\text{-measurable},$$

$$(27) \qquad \zeta : \Theta^T \times Z \to \mathbb{R}, \quad \zeta \theta^t\text{-measurable},$$

$$(28) \qquad K_{t+1} \in \mathbb{R}_+.$$

This minimization problem constructs a class of perturbations around the optimum (c^*, y^*, K^*). The perturbations lower utility in period $(t + 1)$ by a (θ^t, z_{t+1})-contingent amount ζ. This decrease is corrected by raising utility in period t by the expected value of ζ.

I claim that if (c^*, y^*, K^*) is optimal, a solution to *MIN* is

$$(29) \qquad c_t(\theta^T) = c_t^*(\theta^T, \overline{z}^t) \quad \text{a.e.,}$$

$$(30) \qquad c_{t+1}(\theta^T, z_{t+1}) = c_{t+1}^*(\theta^T, \overline{z}^t, z_{t+1}) \quad \text{a.e.,}$$

$$(31) \qquad K_{t+1} = K_{t+1}^*(\overline{z}^t),$$

$$(32) \qquad \zeta(\theta^T, z_{t+1}) = 0 \quad \text{a.e.}$$

Suppose instead that the solution to *MIN* is $(c_t', c_{t+1}', K_{t+1}', \zeta')$. Let B be the Borel subset of Θ^T with measure 1 on which the constraints in *MIN* are valid. Define (c^{**}, K^{**}) by

$$c_t^{**}(\theta^T, \overline{z}^t, z_{t+1}, \ldots, z_T) = c_t'(\theta^T)$$

$$\text{for all } \theta^T \text{ in } B \text{ and all } (z_{t+s})_{s=1}^{T-t} \text{ in } Z^{T-t},$$

$$c_{t+1}^{**}(\theta^T, \overline{z}^t, z_{t+1}, z_{t+2}, \ldots, z_T) = c_{t+1}'(\theta^T, z_{t+1})$$

$$\text{for all } \theta^T \text{ in } B \text{ and all } (z_{t+s})_{s=1}^{T-t} \text{ in } Z^{T-t},$$

$$c_t^{**}(\theta^T, z^T) = c_t^*(\theta^T, z^T) \quad \text{for all other } t, \theta^T, z^T,$$

$$K_t^{**}(\overline{z}^t, z_{t+1}, z_{t+2}, \ldots, z_T) = K_{t+1}' \quad \text{for all } (z_{t+s})_{s=1}^{T-t} \text{ in } Z^{T-t},$$

$$K_t^{**}(z^T) = K_t^*(z^T) \quad \text{for all other } z^T.$$

Obviously, the planner's objective is the same when evaluated at (c^{**}, y^*, K^{**}) as at (c^*, y^*, K^*). Also, (c^{**}, y^*, K^{**}) does not satisfy the period t resource constraint in history $\overline{z}^t$ with equality (because it uses fewer resources than (c^*, y^*, K^*)).

The crux of the proof is to show that (c^{**}, y^*) is incentive-compatible. Let σ be an arbitrary reporting strategy. Then

$$(33) \qquad W(\sigma; c^{**}, y^*) - W(\sigma; c^*, y^*)$$

$$(34) \qquad = \int_{Z^T} \int_{\Theta^T} \sum_{t=1}^{T} \beta^{t-1} \{u(c_t^{**}(\sigma)) - u(c_t^*(\sigma))\}$$

$$(35) \qquad = \mu(\overline{z}^t)\bigg\{\beta^t \int_B \sum_{z_{t+1} \in Z} \mu(z_{t+1}|\overline{z}^t)\beta\zeta'\big(\sigma_\theta(\theta^T, \overline{z}^t), z_{t+1}\big)\, d\mu_\Theta$$

$$\qquad \qquad - \beta^{t+1} \int_B \sum_{z_{t+1} \in Z} \mu(z_{t+1}|\overline{z}^t)\zeta'\big(\sigma_\theta(\theta^T, \overline{z}^t), z_{t+1}\big)\, d\mu_\Theta\bigg\}$$

$$(36) \qquad = 0,$$

where $\sigma_\theta(\theta^T, z^t)$ is the θ component of $\sigma(\theta^T, z^t)$. It follows that

$$(37) \qquad W(\sigma; c^{**}, y^*) - W(\sigma_{\mathrm{TT}}; c^{**}, y^*)$$

$$(38) \qquad = W(\sigma; c^*, y^*) - W(\sigma_{\mathrm{TT}}; c^*, y^*)$$

$$(39) \qquad \leq 0.$$

Thus, if (c^*, y^*) is incentive-compatible, so is (c^{**}, y^*). Because (c^{**}, y^*) uses fewer resources, this completes the first part of the proof.

Part 2

In this part of the proof, I derive first-order necessary conditions to *MIN*. The basic approach is like GKT (2003). The constraint set of *MIN* is a subset of essentially bounded random variables over $\Theta^T \times Z$. Let L_t^∞ be the set of essentially bounded random variables over Θ^T that are θ^t-measurable. As in GKT (2003), it can be shown that the optimum is a regular point of the constraint set. Hence, from Luenberger (1969, Theorem 1, p. 243), we can conclude that the optimum must satisfy the necessary conditions for *MIN*, which are

$$(40) \qquad \sum_{z_{t+1} \in Z} \big((1 - \delta) - F_K(K_{t+1}^*, Y_{t+1}^*)\big)\gamma_{t+1}^*(z_{t+1}) = 1,$$

$$(41) \qquad \int \eta_t\, d\mu_\Theta - \langle m_1^*, u'(c_t^*)\eta_t\rangle = 0 \quad \text{for all } \eta_t \text{ in } L_t^\infty,$$

$$(42) \qquad \gamma_{t+1}^*(z_{t+1}) \int \varepsilon_{t+1}\, d\mu_\Theta - \big\langle m_2^*(z_{t+1}), u'(c_{t+1}^*(\cdot, z_{t+1}))\varepsilon_{t+1}\big\rangle = 0$$

$$\qquad \text{for all } \varepsilon_{t+1} \text{ in } L_{t+1}^\infty \text{ and all } z_{t+1} \text{ in } Z,$$

$$(43) \qquad 0 = \langle \beta m_1^* \mu(z_{t+1}|\overline{z}^t), \nu_t \rangle - \langle m_2^*(z_{t+1}), \nu_t \rangle$$

for all ν_t in L_t^∞ and all z_{t+1} in Z.

Here, m_1^* is an element of the dual of L_t^∞ and is the Lagrange multiplier on the first constraint of *MIN*; for each value of z_{t+1}, $m_2^*(z_{t+1})$ is an element of the dual of L_t^∞ (NOT L_{t+1}^∞) and is the Lagrange multiplier on the second constraint of *MIN*. As is standard, the notation $\langle z, u \rangle$ represents the outcome of applying the linear operator z to u. Finally, $\gamma_{t+1}^*(z_{t+1})$ is a multiplier on the last constraint in *MIN* for each value of z_{t+1}.

We can rewrite the second first-order condition and combine the latter two to get

$$(44) \qquad \int \frac{\eta_t'}{u'(c_t^*)} d\mu_\Theta - \langle m_1^*, \eta_t' \rangle = 0 \quad \text{for all } \eta_t' \text{ in } L_t^\infty,$$

$$(45) \qquad \gamma_{t+1}^*(z_{t+1}) \int \frac{\nu_t'}{u'(c_{t+1}^*(\cdot, z_{t+1}))} d\mu_\Theta = \langle \beta m_1^* \mu(z_{t+1}|\overline{z}^t), \nu_t' \rangle$$

for all ν_t' in L_t^∞ and all z_{t+1} in Z.

Together, these imply that

$$(46) \qquad \beta \mu(z_{t+1}|\overline{z}^t) \int \frac{\eta_t'}{u'(c_t^*)} d\mu_\Theta$$

$$= \gamma_{t+1}^*(z_{t+1}) \int \frac{\eta_t'}{u'(c_{t+1}^*(\cdot, z_{t+1}))} d\mu_\Theta \quad \text{for all } \eta_t' \text{ in } L_t^\infty.$$

By plugging in $\eta_t' = \mathbb{1}_A$, where A is an arbitrary Borel set in Θ^t, and using the standard definition of a conditional expectation, we get

$$(47) \qquad \beta \frac{\mu(z_{t+1}|\overline{z}^t)}{u'(c_t^*)} = \gamma_{t+1}^*(z_{t+1}) E\left(\frac{1}{u'(c_{t+1}^*(\cdot, z_{t+1}))} \Big| \theta^t \right).$$

Define

$$(48) \qquad \lambda_{t+1}^*(\overline{z}^t, z_{t+1}) = \frac{\gamma_{t+1}^*(z_{t+1})}{\mu(z_{t+1}|\overline{z}^t)}.$$

Then

$$(49) \qquad \lambda_{t+1}^*(\overline{z}^t, z_{t+1}) = \beta \frac{\{E(u'(c_{t+1}^*)^{-1}|\theta^t, z^{t+1})\}^{-1}}{u'(c_t^*)}$$

and

$$(50) \qquad 1 = E\{\lambda_{t+1}^*(1 - \delta + F_{K,t+1})|z^t\}.$$

This proves the proposition.   *Q.E.D.*

APPENDIX B: EXAMPLE OF A FAILURE OF ASSUMPTION 1

In this appendix, I construct an example environment in which the optimal consumption c^* does not depend on θ solely through y^*.

Let $T = 2$ and $\Theta = \{1, 2, 3\}$. Assume that

$$(51) \qquad \phi_1(1, \theta_2) = 5 + h \quad \text{and} \quad \phi_2(1, \theta_2) = 5 \quad \text{for all } \theta_2,$$

$$(52) \qquad \phi_1(2, \theta_2) = 5 \quad \text{and} \quad \phi_2(2, \theta_2) = 5 \quad \text{for all } \theta_2,$$

$$(53) \qquad \phi_1(3, \theta_2) = 5 \quad \text{and} \quad \phi_2(3, \theta_2) = 4.5 \quad \text{for all } \theta_2.$$

Hence, agents know their skill sequences in period 1 itself. There are three types of agents. Type 1 agents have high skills in period 1 and medium skills in period 2. Type 2 agents have medium skills in both periods. Finally, type 3 agents have medium skills in period 1 and low skills in period 2. Later, I describe how the parameter h is chosen.

I assume that $u(c) = c^{1/2}$, $v(l) = l^2$, and $\beta = 1$. Also, I assume that the depreciation rate $\delta = 1$ and $F(K, Y) = K + Y$. All agents are initially endowed with zero units of capital.

I solved numerically for the optimal allocation of consumption and effective labor. I find that if I choose $h = 0.30087$, I get

$$(54) \qquad c_1(1, \cdot) = c_2(1, \cdot) = 8.900,$$

$$(55) \qquad c_1(2, \cdot) = c_2(2, \cdot) = 8.532,$$

$$(56) \qquad c_1(3, \cdot) = c_2(3, \cdot) = 8.497,$$

$$(57) \qquad y_1(1, \cdot) = 9.419; \quad y_2(1, \cdot) = 8.380,$$

$$(58) \qquad y_1(2, \cdot) = y_2(2, \cdot) = 8.559,$$

$$(59) \qquad y_1(3, \cdot) = 9.419; \quad y_2(3, \cdot) = 7.523.$$

I chose h so that in the efficient allocation, $y_1(1, \cdot) = y_1(3, \cdot)$. Hence, we have an example in which consumption in period 1 is different for types 1 and 3, but effective labor is the same. There is no way to implement this outcome using a tax system that depends only on effective labor.

This example is nongeneric: by perturbing h away from 0.30087, we get an allocation in which consumption is a function of effective labor. I suspect that it is possible to construct similar examples in which Θ is an interval that is more robust to perturbing the parameters of the economy.

REFERENCES

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

ATKESON, A., AND R. E. LUCAS, JR. (1995): "Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance," *Journal of Economic Theory*, 66, 64–88.

ATKINSON, A., AND A. SANDMO (1980): "Welfare Implications of the Taxation of Savings," *Economic Journal*, 90, 529–549.

BASSETTO, M., AND N. KOCHERLAKOTA (2004): "On the Irrelevance of Government Debt when Taxes Are Distortionary," *Journal of Monetary Economics*, 51, 299–304.

CAREY, D., AND J. RABESONA (2002): "Tax Ratios on Labour and Capital Income and on Consumption," Economic Studies 35, OECD.

CHAMLEY, C. (1986): "Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives," *Econometrica*, 54, 607–622.

CHARI, V. V., L. CHRISTIANO, AND P. KEHOE (1994): "Optimal Fiscal Policy in a Business Cycle Model," *Journal of Political Economy*, 102, 617–652.

CHIAPPORI, P.-A., I. MACHO, P. REY, AND B. SALANIE (1994): "Repeated Moral Hazard: The Role of Memory, Commitment, and the Access to Credit Markets Source," *European Economic Review*, 38, 1527–1553.

DOEPKE, M., AND R. TOWNSEND (2005): "Dynamic Mechanism Design with Hidden Income and Hidden Actions," *Journal of Economic Theory*, forthcoming.

FERNANDES, A., AND C. PHELAN (2000): "A Recursive Formulation for Repeated Agency with History Dependence," *Journal of Economic Theory*, 91, 223–247.

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

GOLOSOV, M., AND A. TSYVINSKI (2004): "Designing Optimal Disability Insurance," Working Paper, MIT.

——— (2005): "Optimal Taxation with Endogenous Insurance Markets," Working Paper, MIT.

GUESNERIE, R. (1995): *A Contribution to the Theory of Pure Taxation*. Cambridge, U.K.: Cambridge University Press.

HAMMOND, P. (1987): "Markets as Constraints: Multilateral Incentive Compatibility in Continuum Economies," *Review of Economic Studies*, 54, 399–412.

KOCHERLAKOTA, N. (2004): "Wedges and Taxes," *The American Economic Review*, 1 May 2004, 94, No. 2, 109–113.

LUCAS, R. E., JR. (2003): "Macroeconomic Priorities," *American Economic Review*, 93, 1–14.

LUENBERGER, D. (1969): *Optimization by Vector Space Methods*. New York: John Wiley and Sons.

MEGHIR, C., AND L. PISTAFERRI (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72, 1–32.

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

ROGERSON, W. (1985): "Repeated Moral Hazard," *Econometrica*, 53, 69–76.

STORESLETTEN, K., C. TELMER, AND A. YARON (2001): "How Important Are Idiosyncratic Shocks? Evidence from Labor Supply," *American Economic Review Papers and Proceedings*, 91, 413–417.