

Federal Reserve Bank of Minneapolis
Research Department Staff Report

February 2004

Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation

Narayana R. Kocherlakota*

Stanford University, Federal Reserve Bank of Minneapolis, and NBER

ABSTRACT

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 in an arbitrary fashion. I construct an optimal tax system that is restricted 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 rate in the following period is zero. As well, the government never collects any net revenue from wealth taxes.

*This version fixes some typos in the December 2003 version. The material in this paper previously circulated as part of the manuscript, "A Mirrlees Approach to Dynamic Optimal Taxation: Implications for Wealth Taxes and Asset Prices." Feel free to contact me with comments or questions via email at nkocher@stanford.edu. I acknowledge the support of NSF SES-0076315. This work grew out of my joint paper with Mikhail Golosov and Aleh Tsyvinski and owes a large intellectual debt to them. I thank Barbara McCutcheon and Ivan Werning for helpful comments; the paper has also benefited from comments from participants in seminars at Penn State University, University of Iowa, Federal Reserve Bank of Chicago, University of British Columbia, Federal Reserve Bank of Minneapolis, Stanford University, and Federal Reserve Bank of Cleveland. The views expressed herein are those of the author and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

In this paper, I consider the following question. Suppose a government needs to finance a given stochastic process of purchases using wealth taxes and labor income taxes. What are the properties of the optimal wealth taxes?

There is a great deal of literature that addresses these and related questions using a Ramsey approach: the government is assumed to be able to use only *linear* taxes on wealth and/or labor income. (See Chari and Kehoe (1999) for an excellent survey.) I instead use what I term a Mirrlees approach. Like Mirrlees (1971), I assume that agents differ in skills (that is, labor productivities), and that a given agent is privately informed about his skill. The government's tax code is restricted only by the government's informational limitations. Both the Mirrlees and Ramsey approaches are motivated by the fact that modern societies rarely use lump-sum taxes, but they differ dramatically in the way that they deal with this fact. Under the Ramsey approach, the government *cannot* use lump-sum taxes. Under the Mirrlees approach, the government *chooses* not to use lump-sum taxes.

My analysis builds off a recent paper by Golosov, Kocherlakota, and Tsyvinski (2003) (henceforth, GKT). GKT consider a dynamic economy in which individual skills are private information. Skills are allowed to follow arbitrary stochastic processes; however, preferences are restricted to be additively separable between consumption and leisure. In this setting, GKT provide a partial characterization of Pareto optimal allocations. They show that in all periods, any individual's shadow interest rate is no higher than, and typically strictly less than, the rate of return to capital. In other words, it is Pareto optimal to have a *wedge* between individual shadow interest rates and social shadow interest rates.

GKT's results are about wedges in Pareto optima, not taxes in an economy with decentralized trade. In this paper, I provide a partial characterization of optimal taxes in a

version of GKT's model economy. Unlike GKT, I allow for publicly observable aggregate shocks (including government purchases shocks). As in GKT, agents' preferences are additively separable over time and between consumption and leisure. I adopt the following model of trade. I assume that agents can sell units of effective labor and rent capital to a representative firm, subject to taxes that are allowed to be arbitrarily nonlinear functions of current and past labor income, but are restricted to be linear in wealth.

I construct a class of such tax systems that weakly implement the optimal allocation. The main result in my paper concerns the nature of the wealth taxes in these optimal systems. It would seem natural in this setting to use a tax system in which the wealth tax rate levied in period $(t + 1)$ is equal to the socially optimal wedge between private and social shadow interest rates between period t and period $(t + 1)$. We know from the work of GKT that under such a system, the wealth tax would typically be positive on each person. However, following Golosov and Tsyvinski (2003a) and Albanesi and Sleet (2003), I show that this kind of tax system is suboptimal. The problem is that such a system does not have enough instruments to prevent individuals from doing a *joint* deviation of saving and lying.

I instead design an optimal tax system that uses period $(t + 1)$ wealth taxes that depend on period $(t + 1)$ labor income (as well as prior labor incomes). I find that under the optimal system, an individual's expected wealth tax rate in period $(t + 1)$, conditional on his period t information and on the period $(t + 1)$ history of public shocks, is zero. Individuals who are surprisingly highly skilled in period $(t + 1)$ receive a subsidy that is a linear function of their wealths. Individuals who are surprisingly unskilled in period $(t + 1)$ are taxed on their wealths. Intuitively, society needs income-contingent wealth taxes to deter the joint deviation of an individual's accumulating too much wealth from period t to period $(t + 1)$

and then not working enough in period $(t + 1)$. For this reason, it is optimal to have higher wealth taxes on those who generate less labor income.

The above result has strong implications for aggregate wealth tax revenue. Consider a group of agents in period t who have the same history of individual shocks through period t . Because their histories are the same, all of these agents choose the same level of capital to hold from period t to period $(t + 1)$. In period $(t + 1)$, their wealth tax rates differ, depending on their labor income realization in that period. But, as stated in the above paragraph, the average optimal wealth tax rate across the agents in this group is zero. Since they all have the same capital, it follows that the net wealth tax revenue in this group is zero. By adding up across all such groups, we can conclude that aggregate net wealth tax collections are always zero. The optimal wealth taxes are purely redistributive in every date and state.

This paper is not the first to look at optimal taxes (as opposed to wedges) in a dynamic version of Mirrlees' model. Golosov and Tsyvinski (2003a) construct an optimal tax system in the Diamond and Mirrlees (1978) disability insurance model. Albanesi and Sleet (2003) design an optimal tax system for the case in which individual skills are independently and identically distributed over time. In both of these papers, the optimal tax systems are considerably simpler than the one that I design. Golosov and Tsyvinski construct a system in which agents face an 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). But this simplicity comes at a cost: both papers restrict attention to a narrow class of stochastic processes for skills, and disregard the possibility of aggregate shocks.

Chari, Christiano and Kehoe (CCK) (1994) consider the same question as I do: what

is the structure of optimal capital income taxes when government purchases are stochastic? However, they use 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. They find that the period t conditional expectation of period $(t + 1)$ tax rates is near zero, but the conditional variance is large. These results may seem to be the same as mine.

However, CCK's expectations are over different random variables. Their results imply that the conditional variance of capital income tax rates over *aggregate* states is non-zero. According to my analysis, the aggregate net wealth tax collections are zero. Any uncertainty over optimal wealth tax rates is generated by individual-level risk, not aggregate risk.

This difference in results is hardly surprising, because entirely different forces are at work in the two kinds of analyses. In the Ramsey approach, the goal is to minimize the deadweight loss associated with the distortions generated by the linearity of taxes. (Sufficiently nonlinear taxes are non-distorting, because they are lump-sum.) Under the Mirrlees approach, 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.

The economy lasts for T periods, where T is finite, 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:

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

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 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, and let μ_Z be a probability measure over the subsets of Z^T that assigns positive probability to all elements 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 $\mu_\Theta(B)$.

Any given agent learns the realization of z_t and his own 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 $\theta^t = (\theta_1, \dots, \theta_t)$ and the history of public shocks $z^t = (z_1, \dots, z_t)$. 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:

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

$$\phi_t : \Theta^T \times Z^T \rightarrow (0, \infty)$$

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

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:

$$K : Z^T \rightarrow R_+^{T+1}$$

$$c : \Theta^T \times Z^T \rightarrow R_+^T$$

$$y : \Theta^T \times Z^T \rightarrow [0, \bar{y}]$$

$$K_{t+1} \text{ is } z^t\text{-measurable}$$

$$(c_t, y_t) \text{ is } (\theta^t, z^t)\text{-measurable}$$

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 amount of capital carried over 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 purchasing needs $G_t : Z^T \rightarrow R_+$ 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 :

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

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

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

$$K_1 \leq K_1^*$$

Here, $F_t : R_+^2 \times Z^T \rightarrow R_+$ 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. Note that (C_t, Y_t) are z^t -measurable.

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* $\sigma : \Theta^T \times Z^T \rightarrow \Theta^T \times Z^T$, where σ_t is (θ^t, z^t) -measurable and $\sigma(\theta^T, z^T) = (\theta^T, z^T)$. Let Σ be the set of all possible reporting strategies, and define:

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

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

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:

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

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:

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

subject to (c, y, K) being incentive-feasible. The idea here is that all agents are treated symmetrically. There is an optimal allocation (the constraint set is compact in the product topology and the objective continuous in the same topology)

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 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:

λ_{t+1}^* is z^{t+1} -measurable

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

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

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. In appendix.

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 .

Second, the theorem establishes that this harmonic conditional 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.) But 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 $\theta^t, \theta^{t'}$:

$$u(c_t(\theta^t, z^t) + \eta(\theta^t)) - u(c_t(\theta^{t'}, z^t) + \eta(\theta^{t'})) = 0$$

or, using a first order approximation:

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

for some B . We can solve for B using:

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

so that the marginal gain to the planner is given by:

$$\begin{aligned} & \sum_{\theta^t \in \Theta^t} \mu_{\Theta}(\theta^t) u'(c_t(\theta^t, z^t)) \eta(\theta^t) \\ &= B \\ &= \varepsilon \left[\sum_{\theta^t \in \Theta^t} \mu_{\Theta}(\theta^t) / u'(c_t(\theta^t, z^t)) \right]^{-1} \end{aligned}$$

The shadow value of resources in a history z^t is given by the harmonic mean of marginal utilities, not the arithmetic mean.¹

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

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

with positive probability if $Var(u'(c_{t+1}^*)|z^{t+1}, \theta^t) > 0$. 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. But what does this result say about taxes?

A. A Problem with the Natural Tax/Wedge Connection

It is natural to think of the intertemporal wedge as telling us that optimal capital taxes should be positive. Why? If agents can buy and sell capital in a competitive market subject to a linear tax, they face the following first order condition:

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

¹Note 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. GKT's 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 Pareto optimal.

But it may be suboptimal for τ_{t+1}^k to be (θ^t, z^t) -measurable.² To see this, consider the following example (which is similar to ones described in Albanesi and Sleet (2003) and Golosov and Tsyvinski (2003)). Let $u(c) = \ln(c)$, and $v(l) = l^2/2$. Suppose $T = 2$, $\Theta = \{0, 1\}$, $Z = \{1\}$, $\pi_h = 1/2$, and $F(K, Y) = rK + wY$. As well, suppose $\phi_1(\theta) = 1$, $\phi_2(\theta, z) = \theta$, and $v(l) = l^2/2$. Set $G = 0$. Then, we can re-write the planner's problem as:

$$\begin{aligned} & \max_{c_1, c_{2h}, c_{2l}, y_1, y_{2h}, K_2} \ln(c_1) - y_1^2/2 + \ln(c_{2h})/2 + \ln(c_{2l})/2 - y_{2h}^2/4 \\ & \text{s.t. } c_1 + K_2 = rK_1 + wy_1 \\ & c_{2h}/2 + c_{2l}/2 = rK_2 + wy_{2h}/2 \\ & \ln(c_{2h}) - y_{2h}^2/2 \geq \ln(c_{2l}) \\ & c_{2h}, c_{2l}, y_{2h}, K_2, y_1 \geq 0 \end{aligned}$$

The solution to this problem must satisfy the following first order conditions:

$$\begin{aligned} c_1^* + K_2^* &= rK_1 + wy_1^* \\ c_{2h}^*/2 + c_{2l}^*/2 &= rK_2^* + wy_{2h}^*/2 \\ \ln(c_{2h}^*) - y_{2h}^2/4 &= \ln(c_{2l}^*) \\ 1/c_1^* &= r/[0.5c_{2h}^* + 0.5c_{2l}^*] \\ w/c_{2h}^* &= y_{2h}^* \\ y_1^* &= w/c_1^* \end{aligned}$$

²This argument is also similar to the one used by Chiappori, et al. (1994).

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 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} (which may be negative). If they earn positive income, they get a handout α_{2h} . 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:

$$\begin{aligned} & \max_{c_1, y_1, c_{2h}, c_{2l}, y_{2h}, k_2} \ln(c_1) - y_1^2/2 + \ln(c_{2h})/2 + \ln(c_{2l})/2 - y_{2h}^2/4 \\ & s.t. \quad c_1 + k_2 = rk_1 + wy_1 \\ & \quad c_{2h} = r(1 - \tau_k)k_2 + wy_{2h} + \alpha_{2h}, y_{2h} > 0 \\ & \quad c_{2h} = r(1 - \tau_k)k_2 + \alpha_{2l} \text{ if } y_{2h} = 0 \\ & \quad c_{2l} = r(1 - \tau_k)k_2 + \alpha_{2l} \\ & \quad k_2, c_{2h}, c_{2l}, y_{2h}, y_1 \geq 0 \end{aligned}$$

and markets clear:

$$\begin{aligned} c_1 + k_2 &= rk_1 + wy_1 \\ c_{2h}/2 + c_{2l}/2 &= rk_2 + wy_{2h}/2 \end{aligned}$$

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:

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

$$y_1 = w/c_1$$

$$w/c_{2h} = y_{2h}$$

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

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

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

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:

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

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

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

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

But there's 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:

$$1/c_1^* \leq r(1 - \tau_k)/c_{2l}^*$$

because

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

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

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:

$$\begin{aligned} & \ln(c_1^*) + \ln(c_{2h}^*)/2 - y_{2h}^{*2}/2 + \ln(c_{2l}^*)/2 \\ &= \ln(c_1^*) + \ln(c_{2l}^*) \end{aligned}$$

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

B. Fixing the Problem

The above problem came 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:

$$\begin{aligned} \max_{c_1, y_1, c_{2h}, c_{2l}, y_{2h}, k_2} \quad & \ln(c_1) - y_1^2/2 + \ln(c_{2h})/2 + \ln(c_{2l})/2 - y_{2h}^2/4 \\ \text{s.t.} \quad & c_1 + k_2 = rk_1 + wy_1 \\ & c_{2h} = r(1 - \tau_{kh})k_2 + wy_{2h} + \alpha_{2h}, y_{2h} > 0 \\ & c_{2h} = r(1 - \tau_{kl})k_2 + \alpha_{2l} \text{ if } y_{2h} = 0 \\ & c_{2l} = r(1 - \tau_{kl})k_2 + \alpha_{2l} \\ & k_2, c_{2h}, c_{2l}, y_{2h}, y_1 \geq 0 \end{aligned}$$

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

$$\begin{aligned} (1 - \tau_{kl})r/c_{2l}^* &= 1/c_1^* \\ (1 - \tau_{kh})r/c_{2h}^* &= 1/c_1^* \\ \alpha_{2i} &= c_{2i}^* - r(1 - \tau_{ki}^*)k_2^*, i = h, l \end{aligned}$$

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:

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

$$w/c_1 = y_1$$

$$c_1 + k_2 = rk_1 + wy_1$$

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

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

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:

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

$$w/c_1 = y_1$$

$$c_1 + k_2 = rk_1 + wy_1$$

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

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

Again, the starred allocations satisfy these first order conditions. The agent is indifferent between working y_{2h}^* in period 2 (when abled) 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 don't 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

I now use the above two-period analysis to build a general implementation of a solution to the planner's problem. I make one assumption: c^* depends on (θ^T, z^T) only through $(y^*(\theta^T, z^T), z^T)$. This assumption allows me to implement the optimal allocation using a tax schedule that is written in terms of effective labor, not in terms of θ^t .

To this end, let (c^*, y^*, K^*) be an optimal allocation. Let:

$$DOM = \{(y^T, z^T) \in R_+^T \times Z^T \mid y^T = y^*(\theta^T, z^T) \text{ for some } \theta^T \text{ in } \Theta^T\}$$

I assume that there exists a function $\tilde{c}^* : DOM \rightarrow \Theta^T \times Z^T$ where \tilde{c}_t^* is (y^t, z^t) -measurable and:

$$\tilde{c}^*(y^*(\theta^T, z^T), z^T) = c^*(\theta^T, z^T)$$

for all (θ^T, z^T) . This assumption guarantees that an agent's consumption depends on his private information only through his history of effective labor. It is an assumption about endogenous variables, and I do not know how to map it into an assumption about the model's exogenous elements.

This assumption seems somewhat innocuous. After all, how could one get effective labor supply to vary across two different types without offering different levels of consumption

to them? Indeed, in a static setting, the assumption is a trivial consequence of incentive-compatibility. But in a dynamic setting, it is possible to construct examples of environments in which the assumption is not satisfied. I present one in Appendix B.³

A. Elements of the Implementation

Given the optimal allocation (c^*, y^*, K^*) , we know from Proposition 1 that there exists $\lambda_{t+1}^* : Z^T \rightarrow R_+$ such that λ_{t+1}^* is z^{t+1} -measurable and:

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

Let $MPK_t^*(z^T) = F_{K_t}(K_t^*(z^T), Y_t^*(z^T), z^T)$ and $MPL_t^*(z^T) = F_{Y_t}(K_t^*(z^T), Y_t^*(z^T), z^T)$ be the marginal product of capital and labor respectively in the optimal allocation.

The key elements of the implementation will be a tax on wealth and a tax on labor income. Define $\tau_{t+1}^* : DOM \rightarrow R_+$ to satisfy:

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

for all (y^T, z^T) in DOM . Verbally, τ_{t+1}^* equates the ex-post individual marginal rate of substitution with the ex-post societal marginal rate of transformation. Note that τ_{t+1}^* is (y^{t+1}, z^{t+1}) measurable. We will use τ^* as a wealth tax.

Finally, let $(\psi^*, \hat{k}^*) : DOM \rightarrow R_+^{2T}$, where ψ_t^* and \hat{k}_{t+1}^* are (y^t, z^t) -measurable and satisfy:

$$\begin{aligned} \hat{c}_t^*(y^T, z^T) + \hat{k}_{t+1}^*(y^T, z^T) &= (1 - \tau_{t+1}^*(y^T, z^T))(1 - \delta + MPK_t^*(z^T))\hat{k}_t^*(y^T, z^T) \\ &\quad + MPL_t^*(z^T)y_t - \psi_t^*(y^T, z^T) \end{aligned}$$

³I used an insight of Ivan Werning's to construct this example.

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

$$\widehat{k}_1^* = K_1^*$$

for all t and for all (y^T, z^T) . Here, the idea is that an agent who earns y_t in period t and has history y^{t-1} loses taxes to the government in period t as prescribed by ψ^* . The process \widehat{k}^* describes how much of the capital stock is owned by an agent with labor income history y^T , when the public history is z^T .

B. A Sequential Markets Economy and Its Equilibrium

We use these definitions to construct a sequential markets economy with nonlinear taxes in which there exists an equilibrium that implements the optimum. 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. The agents pay wealth taxes τ^* on their undepreciated physical capital holdings; they also pay labor income taxes ψ^* on their labor income $w_t y_t$. They can split their after-tax wealth among consumption and capital for next period.

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

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

$$\begin{aligned}
s.t. c_t(\theta^T, z^T) + k_{t+1}(\theta^T, z^T) &= (1 - \tau_t^*(y(\theta^T, z^T), z^T))(1 - \delta + r_t(z^T))k_t(\theta^T, z^T) \\
&\quad + w_t(z^T)y_t(\theta^T, z^T) - \psi_t^*(y(\theta^T, z^T), z^T) \text{ for all } (\theta^T, z^T)
\end{aligned}$$

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

$(y(\theta^T, z^T), z^T)$ is in DOM for all (θ^T, z^T)

$$k_1 \leq K_1^*$$

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 :

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

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

$$\begin{aligned}
G_t(z^T) &= \int_{\theta^T \in \Theta^T} \tau_t^*(y(\theta^T, z^T), z^T) \{(1 - \delta + r_t(z^T))k_t(\theta^T, z^T)\} d\mu_{\Theta} \\
&\quad + \int_{\theta^T \in \Theta^T} \psi^*(y(\theta^T, z^T), z^T) d\mu_{\Theta}
\end{aligned}$$

C. A Second Welfare Theorem

I now claim that the optimal allocation is an equilibrium allocation. 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 a tax system (ψ^*, ϕ^*, D) , the allocation (c^*, y^*, K^*) is optimal for an agent in the economy. To prove this claim, we need the following proposition.

PROPOSITION 2. *Given prices (r^*, w^*) and a tax system (ψ^*, ϕ^*, D) , and given that the typical agent chooses 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. Define $\tau' : \Theta^T \times Z^T \rightarrow R$ 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:

$$\begin{aligned} & \max_{c, K} \sum_{t=1}^T \beta^{t-1} \int_{\theta^T \in \Theta^T} \int_{z^T \in Z^T} u(c_t) d\mu_{\Theta} d\mu_Z \\ & s.t. c_t + k_{t+1} = (1 - \tau_t^*)(1 - \delta + r_t)k_t + w_t y'_t - \psi'_t \\ & c_t, k_{t+1} \text{ are } (\theta^t, z^t)\text{-measurable and non-negative} \\ & k_1 \leq K_1^* \end{aligned}$$

The first order conditions to this problem are:

$$\beta E\{(1 - \tau'_{t+1})u'(c_{t+1})(1 - \delta + r_{t+1})|\theta^t, z^t\} = u'(c_t)$$

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

The first order conditions are necessary and sufficient (once y' is fixed).

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 D :

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

and so for all (θ^T, z^T) :

$$\begin{aligned} & (1 - \tau'_{t+1}(\theta^T, z^T))\lambda_{t+1}^*(z^T)^{-1} \\ &= (1 - \tau_t^*(y'(\theta^T, z^T), z^T))\lambda_{t+1}^*(z^T)^{-1} \\ &= \beta^{-1}u'(\widehat{c}_t^*(y'(\theta^T, z^T), z^T))/u'(\widehat{c}_{t+1}^*(y'(\theta^T, z^T), z^T)) \\ &= \beta^{-1}u'(c'_t(\theta^T, z^T))/u'(c'_{t+1}(\theta^T, z^T)) \end{aligned}$$

Hence:

$$\begin{aligned} & \beta E\{(1 - \tau'_{t+1})u'(c'_{t+1})(1 - \delta + r_{t+1})|\theta^t, z^t\} - u'(c'_t) \\ &= [E\{\lambda_{t+1}(1 - \delta + r_{t+1})|\theta^t, z^t\} - 1]u'(c'_t) \\ &= 0 \end{aligned}$$

This proves the proposition. QED

Proposition 2 considers an agent who chooses an arbitrary effective labor strategy y' where $(y'(\theta^T, z^T), z^T) \in D$ for all (θ^T, z^T) . Because $(y'(\theta^T, z^T), z^T) \in D$ for all (θ^T, z^T) , there

exists a reporting strategy $\sigma' : \Theta^T \times Z^T \rightarrow \Theta^T \times Z^T$ that satisfies:

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

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) :

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

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 $k^*(\theta^T, z^T) = \hat{k}^*(y^*(\theta^T, z^T), z^T)$. 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. But 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^*, D) . 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

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:

$$\tau_{t+1}^{**}(\theta^T, z^T) = \tau_{t+1}^*(y^*(\theta^T, z^T), z^T) \text{ for all } (\theta^T, z^T)$$

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:

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

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

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.

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 redistributive 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:

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

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.

It is important to note that the labor income taxes ψ^* are indeterminate. There is a large set of labor income tax schedules and individual capital-holdings (ψ^*, \hat{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 only earn labor income 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 \$10000 in period 1, and not tax agents with low

income. Then, we can construct another optimal tax system by taxing high-income agents \$1000 in period 1, and $\$9000(1+r)$ in period 2, while transferring \$9000 in period 1 to low-income agents and then taxing them $\$9000(1+r)$ 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. But individual-capital holdings in equilibrium change (high-income agents hold less capital under the second system, while low-income agents hold more).

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.

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.⁴

6. Conclusion

In this paper, I describe a general implementation for the Pareto optima in a dynamic Mirrlees economy. The implementation relies on a tax system that is nonlinear in labor

⁴Like I do, Golosov and Tsyvinski (2003b) 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. The optimal tax is typically non-zero in their setting. Its sign depends on details of the data generation process for skills.

income and linear in wealth. As in GKT (2003), it is Pareto 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.

In this 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 (1992)). 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.⁵

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.

⁵Golosov and Tsyvinski (2003b) provide a formal justification of this basic intuition.

7. Appendix A

In this appendix, I prove 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.

A. 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 $c_t^*(\theta^T, \bar{z}^t, z_{t+1}, z_{t+2}, \dots, z_T)$.)

$$\begin{aligned} & \min_{c_t, c_{t+1}, K_{t+1}, \zeta} \int_{\theta^T \in \Theta^T} c_t(\theta^T) d\mu_{\Theta} + K_{t+1} \\ & s.t. \ u(c_t(\theta^T)) = u(c_t^*(\theta^T, \bar{z}^t)) + \beta \sum_{z_{t+1} \in Z} \zeta(\theta^T, z_{t+1}) \mu(z_{t+1}|\bar{z}^t) \text{ for almost all } \theta^T \text{ in } \Theta^T \\ & \quad u(c_{t+1}(\theta^T, z_{t+1})) = u(c_{t+1}^*(\theta^T, \bar{z}^t, z_{t+1})) - \zeta(\theta^T, z_{t+1}) \text{ for all } z_{t+1} \text{ in } Z \\ & \quad \text{and almost all } \theta^T \text{ in } \Theta^T \end{aligned}$$

$$\int_{\theta^T \in \Theta^T} c_{t+1}(\theta^T) d\mu_{\Theta} - F_{t+1}(K_{t+1}, Y_{t+1}^*(\bar{z}^T), \bar{z}^T) - (1 - \delta)K_{t+1} = -K_{t+2}^*(\bar{z}^T) - G_{t+1}(\bar{z}^T)$$

$$c_t : \Theta^T \rightarrow R_+, c_t \text{ } \theta^t\text{-measurable}$$

$$c_{t+1} : \Theta^T \times Z \rightarrow R_+, c_{t+1} \text{ } \theta^{t+1}\text{-measurable}$$

$$\zeta : \Theta^T \times Z \rightarrow R, \zeta \text{ } \theta^t\text{-measurable}$$

$$K_{t+1} \in R_+$$

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 increase is corrected by lowering utility in period t by the expected value of ζ .

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

$$c_t(\theta^T) = c_t^*(\theta^T, \bar{z}^t) \text{ a.e.}$$

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

$$K_{t+1} = K_{t+1}^*(\bar{z}^t)$$

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

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:

$$c_t^{**}(\theta^T, \bar{z}^t, z_{t+1}, \dots, z_T) = c'_t(\theta^T) \text{ for all } \theta^T \text{ in } B \text{ and all } (z_{t+s})_{s=1}^{T-t} \text{ in } Z^{T-t}$$

$$c_{t+1}^{**}(\theta^T, \bar{z}^t, z_{t+1}, z_{t+2}, \dots, z_T) = c'_{t+1}(\theta^T, z_{t+1}) \text{ for all } \theta^T \text{ in } B \text{ and all } (z_{t+s})_{s=1}^{T-t} \text{ in } Z^{T-t}$$

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

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

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

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 \bar{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:

$$\begin{aligned}
& W(\sigma; c^{**}, y^*) - W(\sigma; c^*, y^*) \\
&= \int_{Z^T} \int_{\Theta^T} \sum_{t=1}^T \beta^{t-1} \{u(c_t^{**}(\sigma)) - u(c_t^*(\sigma))\} \\
&= \mu(\bar{z}^t) \{ \beta^t \int_B \sum_{z_{t+1} \in Z} \mu(z_{t+1} | \bar{z}^t) \beta \zeta'(\sigma_\theta^t(\theta^T, \bar{z}^t), z_{t+1}) d\mu_\Theta \\
&\quad - \beta^{t+1} \int_B \sum_{z_{t+1} \in Z} \mu(z_{t+1} | \bar{z}^t) \zeta'(\sigma_\theta^t(\theta^T, \bar{z}^t), z_{t+1}) d\mu_\Theta \} \\
&= 0
\end{aligned}$$

where $\sigma^t(\theta^T, z^t) \equiv (\sigma_1(\theta^T, z^t, z_{t+1}, z_{t+2}, \dots, z_T), \dots, \sigma_t(\theta^T, z^t, z_{t+1}, \dots, z_T))$ for arbitrary (z_{t+1}, \dots, z_T) .

It follows that:

$$\begin{aligned}
& W(\sigma; c^{**}, y^*) - W(\sigma_{TT}; c^{**}, y^*) \\
&= W(\sigma; c^*, y^*) - W(\sigma_{TT}; c^*, y^*) \\
&\leq 0
\end{aligned}$$

Thus, if (c^*, y^*) is incentive-compatible, so is (c^{**}, y^*) . This completes the first part of the proof.

B. 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. The necessary conditions for *MIN* then are:

$$\begin{aligned} \sum_{z_{t+1} \in Z} ((1 - \delta) - F_K(K_{t+1}^*, Y_{t+1}^*)) \gamma_{t+1}^*(z_{t+1}) &= 1 \\ \int \eta_t d\mu_\Theta - \langle m_1^*, u'(c_t^*) \eta_t \rangle &= 0 \text{ for all } \eta_t \text{ in } L_t^\infty \\ \gamma_{t+1}^*(z_{t+1}) \int \varepsilon_{t+1} d\mu_\Theta - \langle m_2^*(z_{t+1}), u'(c_{t+1}^*(\cdot, z_{t+1})) \varepsilon_{t+1} \rangle &= 0 \\ &\text{for all } \varepsilon_{t+1} \text{ in } L_{t+1}^\infty \text{ and all } z_{t+1} \text{ in } Z \\ 0 = \langle \beta m_1^* \mu(z_{t+1} | \bar{z}^t), \nu_t \rangle - \langle m_2^*(z_{t+1}), \nu_t \rangle \\ &\text{for all } \nu_t \text{ in } L_t^\infty \text{ and all } z_{t+1} \text{ in } Z \end{aligned}$$

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*. 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:

$$\begin{aligned} \int \eta'_t / u'(c_t^*) d\mu_\Theta - \langle m_1^*, \eta'_t \rangle &= 0 \text{ for all } \eta'_t \text{ in } L_t^\infty \\ \gamma_{t+1}^*(z_{t+1}) \int v'_t / u'(c_{t+1}^*(\cdot, z_{t+1})) d\mu_\Theta &= \langle \beta m_1^* \mu(z_{t+1} | \bar{z}^t), \nu'_t \rangle \\ &\text{for all } \nu'_t \text{ in } L_t^\infty \text{ and all } z_{t+1} \text{ in } Z \end{aligned}$$

Together, these imply that:

$$\beta \mu(z_{t+1} | \bar{z}^t) \int \eta'_t / u'(c_t^*) d\mu_\Theta = \gamma_{t+1}^*(z_{t+1}) \int \eta'_t / u'(c_{t+1}^*(\cdot, z_{t+1})) d\mu_\Theta \text{ for all } \eta'_t \text{ in } L_t^\infty$$

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

$$\beta\mu(z_{t+1}|\bar{z}^t)/u'(c_t^*) = \gamma_{t+1}^*(z_{t+1})E(1/u'(c_{t+1}^*(., z_{t+1}))|\theta^t)$$

Define:

$$\lambda_{t+1}^*(\bar{z}^t, z_{t+1}) = \gamma_{t+1}^*(z_{t+1})/\mu(z_{t+1}|\bar{z}^t)$$

Then:

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

and:

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

This proves the proposition.

8. Appendix B

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:

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

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

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

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 period. 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:

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

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

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

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

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

$$y_1(3, \cdot) = 9.419; y_2(3, \cdot) = 7.523$$

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

References

- [1] ALBANESI, S., and SLEET, C. (2003), “Dynamic Optimal Taxation with Private Information,” Federal Reserve Bank of Minneapolis Discussion Paper 140.
- [2] BASSETTO, M., and KOCHERLAKOTA, N. (2004), “On the Irrelevance of Government Debt when Taxes are Distortionary,” forthcoming, *Journal of Monetary Economics*.
- [3] CHARI, V. V., CHRISTIANO, L., and KEHOE, P. (1994), “Optimal Fiscal Policy in a Business Cycle Model,” *Journal of Political Economy* 102, 617-52.
- [4] 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).
- [5] CHIAPPORI, P.-A., MACHO, I., REY, P., and SALANIE, B. (1994), “Repeated Moral Hazard: The Role of Memory, Commitment, and the Access to Credit Markets Source,” *European Economic Review* 38, 1527-53.
- [6] DIAMOND, P., and MIRRLEES, J. A. (1978), “A Model of Social Insurance with Variable Retirement,” *Journal of Public Economics*, 10, 295-336.
- [7] GOLOSOV, M., KOCHERLAKOTA, N., and TSYVINSKI, A. (2003), “Optimal Indirect and Capital Taxation,” *Review of Economic Studies* 70, 569-587.
- [8] GOLOSOV, M., and TSYVINSKI, A. (2003a), “Designing Optimal Disability Insurance,” University of Minnesota working paper.
- [9] GOLOSOV, M., and TSYVINSKI, A. (2003b), “Optimal Taxation with Endogenous Insurance Markets,” University of Minnesota working paper.

- [10] MIRRLEES, J. (1971), “An Exploration in the Theory of Optimum Income Taxation,” *Review of Economic Studies*, 38, 175-208.
- [11] ROGERSON, W. (1985), “Repeated Moral Hazard,” *Econometrica*, 53, 69-76.
- [12] STORESLETTEN, K., TELMER, C., and YARON, A. (2001), “How Important are Idiosyncratic Shocks? Evidence from Labor Supply,” *American Economic Review Papers and Proceedings* 91, 413-417.