

Federal Reserve Bank of Minneapolis
Research Department Staff Report 293

September 2001

Optimal Indirect and Capital Taxation

Mikhail Golosov*

University of Minnesota
and Federal Reserve Bank of Minneapolis

Narayana Kocherlakota*

University of Minnesota
and Federal Reserve Bank of Minneapolis

Aleh Tsyvinski*

University of Minnesota
and Federal Reserve Bank of Minneapolis

ABSTRACT

In this paper, we consider an environment in which agents' skills are private information, are potentially multi-dimensional, and follow arbitrary stochastic processes. We allow for arbitrary incentive-compatible and physically feasible tax schemes. We prove that it is typically Pareto optimal to have positive capital taxes. As well, we prove that in any given period, it is Pareto optimal to tax consumption goods at a uniform rate.

*Kocherlakota acknowledges the support of NSF SES-0076315. For comments and questions, email nkocher@econ.umn.edu. Versions of this paper were presented at the Minnesota Workshop in Macroeconomic Theory, New York University, UCLA, and UCSD; we thank the seminar participants for their comments. We thank Yan Bai, Marco Bassetto, Florin Bidian, Harold Cole, Larry Jones, Patrick Kehoe, Chris Phelan, Jing Zhang, Rui Zhao, and especially V. V. Chari for their comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

The modern economic analysis of optimal taxation takes two distinct forms. One line of research emphasizes the effects of taxation on capital accumulation (see Chari and Kehoe (1999) for an excellent survey). The basic assumption is that a government faces a dynamic Ramsey problem: it needs to fund a stream of purchases over time using linear taxes on capital and labor income. The hallmark result of this literature is that it is optimal for the government to set capital income tax rates to zero in the long run (Chamley (1986), Judd (1985)).

A second branch of the literature is based on the work of Mirrlees (1971, 1976). Here, the government has access to nonlinear taxation. However, agents have fixed heterogeneous skill levels that are unobservable to others. The goal of taxation in this setting becomes (in part) one of transferring resources from the highly skilled to the less skilled in an efficient way, given that incomes but not skills are observable. An important lesson of this literature is the *uniform commodity taxation theorem* of Atkinson and Stiglitz (1976, 1980). It states that if utility is weakly separable between consumption and leisure, then, despite the presence of the incentive problem, it is socially optimal for all consumption goods to be taxed at the same rate.

In this paper, we re-examine the zero capital taxation and uniform commodity taxation theorems in the context of a large class of dynamic economies. We enlarge the class of economies previously studied in two ways. We allow for multiple types of labor; correspondingly, agents' skills are multi-dimensional. More importantly, we allow skills to evolve stochastically over time. We impose *no* restriction on the evolution of skills except that it must be independent across agents.

Besides enlarging the class of economies, we enlarge the choice set of the taxation authority. We do not restrict attention to linear tax schemes (a la Ramsey) or piecewise differentiable schemes (a la Mirrlees). Instead, we allow the taxation authority to use arbitrary nonlinear tax schemes; in other words, it can achieve any incentive-compatible and physically feasible allocation.

This general class of environments is technically challenging: it features both dynamically evolving private information, and a multiple-dimensional type space. There is no known way to develop a full characterization of the socially optimal allocations in this environment. In particular, we might well obtain misleading answers if we were to simply substitute first-order conditions for the large number of incentive constraints, and then apply Lagrangian methods.¹

In the first part of the paper, we reconsider the zero capital income taxation theorem. We specialize the environment to have only one consumption good. We assume also that utility is additively separable in consumption and leisure. We prove that in a Pareto optimal² allocation, individual consumption satisfies a “reciprocal” intertemporal first order condition of the kind derived by Rogerson (1985a):

$$1/u'(c_t) = (\beta R_{t+1})^{-1} E_t\{1/u'(c_{t+1})\}$$

Here, R_{t+1} is the marginal return to investment, u is the agent’s momentary utility function, and β is the individual discount factor.

¹Rogerson (1985b) provides sufficient conditions for the validity of the first-order approach in a static principal-agent context. However, there are no known generalizations of his conditions in dynamic settings.

²By Pareto optimal, we mean Pareto optimal relative to the set of all allocations that are both incentive-compatible and physically feasible.

This “reciprocal first order condition” has an important consequence. If individual marginal utility $u'(c_{t+1})$ in a Pareto optimum is random from the point of view of period t , then from Jensen’s inequality we know that:

$$u'(c_t) < \beta R_{t+1} E_t u'(c_{t+1}) \tag{1}$$

(The incentive problem means that it is typically efficient for individual consumption to be stochastic: The planner needs to offer more consumption to high skill types to get them to work more.) We prove that (1) implies that if agents trade capital and consumption in a sequence of competitive markets, it is optimal for tax rates on capital income to be positive.³

The intuition behind the inequality (1) (and the associated capital income tax result) is as follows. Suppose society considers increasing investment by lowering an individual’s period t consumption by ε and raising an individual’s period $(t + 1)$ consumption by εR_{t+1} . Doing so has two immediate consequences on social welfare (measured in utils): there is a cost $u'(c_t)\varepsilon$ and a benefit $\beta\varepsilon R_{t+1} E_t u'(c_{t+1})$. However, there is an additional adverse incentive effect. If u is strictly concave, increasing c_{t+1} by εR_{t+1} reduces the correlation between $u(c_{t+1})$ and productivity. This correlation exists to provide incentives; reducing the correlation means that effort and output both fall in period $(t + 1)$.

Thus, lowering consumption in period t and raising consumption in period $(t + 1)$ has

³Actually, the analysis only implies that any optimal tax sequence must be consistent with (1); the analysis leaves indeterminate the actual sequence of taxes necessary to generate (1). In particular, if it is possible to use consumption taxes, any path of consumption and capital taxes consistent with (1) is optimal. Some of these paths may feature negative capital taxes as long as consumption taxes are growing at a sufficiently fast rate.

Of course, this point is hardly unique to our paper. In particular, it applies to the original Chamley-Judd analysis.

an extra adverse effect on incentives. In a social optimum, marginal social costs and the marginal social benefit are equated, which implies that the partial marginal cost $u'(c_t)$ is less than the total marginal benefit $\beta R_{t+1} E_t u'(c_{t+1})$.⁴

We go on to reconsider the uniform commodity taxation theorem. We revert to the general assumption of multiple consumption goods, and assume that utility is weakly separable between consumption and labor. We prove that any Pareto optimal allocation has the property that within a period, the marginal rate of substitution between any two consumption goods, for any agent, equals the marginal rate of transformation between those goods. This result implies that if agents can trade consumption goods in a spot market, all consumption goods should be taxed uniformly.

The idea behind the proof of the uniform commodity taxation theorem is as follows. Because utility is weakly separable, consumption only affects the incentive constraints and the planner's objective function through the amount of sub-utility derived from consumption. Hence, as long as resources are scarce, the planner wants to find a way to deliver these sub-utilities that minimizes the resource cost of doing so. This immediately implies the uniform commodity taxation theorem.

We make two distinct contributions to public finance. The first contribution is that we find a general role for positive capital income taxes in a Pareto optimum.⁵ Here, we find that thinking based only on representative agent models can be misleading. It is the dynamic

⁴See Kocherlakota (1998) and Mulligan and Sala-i-Martin (1999) for a similar intuition in a two-period context.

⁵Aiyagari (1995) argues that positive capital income taxes are optimal in an incomplete markets setting. However, he considers only steady-states, rules out markets in an ad hoc basis, and allows only for linear taxes. In contrast, we consider all possible allocations that are feasible and incentive-compatible in a given environment, and thus allow for all possible taxation schemes.

Garriga (2001) shows that in overlapping generations contexts, the optimal linear tax on capital income may be non-zero.

evolution of idiosyncratic shocks that makes positive capital income taxes optimal.

The second is that we greatly generalize the applicability of the uniform commodity taxation theorem. The standard proof of this result is based on much stronger assumptions. Atkinson and Stiglitz's (1976) argument is made in a setting without capital or informational evolution. Moreover, the argument is made under restrictive assumptions: optimal taxes are differentiable and a first-order approach is valid. Both assumptions are typically satisfied only under highly restrictive conditions. We simplify the proof and thereby greatly broaden the range of environments to which it applies.

The rest of the paper is structured as follows. In the next section, we describe the class of model environments. In Section 3, we demonstrate the optimality of positive capital income taxation. In Section 4, we generalize the uniform commodity taxation theorem. We defer a complete discussion of the related literature until Section 5; the discussion clarifies why we are able to prove our results in such generality. Finally, we conclude in Section 6.

2. Setup

The economy lasts for T periods, where T may be infinity, and has a unit measure of agents. The economy is endowed with K_1^* units of the single capital good. There are J consumption goods, which are produced by capital and labor at N different tasks. The agents have identical preferences. The preferences of a given agent are von Neumann-Morgenstern, with cardinal utility function:

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

where $c_t \in R_+^J$ is the agent's consumption in period t , and $l \in R_+^N$ is the amount of time spent working in period t by the agent at the N different tasks. We assume that U is bounded from above or bounded from below; this guarantees that the utility from any consumption/labor process is well-defined as an element of the extended reals.

The agents' skills at the N different tasks differ across agents and over time. We model this cross-sectional and temporal heterogeneity as follows. Let Θ be a Borel set in R_+^N , and let μ be a probability measure over the Borel sets that are subsets of Θ^T . At the beginning of time, an element θ^T of Θ^T is drawn for each agent according to the measure μ ; the draws are independent across agents. This random vector θ^T is the agent's type; its t -th component θ_t is the agent's skill vector in period t . We assume that a law of large numbers applies: the measure of agents in the population with type θ^T in Borel set B is given by $\mu(B)$.

What makes the information problem dynamic is that a given agent privately learns his θ_t at the beginning of period t and not before. Thus, at the beginning of period t , an agent knows his history θ^t of current and past skill vectors but not his future skill vectors. We represent this information structure formally as follows. Define $P_t : \Theta^T \rightarrow \Theta^t$ to be the projection operator: $P_t(\theta_1, \dots, \theta_T) = (\theta_1, \dots, \theta_t)$. Then, define a σ -algebra $\Omega_t = \{P_t^{-1}(B) | B \subset \Theta^t \text{ is Borel}\}$. An agent's information evolution can then be represented by the sequence $(\Omega_1, \Omega_2, \dots, \Omega_T)$ of σ -algebras.

Notice that this stochastic specification allows for virtually arbitrary dynamic evolution of an agent's skills. For example, the agent's skills could be constant over time (which is the traditional public finance assumption). Alternatively, the skills could follow stationary or nonstationary stochastic processes over time. The only real restriction is that the skill processes are independent across agents.

What is the economic impact of these skill vectors? An agent with type θ_t produces effective labor y_{nt} in task n according to the function:

$$y_{nt} = \theta_{nt} l_{nt}$$

where l_{nt} is the amount of time spent working at task n . Effective labor y_{nt} is observable, but actual labor l_{nt} is not.

Along with the consumption goods, there is an accumulable capital good. We define an allocation in this society to be $(c, y, K) = (c_t, y_t, K_{t+1})_{t=1}^T$ where for all t :

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

$$c_t : \Theta^T \rightarrow R_+^J$$

$$y_t : \Theta^T \rightarrow R_+^N$$

$$(c_t, y_t) \text{ is } \Omega_t\text{-measurable}$$

Here, $y_{nt}(\theta^T)$ is the amount of effective labor at task n produced by a type θ^T agent in period t , $c_{jt}(\theta^T)$ is the amount of the j th consumption good given to a type θ^T agent in period t , and K_{t+1} is the amount of capital carried over from period t into period $(t + 1)$.

Let $G : R_+^{J+2+N} \rightarrow R$ be strictly increasing and continuously differentiable with respect to its first $(J + 1)$ arguments, and strictly decreasing and continuously differentiable with respect to its $(J + 2)$ th argument. This function tells us which vectors of capital input, labor inputs and consumption outputs are technologically available. Specifically, we assume that

the initial endowment of capital is K_1^* , and define an allocation (c, y, K) to be *feasible* if:

$$\begin{aligned} \left(\int c_t d\mu, \int y_t d\mu \right) &\in R_+^{J+N} \text{ for all } t \\ G\left(\int c_t d\mu, K_{t+1}, K_t, \int y_t d\mu \right) &\leq 0 \text{ for all } t \\ K_1 &= K_1^* \end{aligned}$$

The first requirement is that c_t and y_t be integrable for all t .

Because θ^T is unobservable, allocations must respect incentive-compatibility conditions. A *reporting strategy* σ is a mapping from Θ^T into Θ^T such that for all t , σ_t is Ω_t -measurable. Let Σ be the set of all possible reporting strategies, and define:

$$\begin{aligned} W(\cdot; c, y) &: \Sigma \rightarrow R \\ W(\sigma; c, y) &= \sum_{t=1}^T \beta^{t-1} \int U(c_t(\sigma), (y_{nt}(\sigma)/\theta_{nt})_{n=1}^N) d\mu \end{aligned}$$

to be the utility from reporting strategy σ , given an allocation (c, y) . Let σ^* be the truth-telling strategy ($\sigma^*(\theta^T) = \theta^T$ for all θ^T). Then, an allocation (c, y, K) is *incentive-compatible* if:

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

An allocation which is incentive-compatible and feasible is said to be incentive-feasible.⁶

We allow for the possibility that the planner weights agents differently based on their

⁶We restrict attention to direct mechanisms. By the Revelation Principle, this is without loss of generality. As well, we restrict attention to mechanisms in which an individual's consumption and output depend only on his own announcements. This is without loss of generality because there is a continuum of agents with independent shock processes.

initial skill levels. Specifically, let $\chi_1 : \Theta^T \rightarrow R_+$ be Ω_1 -measurable, and suppose that $\int \chi_1 d\mu = 1$. Then, we define the following programming problem, $P1(K_1)$, for an arbitrary level K_1 of initial capital:

$$\begin{aligned}
V^*(K_1) &= \sup_{c,y,K} \sum_{t=1}^T \beta^{t-1} \int U(c_t, (y_{nt}/\theta_{nt})_{n=1}^N) \chi_1 d\mu \\
&\text{s.t. } G\left(\int c_t d\mu, K_{t+1}, K_t, \int y_t d\mu\right) \leq 0 \text{ for all } t \\
&W(\sigma^*; c, y) \geq W(\sigma; c, y) \text{ for all } \sigma \text{ in } \Sigma \\
&K_1 \text{ given} \\
&c_t \geq 0, y_t \geq 0, K_t \geq 0 \text{ for all } t \text{ and almost all } \theta^T
\end{aligned}$$

We say that (c^*, y^*, K^*) solves $P1(K_1)$ if (c^*, y^*, K^*) lies in the constraint set of $P1(K_1)$ and:

$$V^*(K_1) = \sum_{t=1}^T \beta^{t-1} \int U(c_t^*, (y_{nt}^*/\theta_{nt})_{n=1}^N) \chi_1 d\mu$$

In the actual model economy, there are initially K_1^* units of capital. Hence, the planner's problem is to solve $P1(K_1^*)$. We assume throughout that there is a solution to $P1(K_1^*)$ and that $|V^*(K_1^*)| < \infty$. Any solution to $P1(K_1^*)$ is a Pareto optimum.⁷

Note that the planner's maximized objective V^* is weakly increasing. In our analysis, we will often require that V^* is strictly increasing. The following lemma shows that, under a mild regularity condition, V^* is strictly increasing if U is additively separable between consumption and leisure. (In the remainder of the paper, as is standard, we use the terms

⁷Specifically, any solution to $P1(K_1^*)$ is interim Pareto optimal, conditional on the realization of θ_1 . If $\chi = 1$, the solutions to $P1(K_1^*)$ are symmetric ex-ante Pareto optima.

for almost all θ^T and almost everywhere (or a.e.) equivalently.)

LEMMA 1. Let $U(c, l) = u(c) - v(l)$, where u is strictly increasing and continuously differentiable. Suppose that for any (c^*, y^*, K^*) that solves $P1(K_1^*)$, there exists some t and positive scalars c^+, c_+ such that $c^+ \geq c_{jt}^* \geq c_+$ a.e. for all j . Then, $V^*(K_1) < V^*(K_1^*)$ for all $K_1 < K_1^*$.

Proof. In Appendix. ■

The proof of the lemma works as follows. Suppose the planner has not used up all initial capital. Because utility is additively separable, the planner can distribute the extra resources across agents so as to add the same amount of utility to every type. Thus, if initial capital is not exhausted, the planner can construct a welfare-improving incentive-compatible redistribution of the extra resources.

3. Capital Income Taxes

To obtain results about the intertemporal characteristics of optimal taxation, we simplify the model. We set the number of consumption goods $J = 1$, and set:

$$G(C_t, Y_t, K_t, K_{t+1}) = C_t + K_{t+1} - K_t(1 - \delta) - F(K_t, Y_t)$$

where F is strictly increasing and continuously differentiable in its first argument. (These restrictions on J and G do not apply in the next section.) Throughout the section, we assume that the partial derivative U_c exists and is continuous in its first argument over the positive reals. We proceed by first providing a partial characterization of Pareto optima, and then establishing the implications of this characterization for capital income tax rates.

A. Characterizing Pareto Optima

The main result in this section is a restriction on the intertemporal behavior of individual consumption. The result is similar to (but much more general than) that derived by Rogerson (1985a) for optimal contracts in relationships with repeated moral hazard.

We begin by stating the result. We use the notation $E\{.\mid\Omega_t\}$ to denote the conditional expectation operator.

THEOREM 1. *Let $U(c, l) = u(c) - v(l)$. Suppose (c^*, y^*, K^*) solves $P1(K_1^*)$, and that there exist $t < T$ and scalars c^+, c_+ such that $c^+ \geq c_t^*, c_{t+1}^*, K_{t+1}^* \geq c_+ > 0$ a.e. Then:*

$$\beta(1 - \delta + F_K(K_{t+1}^*, \int y_{t+1}^* d\mu)) = E\{u'(c_t^*)/u'(c_{t+1}^*)\mid\Omega_t\}$$

Proof. In Appendix. ■

The proof of the theorem runs roughly as follows. We know from Lemma 1 that any solution (c^*, y^*, K^*) to $P1(K_1^*)$ must also solve a dual problem, in which the planner chooses an incentive-feasible allocation so as to minimize the initial resources necessary to deliver a given ex-ante objective value. This implies, a fortiori, that (c^*, y^*, K^*) must also solve any version of the dual problem which has a strictly smaller constraint set that includes (c^*, y^*, K^*) .

We construct a particular constraint set reduction. First, we fix y^* . Then, we include only the feasible consumption/capital allocations (c', K') such that:

$$\sum_{t=1}^{\infty} \beta^{t-1} u(c'_t(\theta^T)) = \sum_{t=1}^{\infty} u(c_t^*(\theta^T)) \text{ for all } \theta^T$$

Note the absence of any expectation operators: in words, we look at allocations (c', K') that deliver the same *ex-post* lifetime utility as c^* to *all* possible types θ^T . The crux of the proof lies in showing that this is in fact a reduction of the constraint set - that is, in proving that (c', y^*) is incentive-compatible. We can then derive the theorem by using the first-order necessary conditions to this version of the dual with a smaller constraint set.

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

$$u'(c_t^*) = \beta(1 - \delta + F_K(K_{t+1}^*, \int y_{t+1}^* d\mu))u'(c_{t+1}^*)$$

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

This kind of thinking informs the next two corollaries. The first concerns the (typical) case in which $u'(c_t^*)$ is not perfectly predictable.

COROLLARY 1. *Let $U(c, l) = u(c) - v(l)$. Suppose (c^*, y^*, K^*) solves $P1(K_1^*)$, and that there exist $t < T$ and scalars c^+, c_+ such that $c^+ \geq c_t^*, c_{t+1}^*, K_{t+1}^* \geq c_+ > 0$ a.e. Suppose also that $\int [Var(u'(c_{t+1}^*)|\Omega_t)]d\mu > 0$. Then with positive probability:*

$$u'(c_t^*) < \beta(1 - \delta + F_K(K_{t+1}^*, \int y_{t+1}^* d\mu))E\{u'(c_{t+1}^*)|\Omega_t\}$$

Proof. Simply apply Jensen's inequality to the condition in Theorem 1. ■

This corollary says that if $u'(c_{t+1}^*)$ is not predictable given Ω_t , the expected marginal utility of investing in capital is higher than the marginal utility of current consumption. Note that this lack of predictability is to be expected in general because the planner wants to elicit high labor from high skill types.

It is interesting to contrast Corollary 1 with the results concerning optimal linear taxation of capital and labor income in a representative agent economy. Chamley (1986) and Judd (1985) prove for a general specification of u that it is optimal in the long run to eliminate the wedge between expected marginal utility of investing in capital and the marginal utility of current consumption. Indeed, when $u(c) = c^{1-\sigma}/(1-\sigma)$, Chamley proves an even stronger result: it is optimal for the wedge to be zero for all t , not just in the long run. In contrast, we find that for any specification of u , as long as $u'(c_{t+1}^*)$ is not predictable given Ω_t , the wedge in period t should be non-zero.

There are special circumstances in which the inequality in Corollary 1 becomes an equality instead. In particular, if agents have fixed skills over time, then the Pareto optimal allocations display no wedge between the marginal utility of consumption and the expected marginal utility of investment.

COROLLARY 2. *Suppose that $\mu(B) > 0$ only if $\mu(B) = \mu\{\theta^T \in B | \theta_t = \theta_1 \text{ for all } t\}$. Let $U(c, l) = u(c) - v(l)$. Suppose (c^*, y^*, K^*) solves $P1(K_1^*)$, and that there exist $t < T$ and scalars c^+, c_+ such that $c^+ \geq c_t^*, c_{t+1}^*, K_{t+1}^* \geq c_+ > 0$ a.e. Then:*

$$\beta u'(c_{t+1}^*)(1 - \delta + F_K(K_{t+1}^*, \int y_{t+1}^* d\mu))/u'(c_t^*) = 1 \text{ a.e.}$$

This corollary follows from the fact that θ_t is perfectly predictable, given θ_1 . In fact,

using a similar approach as in Theorem 1, we can prove (at least when Θ is finite) that even if preferences are non-separable between consumption and labor, we obtain a version of Chamley-Judd's classic result for this case of fixed skills.

PROPOSITION 1. *Suppose $T = \infty$, Θ is finite, and that $\mu\{\theta^\infty\} > 0$ iff $\theta_t = \theta_1$ for all t . Suppose that $V^*(K_1) < V^*(K_1^*)$ for all $K_1 < K_1^*$. Let a strictly positive allocation (c^*, y^*, K^*) solve $P1(K_1^*)$, and suppose that for all θ_1 , the sequence $\{c_t^*(\theta_1), y_t^*(\theta_1), K_t^*\}_{t=1}^\infty$ converges to a positive limit $(c_{ss}(\theta_1), y_{ss}(\theta_1), K_{ss})$. Then:*

$$\beta^{-1} = 1 + F_K(K_{ss}, \int y_{ss} d\mu) - \delta$$

Proof. We claim that (c^*, K^*) solves the following minimization problem:

$$\begin{aligned} & \min_{c, (K_t)_{t=1}^\infty} K_1 \\ & \text{s.t. } \int c_t d\mu + K_{t+1} = K_t(1 - \delta) + F(K_t, \int y_t^* d\mu) \text{ for all } t \\ & \sum_{t=1}^\infty \beta^{t-1} U(c_t(\theta_1), (\frac{y_{nt}^*(\theta_1)}{\hat{\theta}_{n1}})_{n=1}^N) = \sum_{t=1}^\infty \beta^{t-1} U(c_t^*(\theta_1), (\frac{y_{nt}^*(\theta_1)}{\hat{\theta}_{n1}})_{n=1}^N) \text{ for all } \theta_1, \hat{\theta}_1 \\ & K_t \in R_+, c_t \geq 0 \text{ for all } t \end{aligned}$$

Suppose not. Then, there exists nonnegative (c', K') such that $K'_1 < K_1^*$ and:

$$\begin{aligned} \int c'_t d\mu + K'_{t+1} &= K'_t(1 - \delta) + F(K'_t, \int y_t^* d\mu) \text{ for all } t \\ \sum_{t=1}^\infty \beta^{t-1} U(c'_t(\theta_1), (\frac{y_{nt}^*(\theta_1)}{\hat{\theta}_{n1}})_{n=1}^N) &= \sum_{t=1}^\infty \beta^{t-1} U(c_t^*(\theta_1), (\frac{y_{nt}^*(\theta_1)}{\hat{\theta}_{n1}})_{n=1}^N) \text{ for all } \theta_1, \hat{\theta}_1 \end{aligned}$$

It is clear that (c', y^*, K') is feasible; (c', y^*) is incentive-compatible because we have kept

the utility of all announcement/true type pairs the same. This allocation solves $P1(K_1)$, for $K_1 < K_1^*$, which violates the assumption that V^* is strictly increasing.

Now, we can characterize (c^*, y^*, K^*) using the first order conditions to this problem. Let λ_t be the multiplier on the period t feasibility constraint and let $\gamma(\theta_1, \hat{\theta}_1)$ be the multiplier on the appropriate utility constraint.

Abusing notation slightly, we use $\mu(\theta_1)$ to denote $\mu\{(\theta_1, \theta_1, \theta_1, \dots)\}$. Differentiating with respect to $c_t(\theta_1)$ for any θ_1 , we obtain:

$$\sum_{\hat{\theta}_1} \gamma(\theta_1, \hat{\theta}_1) \beta^{t-1} U_c(c_t^*(\theta_1), (\frac{y_{nt}^*(\theta_1)}{\hat{\theta}_{n1}})_{n=1}^N) = \lambda_t \mu(\theta_1)$$

where U_c is the partial derivative of U with respect to c . Differentiating with respect to K_{t+1} we obtain:

$$\lambda_t = \lambda_{t+1} (1 + F_K(K_{t+1}, \int y_{t+1}^* d\mu) - \delta)$$

The assumption that $(c_t(\theta_1), y_t(\theta_1), K_t)$ converges to a positive limit for all θ_1 guarantees that:

$$\begin{aligned} \lim_{t \rightarrow \infty} \lambda_t / \lambda_{t+1} &= 1/\beta \\ \lim_{t \rightarrow \infty} \lambda_t / \lambda_{t+1} &= (1 + F_K(K_{t+1}, \int y_{t+1}^* d\mu) - \delta) \end{aligned}$$

This implies the proposition. ■

B. Capital Trading and Capital Income Taxes

The above results concern the wedges (or lack thereof) between marginal rates of substitution and transformation in Pareto optima. We now want to translate these results about wedges into results about taxes. Many of the mechanisms that implement the Pareto optimum operate by requiring agents to sign exclusive contracts with a planner or intermediary (see Prescott and Townsend (1984) or Atkeson and Lucas (1992)). In these kinds of mechanisms, the implication for taxes is that agents should face an infinite tax if they engage in any side-trading of capital and consumption.

We instead allow agents to (non-exclusively) trade consumption and capital in a sequence of competitive markets. We prove that in this sequential markets setting, our previous results about wedges translate directly into conclusions about capital income taxes (as long as utility is additively separable). We assume throughout this subsection that F is strictly concave in its first argument, that T is finite, and that Θ is finite. (We believe, though, that the results are robust to relaxing the latter two assumptions.)

We do not address the question of how to design a labor income tax schedule that supports the socially optimal allocation. The obvious construction would involve setting a marginal tax rate for each agent that equates his marginal rate of substitution between consumption and time to his marginal rate of transformation. There are two problems with this approach. The first is that the resultant tax schedule may give rise to a non-convex decision problem for the agent. This means that even though his first order conditions are satisfied by the social optimum, he may not find it optimal to make choices consistent with the social optimum.

The second problem is peculiar to the dynamic setting. It is conceivably optimal for

the planner to condition an agent's second period effective labor on the agent's report θ_1 , but not condition the agent's first period effective labor on that report. This would mean that the tax schedule must be a function of reports, not of effective labor.

For these reasons, it is useful to isolate our questions about optimal capital income taxes from questions about optimal labor income taxes. To do so, we consider a class of *capital-trading* mechanisms that work as follows. In each period, each agent makes a report from the set Θ to a social planner. Based on the history of these reports, each agent receives some amount of consumption as after-tax income and is told what vector of effective labor to provide.

Up until this point, the capital-trading mechanisms are standard direct mechanisms. The difference is that agents need not consume their income processes. Instead, they can exchange capital and consumption, and rent out capital services, in a sequence of competitive markets. In each period, an agent faces a linear tax on his capital rental income; the tax rate may be a function of his history of reports.

The other side of the capital rental market is assumed to be a single representative firm. The firm is also partially centralized, because it is simply endowed with a sequence of effective labor which it cannot alter. However, the firm can freely rent capital from the agents; firm profits are split evenly among the agents in the economy.

Thus, under a capital-trading mechanism, labor and after-tax income are allocated according to a direct mechanism. However, agents are allowed to engage in decentralized trade in capital markets. The only restriction is that they face (possibly report-contingent) tax rates on their capital income.

Formally, a capital-trading mechanism is a specification $(z, y, \tau) = (z_t, y_t, \tau_t)_{t=1}^T$ such

that:

$$z_t : \Theta^T \rightarrow R_+$$

$$y_t : \Theta^T \rightarrow R_+^N$$

$$\tau_t : \Theta^T \rightarrow R$$

$$(z_t, y_t, \tau_{t+1}) \text{ is } \Omega_t\text{-measurable}$$

Here, we interpret z as an after-tax income process, y as an effective labor process, and τ as the tax rate on capital income. Thus, given (z, y, τ) , and a rental rate sequence $r \in R_+^T$, a typical agent, initially endowed with K_1^* units of capital, solves the problem:

$$\max_{(c, k, \sigma)} \int \sum_{t=1}^T \beta^{t-1} U(c_t, (y_{nt}(\sigma)/\theta_{nt})_{n=1}^N) d\mu$$

$$s.t. \ c_t + k_{t+1} \leq k_t(1 + r_t(1 - \tau_t(\sigma)) - \delta) + z_t(\sigma)$$

$$c_t \geq 0, k_t \geq 0, k_1 = K_1^*$$

$$c_t, k_{t+1} \text{ } \Omega_t\text{-measurable}$$

$$\sigma \in \Sigma$$

Note that agents take into account their ability to trade in the sequential capital markets when they are making their reports about their types. Their after-tax incomes and their capital income tax rates depend on their reports.

There is a representative firm which operates every period. Given y , and a rental rate

sequence r , the firm solves the following deterministic maximization problem:

$$\max_{K_t \geq 0} F(K_t, \int y_t d\mu) - r_t K_t$$

in each period. We assume that firm profits are split evenly among the agents, and so are embedded directly into z_t .

Given a capital-trading mechanism (z, y, τ) , (c, k, r, K) is a sequential markets equilibrium if it satisfies three conditions. First, (c, k, σ^*) solves the agent's problem given (z, y, τ, r) . (Recall that σ^* is the truth-telling strategy in Σ .) Second, K solves the firm's problem, given (y, r) . Finally, markets clear in every period:

$$\begin{aligned} \int c_t d\mu + K_{t+1} &= F(K_t, \int y_t d\mu) + (1 - \delta)K_t \\ \int k_t d\mu &= K_t \end{aligned}$$

We now prove two results about capital-trading mechanisms. Both require the assumption that utility is additively separable. The first result is that any incentive-feasible allocation is a sequential markets equilibrium of some capital-trading mechanism. The key to the result is that all agents, regardless of their type, have the same preferences over consumption processes.

PROPOSITION 2. *Let $U(c, l) = u(c) - v(l)$, where $u', -u'' > 0$. Suppose (c^*, y^*, K^*) is incentive-feasible and $(c_t^*, K_{t+1}^*) > 0$ for all t . Then, there exists (k, r, z, τ) such that (c^*, k, r, K^*) is an equilibrium of a capital-trading mechanism (z, y^*, τ) .*

Proof. Given (c^*, y^*, k^*) , define:

$$k_t = K_t^*$$

$$r_t = F_K(K_t^*, \int y_t^* d\mu)$$

$$\tau_{t+1} = 1 - (-1 + \delta + u'(c_t^*) / [\beta E\{u'(c_{t+1}^*) | \Omega_t\}]) / r_{t+1}$$

$$z_t = c_t^* + K_{t+1}^* - K_t^*(1 + r_t(1 - \tau_t) - \delta)$$

K is clearly optimal for the firm given the rental rate sequence r and the (aggregate) effective labor sequence $\int y d\mu$. We need to show that (c^*, k^*) solves the agent's problem given (z, y^*, τ, r) .

To do so, fix any reporting strategy σ . Conditional on this strategy, the agent faces the decision problem:

$$\max_{(c, k)} \sum_{t=1}^T \beta^{t-1} \int u(c_t) d\mu$$

$$s.t. c_t + k_{t+1} = k_t(1 + r_t(1 - \tau_t(\sigma)) - \delta) + z_t(\sigma)$$

$$c_t, k_{t+1} \text{ } \Omega_t\text{-measurable}$$

$$c_t \geq 0, k_t \geq 0, k_1 = K_1^*$$

We claim that the solution to this problem is to set $k_t = K_t^*$ and $c_t = c_t^*$. The choice set is convex. Clearly, these choices satisfy the agent's intertemporal first order conditions. They also satisfy his flow budget constraints because of the definition of $z_t(\sigma)$.

Now, which reporting strategy does the agent use? Conditional on any σ , the agent receives the allocation $(c_t^*(\sigma), y_t^*(\sigma))$. But because (c^*, y^*) is incentive-compatible, it is at

least weakly optimal for the agent to choose σ^* .

Because (c^*, y^*, K^*) is feasible, the sequential markets clear. ■

Proposition 2 demonstrates that when we optimize over incentive-feasible allocations (as in Theorem 1), we are implicitly optimizing over capital-trading mechanisms. The following converse proposition shows that in any sequential markets equilibrium, the sign of the capital income taxes is the same as the sign of the wedge between intertemporal marginal rates of substitution and transformation.

PROPOSITION 3. *Let $U(c, l) = u(c) - v(l)$, where $u', -u'' > 0$. Suppose (c, k, r, K) , $k_t > 0$ for all t , is a sequential markets equilibrium of a capital-trading mechanism (z, y, τ) . Then:*

$$(1 + F_K(K_{t+1}, \int y_{t+1} d\mu)(1 - \tau_{t+1}) - \delta) = u'(c_t) / \beta E\{u'(c_{t+1}) | \Omega_t\}.$$

Proof. Individual optimality and firm optimality imply that:

$$\begin{aligned} r_t &= F_K(K_t, \int y_t d\mu) \\ u'(c_t) &= (1 + r_{t+1}(1 - \tau_{t+1}) - \delta) \beta E\{u'(c_{t+1}) | \Omega_t\} \end{aligned}$$

which in turn implies the proposition. ■

Combining Propositions 2 and 3 with Corollary 1, we conclude that it is typically Pareto optimal for capital income taxes to be positive.

4. Uniform Commodity Taxation

In this section, we prove the uniform commodity taxation theorem. We return to the general setup described in the first section (with multiple commodities and a general production structure), except that we assume that utility is weakly separable:

$$U(c, l) = V(u(c), l), \quad u : R_+^J \rightarrow R_+$$

We also assume that u is strictly increasing and is continuously differentiable over the positive orthant of R^J . The notation u_j and G_j represents the partial derivatives of those functions with respect to their j th arguments.

THEOREM 2. *Suppose $V^*(K_1) < V^*(K_1^*)$ for all $K_1 < K_1^*$. Let (c^*, y^*, K^*) solve $P1(K_1^*)$ and suppose that there exist some t and scalars c^+, c_+ such that $c^+ > c_{jt}^*(\theta^T) > c_+ > 0$ for all j and for almost all θ^T . Then, if $J > 1$,*

$$\begin{aligned} & u_j(c_t^*(\theta^T))/u_k(c_t^*(\theta^T)) \\ &= G_j\left(\int c_t^* d\mu, K_{t+1}^*, K_t^*, \int y_t^* d\mu\right)/G_k\left(\int c_t^* d\mu, K_{t+1}^*, K_t^*, \int y_t^* d\mu\right) \end{aligned}$$

for all j, k and almost all θ^T .

Proof. In Appendix. ■

Thus, in a Pareto optimum, the marginal rate of substitution between two consumption goods is equalized to the marginal rate of transformation between those two goods. The key to the proof is that the consumption goods enter both sides of the incentive constraints only through the sub-utility $u(c)$. Hence, it is optimal for the planner to deliver this sub-utility

from consumption in a way that minimizes the resource cost of doing so.

Theorem 2 establishes a result about marginal rates of substitution and transformation. However, we can follow the line of attack in Section 3B to translate it into a statement about taxes. In particular, suppose agents can trade consumption goods in a competitive spot market in each period. Then, Theorem 2 implies that it is suboptimal for them to face taxes or subsidies in those markets that differ across consumption goods.

5. Related Literature

A key property of the model is that the typical agent's willingness to substitute between consumption goods (within a period or over time) is public information. This aspect of the model implies that it is useful to divide the prior literature into two groups of papers. The first group of papers analyze models in which agents have private information about their willingness to substitute between consumption goods (over time or within a period). We show below that our results do not extend into models of this kind. In contrast, the second group of papers is like ours: it analyzes models in which the agents' willingness to substitute between consumption goods is common knowledge. Our results can be viewed as (considerable) generalizations of those in this literature.

A. Privately Known Intertemporal MRS

There are now many papers on efficient dynamic insurance in the presence of hidden idiosyncratic shocks to endowments or marginal utilities of consumption (see, among others, Townsend (1982), Green (1987), Thomas and Worrall (1990), Atkeson and Lucas (1992), Khan and Ravikumar (2001)). These kinds of shocks mean that a typical agent is privately informed about his marginal rate of substitution between period t consumption and period

$(t + 1)$ consumption.

A key result that runs through this dynamic insurance literature is that in Pareto optimal allocations, the typical agent's shadow interest rate is no larger than the societal shadow interest rate. This result is similar to our Corollary 1.

However, unlike our Corollary 1, the result from the dynamic insurance literature depends crucially on the nature of the shock process to endowments or tastes. To see this point, consider a two-period economy with a continuum of agents who have a utility function:

$$u(c_1) + u(c_2)$$

over sequences of consumption. The typical agent's endowment is $((1 + \theta), (1 + \theta)^2)$, where θ is random with positive support; the endowments are private information. The society can borrow and lend from an outside lender at a net rate of return r .

In this economy, agents with high first-period endowments have high growth rates of endowments. One can show that in an optimal allocation, agents' shadow interest rates are *higher* than r . Intuitively, with hidden endowments, the direction of the gap depends on whether the agents who need insurance payments are more or less willing to substitute current for future consumption.

In our model, we are able to implement socially optimal allocations using report-contingent taxes (see Section 3B). This approach does not work when agents are privately informed about their intertemporal marginal rates of substitution. To be concrete, again consider a two-period setting with a continuum of agents who have a utility function $u(c_1) + u(c_2)$. The society faces an outside net rate of return r . In period 1, half of the agents receive an

endowment θ_H , and half of the agents receive an endowment θ_L , where $\theta_H > \theta_L$. These first-period endowments are private information. All agents have endowments $\theta_2 = (\theta_H + \theta_L)/2$ in period 2.

In the Pareto optimal allocation, a type i agent receives a consumption stream (c_{i1}, c_{i2}) , where $i = H, L$. These streams must satisfy three conditions:

$$u'(c_{H1})/u'(c_{H2}) = (1 + r)$$

$$u'(c_{L1})/u'(c_{L2}) > (1 + r)$$

$$u(c_{H1}) + u(c_{H2}) = u(\theta_H - \theta_L + c_{L1}) + u(c_{L2})$$

The last equality is that the type H 's incentive constraint is satisfied with equality.

Can we implement a Pareto optimal allocation using a mechanism akin to that in Section 3B? Suppose that an agent who announces i receives a sequence of transfers $(c_{i1} - \theta_i, c_{i2} - \theta_2)$ and can borrow and lend at a rate $r_i = u'(c_{i1})/u'(c_{i2}) - 1$. Given that the agents report truthfully, the borrowing-lending opportunity is constructed so that they will not deviate. However:

$$\begin{aligned} u(c_{H1}) + u(c_{H2}) &= u(\theta_H - \theta_L + c_{L1}) + u(c_{L2}) \\ &< \max_s u(\theta_H - \theta_L + c_{L1} - s) + u(c_{L2} + (1 + r_L)s) \end{aligned}$$

because $u'(\theta_H - \theta_L + c_{L1})/u'(c_{L2}) < (1 + r_L)$. It is no longer optimal for type H 's to tell the truth, once they are allowed to borrow and lend at a type-specific interest rate. So, the allocation cannot be implemented using this kind of borrowing/lending mechanism.

Why doesn't this analysis apply to our framework? In our setup, agents' true types do not affect their willingness to borrow and lend. Hence, if a type i doesn't want to deviate from a consumption scheme by borrowing and lending at rate r_i , then no other type j will either.

B. Publicly Known Intertemporal MRS

As mentioned above, in our paper, agents' intertemporal marginal rates of substitution are publicly known. There are many other papers which also adopt this modelling strategy. For example, Diamond and Mirrlees (1978, 1986) consider a special case of our general setup. In their model, agents are long-lived and can be disabled or not. Disabled agents are unproductive; able agents have known productivities. Once disabled, the agent stays disabled; the probability of an able agent becoming disabled is exogenous. The informational problem is that the disability status of the agent is known only to the agent. Diamond and Mirrlees prove that in the social optimum, the shadow societal interest rate is higher than the private shadow interest rate. They argue explicitly that this result implies that capital income taxation is socially optimal. Our contribution over their work is that we generalize their positive capital income taxation result to a much larger class of individual skills processes.

There are several papers on the properties of efficient allocations in the presence of repeated moral hazard (see, among others, Rogerson (1985a), Phelan and Townsend (1991), Phelan (1994)). Again, in these settings the optimal allocations have the property that agents' shadow interest rates are higher than the societal shadow interest rate. The intuition behind this result is essentially the same as that behind Corollary 1. However, in this literature, the idiosyncratic output shocks are restricted to be independently and identically distributed; we

instead allow for a much wider range of skill processes.

We were originally motivated to write this paper by the work of da Costa and Werning (2001). They examine optimal monetary policy in two models (a cash-credit good framework and a shopping-time setup) in which agents are privately informed about their fixed skills. In the cash-credit good framework, da Costa and Werning prove that if preferences are weakly separable between consumption and leisure, then the Friedman Rule (zero nominal interest rates) is socially optimal. This is essentially an implication of the uniform commodity taxation theorem, and so we conjecture that this result could be established in our more general setup. They also consider how deviations from weak separability of preferences affect optimal monetary policy.

In a paper written at the same time as ours, but independently, Werning (2001) analyzes the properties of optimal capital income taxes in a model economy with unobservable and heterogeneous fixed skills. Like us (Corollary 2), he finds that it is optimal for capital income taxes to be zero in this setting.

6. Conclusion

In this paper, we consider the problem of optimal taxation when individual skills are unobservable, evolve stochastically over time, and are multi-dimensional. We show that when utility is weakly separable between consumption and leisure, it is optimal to equate the marginal rate of substitution between consumption goods for any agent to the marginal rate of transformation between those goods. It follows that Pareto optimal allocations are consistent with uniform taxation of all consumption goods.

We consider the intertemporal structure of optimal taxation when there is only a sin-

gle consumption good and utility is additively separable between consumption and leisure. In this case, if the optimal allocation requires future consumption to be random given current information, then individuals face distorted consumption paths. We show that these distortions are consistent with the presence of positive capital income taxes.

Given additive separability of preferences between consumption and labor, the uniform commodity taxation theorem is generally valid, but the zero capital income taxation theorem is generally not. The reason for this distinction is that over time, individuals are acquiring information about their types. It is this idiosyncratic uncertainty that generates positive capital income taxes. In particular, if individuals knew their entire sequence of skills in period 1, then we could use exactly the same reasoning as in Theorem 2 (or Corollary 2) to conclude that Pareto optimal allocations are consistent with zero capital income taxation.

We are able to prove the theorems in a highly general setting. We allow for a multi-dimensional specification of skills. Individual skills are independent over a continuum of individuals but follow arbitrary stochastic processes over time. Nonetheless, it is possible to push this generality still further: We can allow any additional private information as long as individuals' willingness to substitute consumption over time is common knowledge. This means, for example, that we could allow agents to secretly accumulate human capital, and thereby endogenize skills.

The paper abstracts from government purchases. This is merely for notational convenience. The results can be easily extended to two kinds of model economies with government purchases. The first is one in which per-capita government purchases are a deterministic stream that the government must fund using taxes. The second is one in which government purchases are a choice variable for the social planner. In both kinds of models, the results

are all valid regardless of how government purchases affect production or enter preferences.

Appendix

In this appendix, we collect the proofs of the main results.

A1. Proof of Lemma 1

Suppose $V^*(K_1) = V^*(K_1^*)$ for some $K_1 < K_1^*$. Let (c^*, y^*, K^*) solve $P1(K_1)$ and also $P1(K_1^*)$. Without loss of generality, assume that c_1^* satisfies the uniform boundedness conditions. Define $c'_{11}(\theta^T, \varepsilon)$ to be the solution to the equation:

$$u(c'_{11}(\theta^T, \varepsilon), (c_{1j}^*(\theta^T))_{j \neq 1}) - u(c_1^*(\theta^T)) = \varepsilon \text{ for all } \theta^T$$

for ε nonnegative. Here, $c'_{11}(\theta^T, \varepsilon)$ is the amount of consumption good 1 that gives a type θ^T agent ε more utility than c_1^* . Clearly, c'_{11} is Ω_1 -measurable with respect to θ^T , and is continuous with respect to ε .

From the mean value theorem, for ε small, we know that:

$$|c'_{11}(\theta^T, \varepsilon) - c_{11}^*(\theta^T)| = \varepsilon / u_1(c'_{11}(\theta^T, \varepsilon'), (c_{1j}^*(\theta^T))_{j \neq 1}), 0 < \varepsilon' < \varepsilon$$

where u_1 is the partial of u with respect to its first argument. From the regularity conditions on c^* , we know that there exists $M > 0$ such that:

$$|c'_{11}(\theta^T, \varepsilon) - c_{11}^*(\theta^T)| < M\varepsilon \text{ for } \varepsilon \text{ small}$$

Hence, for ε small, $c'_{11}(\theta^T, \varepsilon)$ is integrable as a function of θ^T . Moreover, adding ε to initial consumption is feasible for initial capital K_1^* , as long as ε is sufficiently small. That is, for

sufficiently small ε ,

$$G\left(\int c'_1(\theta^T, \varepsilon)d\mu, K_2^*, K_1^*, \int y^* d\mu\right) < 0$$

where $c'_1(\theta^T, \varepsilon) \equiv (c'_{11}(\theta^T, \varepsilon), (c'_{1j}(\theta^T))_{j \neq 1})$. Thus, (c', y^*, K^*) is feasible, given initial capital K_1^* .

For all θ^T ,

$$\begin{aligned} & u(c'_1(\theta^T, \varepsilon)) - v((y_{n1}^*(\theta^T)/\theta_{n1})_{n=1}^N) \\ &= u(c_1^*(\theta^T)) + \varepsilon - v((y_{n1}^*(\theta^T)/\theta_{n1})_{n=1}^N) \\ &\geq u(c_1^*(\theta^{T'}) + \varepsilon - v((y_{n1}^*(\theta^{T'})/\theta_{n1})_{n=1}^N) \\ &= u(c'_1(\theta^{T'}, \varepsilon)) - v((y_{n1}^*(\theta^{T'})/\theta_{n1})_{n=1}^N) \end{aligned}$$

which proves that (c', y^*) is incentive-compatible (the inequality is implied by the incentive-compatibility of (c^*, y^*)). It follows that (c^*, y^*) cannot be a solution to $P1(K_1^*)$. ■

A2. A Technical Lemma

We use the following notation:

$$L^\infty(\Omega_t) = \{x \text{ } \Omega_t\text{-measurable} \mid \exists A \in \Omega_t \text{ such that } \sup_{\theta^T \in A} |x| < \infty, \text{ and } \mu(A) = 1\}$$

Let $\|\cdot\|$ denote the usual ess-sup norm on $L^\infty(\Omega_t)$.

The proofs of Theorems 1 and 2 use two technical results. The first is Theorem 1, p. 243 of Luenberger (1969). This theorem assumes that in an optimization problem with equality constraints, the objective and constraints are continuously Frechet differentiable in

the neighborhood of a local optimum. It then proves that this local optimum must satisfy analogs of the usual Lagrangian first-order conditions.

The second key result is the following lemma. It establishes that as long as c_t^* is bounded from above and below, the constraints in the minimization problems in the proofs of Theorems 1 and 2 are defined by a function that is continuously Frechet differentiable in a neighborhood of c_t^* .

LEMMA 2. Let $u : R_+ \rightarrow R$ be C^1 and let c_t^* be an element of $L^\infty(\Omega_t)$. Suppose there exist scalars c^+ and c_+ such that $c^+ \geq c_t^* \geq c_+ > 0$. Define $U : L^\infty(\Omega_t) \rightarrow L^\infty(\Omega_t)$ by:

$$U(c_t)(\theta^T) = u(c_t(\theta^T))$$

Then U is continuously Frechet differentiable in a neighborhood of c_t^* .

Proof. Note that u' is uniformly continuous over the interval $[c_+/2, 3c^+/2]$. Let $\{\Delta_{nt}\}_{n=1}^\infty$ be an arbitrary sequence in $L^\infty(\Omega_t)$ such that:

$$\lim_{n \rightarrow \infty} \|\Delta_{nt}\| = 0$$

Then:

$$\begin{aligned} & \lim_{n \rightarrow \infty} \|u(c_t^* + \Delta_{nt}) - u(c_t^*) - u'(c_t^*)\Delta_{nt}\|/\|\Delta_{nt}\| \\ &= \lim_{n \rightarrow \infty} \|u'(c_t^* + \Delta'_{nt})\Delta_{nt} - u'(c_t^*)\Delta_{nt}\|/\|\Delta_{nt}\|, 0 \leq \Delta'_{nt} \leq \Delta_{nt} \\ &\leq \lim_{n \rightarrow \infty} \|u'(c_t^* + \Delta'_{nt}) - u'(c_t^*)\|(\|\Delta_{nt}\|/\|\Delta_{nt}\|) \end{aligned}$$

$$\begin{aligned}
&= \lim_{n \rightarrow \infty} \|u'(c_t^* + \Delta'_{nt}) - u'(c_t^*)\| \\
&= 0
\end{aligned}$$

The first step follows from the mean value theorem and the last step from the uniform continuity of u' over $[c_+/2, 3c^+/2]$.

It follows that in a neighborhood of c_t^* , the Frechet derivative of U is well-defined and given by $U'(c_t)(\Delta) = u'(c_t)\Delta$ for all Δ in $L^\infty(\Omega_t)$. The norm of this linear operator is given by $\|u'(c_t)\|$. Let $\|c_t - c_t^*\| < c_+/2$ and let $\{\Delta_{nt}\}_{n=1}^\infty$ be a sequence in L^∞ such that $\lim_{n \rightarrow \infty} \|\Delta_{nt}\| = 0$. Then:

$$\lim_{n \rightarrow \infty} \|u'(c_t + \Delta_{nt}) - u'(c_t)\| = 0$$

because u' is uniformly continuous over $[c_+/2, c^+/2 + c_+/2]$. So U is continuously Frechet differentiable in a neighborhood of c_t^* . ■

We can now turn to the proof of Theorems 1 and 2.

A3. Proof of Theorem 1

The proof has two distinct parts.

Part 1: Constructing a Minimization Problem

In the first part of the proof, we construct a particular class of two-period deviations from the candidate optimum. The class of possible deviations satisfies two requirements. First, the deviations are required to deliver the same utility to all types as does the candidate optimum. Second, the deviations are required to satisfy resource-feasibility in all periods.

Obviously, the first requirement means that all of these deviations provide the same

objective value to the planner. As well, the first requirement implies that all of the deviations are incentive-compatible. Hence, we now have a necessary condition for the candidate optimum: it must use fewer initial resources than any of these possible deviations.

More precisely, consider the following minimization problem *MIN1*:

$$\begin{aligned}
& \min_{\eta_t, \varepsilon_{t+1}, \zeta_t} [\zeta_t + \int \eta_t d\mu] \\
& s.t. \\
& \int \varepsilon_{t+1} d\mu = F(K_{t+1}^* + \zeta_t, \int y_{t+1}^* d\mu) - F(K_{t+1}^*, \int y_{t+1}^* d\mu) + (1 - \delta)\zeta_t \\
& u(c_t^* + \eta_t) + \beta u(c_{t+1}^* + \varepsilon_{t+1}) = u(c_t^*) + \beta u(c_{t+1}^*) \text{ a.e.} \\
& c_t^* + \eta_t \geq 0, c_{t+1}^* + \varepsilon_{t+1} \geq 0, K_{t+1}^* + \zeta_t \geq 0 \text{ a.e.} \\
& \eta_t \in L^\infty(\Omega_t), \varepsilon_{t+1} \in L^\infty(\Omega_{t+1}), \zeta_t \in R
\end{aligned}$$

The objective of this problem is to minimize the resources used in period t . The first constraint requires that feasibility be satisfied in period $(t+1)$. The second constraint requires that utility to all types be kept the same under the deviation plan as under the candidate optimum.

We claim that *MIN1* is solved by setting $(\eta_t, \varepsilon_{t+1}, \zeta_t) = 0$. Suppose not, and that there exists some element $(\eta_t, \varepsilon_{t+1}, \zeta_t)$ of the constraint set which generates a negative value for the objective. There exists a subset B of Θ^T such that $\mu(B) = 1$ and:

$$\begin{aligned}
& u(c_t^*(\theta^T) + \eta_t(\theta^T)) + \beta u(c_{t+1}^*(\theta^T) + \varepsilon_{t+1}(\theta^T)) \\
& = u(c_t^*(\theta^T)) + \beta u(c_{t+1}^*(\theta^T)) \text{ for all } \theta^T \text{ in } B
\end{aligned}$$

Define (c', K') so that $c' = c^*$ and $K' = K^*$ except that:

$$\begin{aligned} c'_t(\theta^T) &= c_t^*(\theta^T) + \eta_t(\theta^T) \text{ for all } \theta^T \text{ in } B \\ c'_{t+1}(\theta^T) &= c_{t+1}^*(\theta^T) + \varepsilon_{t+1}(\theta^T) \text{ for all } \theta^T \text{ in } B \\ K'_{t+1} &= K_{t+1}^* + \zeta_t \end{aligned}$$

We claim that (c', y^*, K') is incentive-feasible, delivers the same value of the planner's objective as (c^*, y^*, K^*) and uses fewer resources. The allocation (c', y^*, K') is obviously feasible because:

$$\begin{aligned} \int c'_t d\mu + K'_{t+1} &= \int c_t^* d\mu + K_{t+1}^* + \zeta_t + \int \eta_t d\mu \\ &< \int c_t^* d\mu + K_{t+1}^* \end{aligned}$$

We next want to show that the allocation (c', y^*, K') is incentive-compatible. By construction:

$$\begin{aligned} &u(c'_t(\theta^T)) + \beta u(c'_{t+1}(\theta^T)) \\ &= u(c_t^*(\theta^T)) + \beta u(c_{t+1}^*(\theta^T)) \text{ for all } \theta^T \end{aligned}$$

(not just θ^T in B). Then, we know that for any σ in Σ and for all θ^T :

$$\begin{aligned} &\sum_{s=1}^T \beta^{s-1} u(c'_s(\sigma(\theta^T))) \\ &= \sum_{s=1}^{t-1} \beta^{s-1} u(c_s^*(\sigma(\theta^T))) + \beta^{t-1} [u(c'_t(\sigma(\theta^T))) + \beta u(c'_{t+1}(\sigma(\theta^T)))] + \sum_{s=t+2}^T \beta^{s-1} u(c_s^*(\sigma(\theta^T))) \\ &= \sum_{s=1}^{t-1} \beta^{s-1} u(c_s^*(\sigma(\theta^T))) + \beta^{t-1} [u(c_t^*(\sigma(\theta^T))) + \beta u(c_{t+1}^*(\sigma(\theta^T)))] + \sum_{s=t+2}^T \beta^{s-1} u(c_s^*(\sigma(\theta^T))) \end{aligned}$$

$$= \sum_{s=1}^T \beta^{s-1} u(c_s^*(\sigma(\theta^T)))$$

This means that for any σ , agents get the same utility from c' as from c^* . It follows that (c', y^*) is incentive-compatible:

$$\begin{aligned} & \int \sum_{t=1}^T \beta^{t-1} [u(c'_t) - v((y_{nt}^*/\theta_{nt})_{n=1}^N)] d\mu \\ = & \int \sum_{t=1}^T \beta^{t-1} [u(c_t^*) - v((y_{nt}^*/\theta_{nt})_{n=1}^N)] d\mu \\ \geq & \int \sum_{t=1}^T \beta^{t-1} [u(c_t^*(\sigma)) - v((y_{nt}^*(\sigma)/\theta_{nt})_{n=1}^N)] d\mu \text{ for any } \sigma \\ = & \int \sum_{t=1}^T \beta^{t-1} [u(c'_t(\sigma)) - v((y_{nt}^*(\sigma)/\theta_{nt})_{n=1}^N)] d\mu \end{aligned}$$

The inequality comes from the fact that (c^*, y^*) is incentive-compatible.

Hence, (c', y^*, K') uses fewer resources, is incentive-compatible, and delivers the same value of the objective to the planner. This violates Lemma 1. We can therefore characterize (c^*, K^*) using the first order conditions of *MIN1*.

Part 2: Deriving the First Order Conditions

The second part of the proof is purely technical: in it, we verify that the theorem's implication is in fact a first-order condition for *MIN1*.

Suppose we enlarge the constraint set by dropping the non-negativity constraints. The non-negative orthant of $L^\infty(\Omega_t)$ has a non-empty interior. Hence, 0 must also be a local minimum of the enlarged minimization problem without the non-negativity constraints.

Note that the Frechet derivative $U'(c_t^*)$ maps $L^\infty(\Omega_t)$ onto $L^\infty(\Omega_t)$. Hence, $(0, 0, 0)$ is a regular point of the constraint set. From Lemma 2 and Luenberger (1969; Theorem 1, page

243), we can conclude that there exist $z_{t+1}^* \in L^{\infty*}(\Omega_{t+1})$ (the dual of $L^\infty(\Omega_{t+1})$) and $\lambda_t^* \in R$ such that 0 is a stationary point of the following Lagrangian.

$$\begin{aligned}
& L(\zeta_t, \eta_t, \varepsilon_{t+1}) \\
&= \zeta_t + \int \eta_t d\mu + \lambda_t^* \left[\int \varepsilon_{t+1} d\mu - (1 - \delta)\zeta_t - F(K_{t+1}^* + \zeta_t, Y_{t+1}^*) \right] \\
&\quad - \langle z_{t+1}^*, u(c_t^* + \eta_t) + \beta u(c_{t+1}^* + \varepsilon_{t+1}) \rangle
\end{aligned}$$

(Here, as is standard, we use the notation $\langle z, u \rangle$ to denote the result of applying a linear operator z to the random variable u .) In other words:

$$\begin{aligned}
& 1 - \lambda_t^*(1 - \delta) - F_K(K_{t+1}^*, Y_{t+1}^*)\lambda_t^* = 0 \\
& \int \eta_t d\mu - \langle z_{t+1}^*, u'(c_t^*)\eta_t \rangle = 0 \text{ for all } \eta_t \text{ in } L^\infty(\Omega_t) \\
& \lambda_t^* \int \varepsilon_{t+1} d\mu - \langle z_{t+1}^*, \beta u'(c_{t+1}^*)\varepsilon_{t+1} \rangle = 0 \text{ for all } \varepsilon_{t+1} \text{ in } L^\infty(\Omega_{t+1})
\end{aligned}$$

It follows that:

$$\begin{aligned}
& \int \eta_t' / u'(c_t^*) d\mu = \langle z_{t+1}^*, \eta_t' \rangle \text{ for all } \eta_t' \text{ in } L^\infty(\Omega_t) \\
& \beta^{-1} \lambda_t^* \int \varepsilon_{t+1}' / u'(c_{t+1}^*) d\mu = \langle z_{t+1}^*, \varepsilon_{t+1}' \rangle \text{ for all } \varepsilon_{t+1}' \text{ in } L^\infty(\Omega_{t+1}) \\
& \lambda_t^* = [1 - \delta + F_K(K_{t+1}^*, Y_{t+1}^*)]^{-1}
\end{aligned}$$

and so:

$$\beta^{-1} [(1 - \delta + F_K(K_{t+1}^*, Y_{t+1}^*))^{-1} \int \eta_t' / u'(c_{t+1}^*) d\mu = \int \eta_t' / u'(c_t^*) d\mu \text{ for all } \eta_t' \text{ in } L^\infty(\Omega_t)$$

Recall that $y = E(x|\Omega_t)$ if y is Ω_t -measurable and $\int x 1_A d\mu = \int y 1_A d\mu$ for all A in Ω_t . Theorem 1 follows. ■

A4. Proof of Theorem 2

We proceed much as in the proof of Theorem 1. Again, we construct a particular class of deviations from the candidate optimum. In particular, we focus on deviant allocations that deliver the same sub-utility in all states as the optimal allocation.

Thus, we claim that c^* solves the following optimization problem *MIN2*:

$$\begin{aligned} \min_c G(\int c_t d\mu, K_{t+1}^*, K_t^*, \int y_t^* d\mu) \\ \text{s.t. } u(c_t) = u(c_t^*) \text{ a.e.} \\ \text{s.t. } c_t \in L^\infty(\Omega_t) \\ \text{s.t. } c_t \geq 0 \text{ a.e.} \end{aligned}$$

Suppose not. Then, there exists a nonnegative c'_t in $L^\infty(\Omega_t)$ such that:

$$G(\int c'_t d\mu, K_{t+1}^*, K_t^*, \int y_t^* d\mu) < 0$$

and $u(c'_t(\theta^T)) = u(c_t^*(\theta^T))$ for all θ^T in $A \subseteq \Theta^T$, where $\mu(A) = 1$. Let $c''_t(\theta^T) = c'_t(\theta^T)$ for all θ^T in A and $c''_t(\theta^T) = c_t^*(\theta^T)$ for all θ^T not in A . Let $c'' = (c''_t, c_{-t}^*)$.

Clearly, (c'', y^*, K^*) is feasible. As in Theorem 1, this allocation is also incentive-compatible because:

$$W(\sigma^*; c'', y^*)$$

$$\begin{aligned}
&= W(\sigma^*; c^*, y^*) \\
&\geq \max_{\sigma \in \Sigma} W(\sigma; c^*, y^*) \\
&= \max_{\sigma \in \Sigma} W(\sigma; c'', y^*)
\end{aligned}$$

Thus, (c'', y^*, K^*) also solves $P1(K_1^*)$. However, because G is strictly increasing in K_{t+1} , and strictly decreasing in K_t , there exists K' such that (c'', y^*, K') solves $P1(K_1)$ for some $K_1 < K_1^*$. But this means that $V^*(K_1) = V^*(K_1^*)$ which is a contradiction.

Thus, c^* solves the above minimization problem. The rest of the proof is simply technical: establishing that the solution to the minimization problem satisfies the first-order conditions in the theorem.

Note that Lemma 2 can easily be extended to the case in which c_t^* is a finite-dimensional random vector. As in the proof of Theorem 1, if we drop the non-negativity constraints from the minimization problem, we know that c_t^* is a local minimum in the resulting problem, and that it is a regular point in the constraint set. From Lemma 2, and Luenberger (1969; Theorem 1, p. 243), we know that there exists $z_t^* \in L^{\infty*}(\Omega_t)$ such that c_t^* is a stationary point of the Lagrangian:

$$L(c_t) = G\left(\int c_t d\mu, K_{t+1}^*, K_t^*, Y_t^*\right) - \langle z_t^*, u(c_t) \rangle$$

In other words:

$$\begin{aligned}
0 &= G_j\left(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*\right) \int \Delta d\mu - \langle z_t^*, u_j(c_t^*) \Delta \rangle \text{ for all } \Delta \text{ in } L^\infty(\Omega_t) \\
0 &= G_k\left(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*\right) \int \Delta d\mu - \langle z_t^*, u_k(c_t^*) \Delta \rangle \text{ for all } \Delta \text{ in } L^\infty(\Omega_t)
\end{aligned}$$

It follows that:

$$0 = G_j(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*) \int \{\Delta'/u_j(c_t^*)\} d\mu - \langle z_t^*, \Delta' \rangle \text{ for all } \Delta' \text{ in } L^\infty(\Omega_t)$$

$$0 = G_k(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*) \int \{\Delta'/u_k(c_t^*)\} d\mu - \langle z_t^*, \Delta' \rangle \text{ for all } \Delta' \text{ in } L^\infty(\Omega_t)$$

and so:

$$0 = \int [G_j(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*)/u_j(c_t^*) - G_k(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*)/u_k(c_t^*)] \Delta' d\mu \text{ for all } \Delta' \text{ in } L^\infty(\Omega_t)$$

The theorem follows by setting:

$$\Delta' = G_j(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*)/u_j(c_t^*) - G_k(\int c_t^* d\mu, K_{t+1}^*, K_t^*, Y_t^*)/u_k(c_t^*)$$

■

References

- [1] Aiyagari, S. R., 1995, Optimal capital income taxation with incomplete markets, borrowing constraints, and constant discounting, *Journal of Political Economy* 103, 1158-1175.
- [2] Atkeson, A., and Lucas, R. E., Jr., 1992, On efficient distribution with private information, *Review of Economic Studies* 59, 427-53.
- [3] Atkinson, A., and Stiglitz, J.E., 1976, The design of tax structure: Direct versus indirect taxation, *Journal of Public Economics* 6, 55-75.
- [4] Atkinson, A., and Stiglitz, J.E., 1980, *Lectures on public economics*, New York: McGraw-Hill.
- [5] Chamley, C., 1986, Optimal taxation of capital income in general equilibrium with infinite lives, *Econometrica* 54, 607-622.
- [6] Chari, V. V., and Kehoe, P., 1999, Optimal fiscal and monetary policy, in *Handbook of Macroeconomics*, ed. Taylor, J., and Woodford, M., New York: Elsevier.
- [7] da Costa, C., and Werning, I., 2001, On the optimality of the Friedman Rule with heterogeneous agents and non-linear income taxation, University of Chicago manuscript.
- [8] Diamond, P., and Mirrlees, J. A., 1978, A model of social insurance with variable retirement, *Journal of Public Economics* 10, 295-336.
- [9] Diamond, P., and Mirrlees, J. A., 1986, Payroll-tax financed social insurance with variable retirement, *Scandinavian Journal of Economics*, 25-50.
- [10] Garriga, C., 2001, Why are capital taxes high?, Universitat de Barcelona working paper.

- [11] Green, E., 1987, Lending and the smoothing of uninsurable income, in *Contractual Arrangements for Intertemporal Trade*, ed. E. Prescott and N. Wallace, Minneapolis: University of Minnesota Press, 3-25.
- [12] Judd, K., 1985, Redistributive taxation in a simple perfect foresight model, *Journal of Public Economics* 28, 59-83.
- [13] Khan, A., and Ravikumar, B., 2001, Growth and risksharing with private information, forthcoming, *Journal of Monetary Economics*.
- [14] Kocherlakota, N., 1998, The effects of moral hazard on asset prices when financial markets are complete, *Journal of Monetary Economics* 41, 39-56.
- [15] Luenberger, D., 1969, *Optimization by vector space methods*, New York: John Wiley and Sons.
- [16] Mirrlees, J., 1971, An exploration in the theory of optimum income taxation, *Review of Economic Studies* 38, 175-208.
- [17] Mirrlees, J., 1976, Optimal tax theory: A synthesis, *Journal of Public Economics* 6, 327-58.
- [18] Mulligan, C., and Sala-i-Martin, X., 1999, Social security in theory and in practice (II): Efficiency theories, narrative theories, and implications for reform, NBER Working Paper 7119.
- [19] Phelan, C., and Townsend, R., 1991, Comparing multi-period information constrained optima, *Review of Economic Studies* 58, 853-881.

- [20] Phelan, C., 1994, Incentives and aggregate shocks, *Review of Economic Studies* 61, 681-700.
- [21] Prescott, E. C., and Townsend, R., 1984, Pareto optima and competitive equilibria with adverse selection and moral hazard, *Econometrica* 52, 21-45.
- [22] Rogerson, W., 1985a, Repeated moral hazard, *Econometrica* 53, 69-76.
- [23] Rogerson, W., 1985b, The first-order approach to principal-agent problems, *Econometrica* 53, 1357-67.
- [24] Thomas, J., and Worrall, T., 1990, Income fluctuation and asymmetric information: An example of a repeated principal-agent problem, *Journal of Economic Theory* 51, 367-90.
- [25] Townsend, R., 1982, Optimal multiperiod contracts and the gain from enduring relationships under private information, *Journal of Political Economy* 90, 1166-85.
- [26] Werning, I., 2001, Optimal dynamic taxation, University of Chicago manuscript.