The Political Economy of Nonlinear Capital Taxation *

Emmanuel Farhi
Harvard
efarhi@harvard.edu

Iván Werning
MIT
iwerning@mit.edu

September 23, 2008 (10:58am)

Abstract

We study efficient nonlinear taxation of labor and capital in a dynamic Mirrleesian model incorporating political economy constraints. Policies are chosen sequentially over time, without commitment, as the outcome of democratic elections. We study the best equilibrium for this dynamic game. Our main result is that the marginal tax on capital income is progressive, in the sense that richer agents face higher marginal tax rates.

1 Introduction

Modern optimal-tax theory is founded on the trade-off between efficiency and redistribution (Mirrlees, 1971). The losses in efficiency from taxation are determined mechanically by the economic environment—preferences, technology and information. In contrast, the desire to redistribute, often modeled by a social welfare function, may implicitly capture the outcome or demands of some political process.

However, if anything, actual policy making not only considers this trade-off, but also constantly reconsiders it: policies chosen at some point, can be reformed or replaced by new ones at any later time. Due to this lack of commitment, the credibility of policies must be judged by projecting their effects into the future. The impact on future wealth inequality is of particular concern. Otherwise, large levels of inequality may create, ex

*This work benefited from useful discussions and comments volunteered by Daron Acemoglu, Philippe Aghion, Alberto Alesina, George-Marios Angeletos, Robert Barro, Marco Battaglini, Roland Benabou, Steve Coate, Michael Golosov, Narayana Kocherlakota, David Laibson, Greg Mankiw, Chris Phelan, Jim Poterba, Aleh Tsyvinski, Pierre Yared, and seminar and conference participants at Chicago, Columbia, Harvard, MIT, Northwestern, Stanford, Duke Macroeconomics Jamboree, NBER Summer Institute, Society of Economic Dynamics (SED) at Boston. All remaining errors are our own.

post, a political demand for reform, towards policies that redistribute wealth.The purpose of this paper is to explore this mechanism and study optimal policy design when the credibility of policies are taken into account.

Our theory blends recent developments in optimal taxation with elements of political economy. We study a dynamic Mirrleesian model where policy is determined sequentially by democratic elections. We allow for the most general non-linear tax schedules for labor and capital income. Our main result shows that progressive capital taxation emerges naturally in this setting.

Our economy is populated by a continuum of agents that are subject to idiosyncratic labor productivity shocks. Productivity is privately observed, precluding the first-best outcome of full insurance. We assume that tax instruments are restricted only by this asymmetry of information, so that, absent political economy considerations, any incentive compatible allocation is implementable. In particular, we allow for nonlinear taxation of labor and capital income. We study two models, the first with a two-period horizon and the second with an infinite horizon.

If tax policy could be chosen once and for all with full commitment, then standard results from the optimal tax literature would apply. In particular, in the two-period version of our model the tax on capital would be zero (Atkinson and Stiglitz, 1976), while in the infinite-horizon version the average tax on capital could be set to zero (Kocherlakota, 2005). In both cases, the tax on capital would not be progressive.

However, in our model, tax policy is not set in stone at the beginning of time. Instead, it is determined sequentially over time, without commitment, through democratic elections. In particular, to model this political process, we adopt a version of the probabilistic voting model (Coughlin, 1986; Lindbeck and Weibull, 1987). Fortunately, this framework remains tractable within our infinite dimensional policy setting.[1,2] As is well known, under appropriate conditions, the outcome of the political process yields policies that maximize a utilitarian social welfare function.

Elections take place repeatedly, selecting policy makers that can only decide taxes and transfers for one term—they cannot bind future policy makers. Thus, the outcome of the political process can be represented as a utilitarian planner lacking commitment. All the results in the paper apply directly to a setting that simply assumes such a time-inconsistent planner, instead of deriving it from the political process as we do here.

The utilitarian representation implies a concern for inequality and a desire to redis-

[1]In contrast, the median voter model requires stronger assumptions on the policy space, such as a one-dimensional representation with single peaked preferences.

[2]Citizen candidate models, introduced by Besley and Coate (1997) and Osborne and Slivinski (1996), appear to be intractable in our policy setting.

tribute. Without commitment, at any point in time, the most tempting deviation is to wipe the slate clean and implement the most extreme redistribution. In particular, this involves an expropriating capital levy. In equilibrium, such a deviation may be prevented in two ways. First, there may be direct costs of reforming tax policy, providing an intermediate form of commitment. We pursue this idea in our two-period model. Second, there may be a concern for the loss of reputation. We pursue this in our infinite horizon economy, which assumes there are no direct costs of reforming tax policy. Reputation there works as follows. Upon observing a deviation, the private sector's expectations may shift, anticipating future governments to behave similarly. This may lead to a bad economic outcome, where agents do not produce to avoid expropriation. This fear of losing reputation may hold back policy makers.

We formalize this as a dynamic game. The reputational mechanism discussed above then corresponds to a trigger-strategy equilibrium, where a deviation is followed by a bad continuation equilibrium. Drawing on Chari and Kehoe (1990), we focus on sustainable plans or policies, a refinement that focuses on symmetric perfect Bayesian equilibria. Sustainable policies need to be credible, ensuring that current and future governments have the incentive to implement them. As shown by Chari and Kehoe, the necessary requirements can be captured by a sequence of constraints that ensure that no policy maker prefers to deviate towards full redistribution, given that such a deviation would be followed by the worst possible continuation equilibrium. The best sustainable equilibrium can be determined as the solution to a social planning problem with these credibility constraints.

As a consequence of these credibility constraints, equilibrium policies deviate from the normative benchmarks provided by the optimal tax literature. Our main result is that capital taxation is progressive in the sense that agents that enjoy higher consumption face higher marginal tax rates on their savings. We show that this feature can be implemented with a tax schedule on wealth that is convex. As for the level, marginal tax rates may be positive over some regions and negative over others. Indeed, in the two period version of the model, the marginal tax rate on capital is always positive at the very top and negative at the very bottom.

The intuition for these results is as follows. The sign and level of the marginal tax rate placed on any agent is determined by the net effect that an extra unit of capital held by this agent has on the credibility constraint. On the one hand, an extra unit of capital in the hands of some particular agent increases the equilibrium value of the utilitarian objective. In fact, it does so according to this agent's marginal utility. On the other hand, more capital also raises the value of a deviating policy towards full redistribution. The sign of the optimal marginal tax depends on the net of these two effects, since this determines whether it is preferable to encourage or discourage savings by any particular agent. For instance, for a very rich agent, with high consumption and low marginal utility, an extra unit of saving has a negligible effect on the equilibrium utilitarian value. However, the extra unit of capital improves non-trivially the value attached to the deviation towards full redistribution. Thus, capital may be positively taxed for rich agents. The reverse may be true for poor enough agents with low consumption and high marginal utility. Capital may be subsidized for these agents.

The same principle explains the progressivity of the marginal tax rate. The value that an extra unit of capital has on the deviation path with full redistribution is independent of who does the extra saving. The difference between this common value of one unit of capital under a deviation, and the value obtained in equilibrium from this extra capital, which equals that agent's marginal utility, is then solely a function of that agent's consumption. Thus, agents with higher consumption face a higher marginal tax on capital.

The progressivity in the taxation of capital reflects an important features of the allocation, that individual consumption is mean reverting. Agents with higher consumption have lower average consumption growth. This requires that they face lower after-tax rates of return, explaining the progressivity in marginal taxes on capital.

It is optimal to have mean reversion in consumption because this makes policies more credible. As Atkeson and Lucas (1992) showed for taste shock model, absent credibility constraints, incentives are most efficiently provided by rewarding and punishing the agent with permanent shifts in consumption. This is optimal because the agent values smoothing consumption. As a result, individual consumption behaves roughly as a random walk.[3] In a cross section, with a continuum of agents, this implies that the distribution over consumption disperses, with inequality growing over time. Indeed, these dynamics may be so drastic that the lead to the so-called "immseration result", where every agent's consumption and welfare eventually converges to zero. This would eventually imply arbitrarily low levels of utilitarian welfare, violating the credibility constraints. In other words, if such a policy were in place, a utilitarian planner would, at some point, deviate towards fully redistribution, regardless of the consequences. The efficient plan is not credible.

Taking credibility into account requires upsetting the random walk ideal. To provide incentives efficiently, while considering the credibility constraint, it is best to front-load rewards and punishments. This imparts a mean-reverting component in individual consumption. In the cross section, this decreases the growth in consumption inequality and

[3] Formally, consumption can be expressed a function of a state variable that is a martingale.

relaxes the credibility constraints. [4]

Although our model is quite simple, we argue that it captures some essential aspects of the policy design problem. In particular, central to our model is the notion that policies should be designed with an eye towards their credibility, and that inequality may be a crucial determinant of the latter. We believe both features are important in modern democratic societies. In our model, progressive taxation of capital emerges to reduce wealth inequality by discouraging the accumulation by the rich and encouraging savings by the poor. Venturing outside our model, perhaps other policies, such as schooling subsidies, that reduce future inequality may also constitute efficient tools to promote credibility.

In this paper we build on our previous research within intergenerational settings. In Farhi and Werning (2007) we studied an endowment economy where altruistic agents face privately observed taste shocks. In this setting, when the welfare of the first generation is maximized, the allocation features immiseration, as in Atkeson and Lucas (1992). Following Phelan (2006) we consider other efficient allocations. In particular, we trace out the Pareto frontier between current and future generations by adding the constraint that the expected welfare of all future generations remain above some *exogenous* level. We showed that immiseration is then overturned, that is, that a non-degenerate invariant distribution for consumption and welfare exists for any such floor. This paper did not consider taxation, because the taste-shock environment without capital was not a natural setting for such a purpose. In Farhi and Werning (2008), we study to a Mirrleesian model with capital and focus on implications for taxation, especially estate taxation. The main result we obtained there is that the optimal marginal estate tax is progressive and negative. That is, intergenerational transfers should be subsidized, but the marginal subsidy should be smaller for larger estates.

The current paper's setup and results build on these two papers, but with important differences. In the present model, there is a single generation of infinitely lived agents. Thus, efficiency unequivocally implies maximizing welfare from the perspective of the initial period. However, due lack of commitment, credible policies must keep future utilitarian welfare in all periods above some level. Unlike our previous normative intergenerational models, this level is now *endogenous* and corresponds to the value attached to deviating towards full redistribution, given that this then triggers the worst equilibrium. In terms of results, our implementation using nonlinear capital taxation is similar to the one for estate taxation in Farhi and Werning (2008). Indeed, the tax schedule shares the progressivity feature in both cases. However, an important difference is that, whereas estate taxes were always negative in our previous work, here we find that positive marginal

[4] In fact, a non-degenerate distribution for consumption and welfare may obtain in the limit.

taxes may be optimal.

Our paper is most closely related to a line of work by Sleet and Yeltekin (2006, 2008a,b). Sleet and Yeltekin (2008b) study an endowment economy with privately observed taste shocks, where policies are chosen through probabilistic voting. They show that the best equilibrium is isomorphic to a planning problem with higher social patience, as in Farhi and Werning (2007). In parallel and independent work, Sleet and Yeltekin (2008a) consider a Mirrlees model with capital as in this paper. They obtain similar results regarding the progressivity of capital taxation. However, some differences in focus remain. In particular, we consider an implementation with nonlinear taxation and our analysis covers cases with persistent productivity shocks. Their paper, instead, considers asset-pricing implications, which we ignore.

We also make contact with a literature on political economy incorporating limited commitment and heterogenous agents. Benhabib and Rustichini (1996) study the link between wealth and investment in a dynamic game where output in every period is split between consumption by two social groups and investment. They focus on the best subgame perfect equilirbium. The most profitable deviations involve one group extracting as much consumption as possible, leaving no resources for investment. They show in examples that this might lead to lower capital accumulation in equilibrium than in the first best. Wether these effects are more pronounced at low or high wealth levels depends on the curvature of the utility and production functions, resulting respectively in growth traps or situations with low growth at high wealth levels. Acemoglu, Golosov and Tsyvinski (2008) and Acemoglu, Golosov and Tsyvinski (2007) study a model where policy is set by a self-interested ruler or dictator who derives utility from private consumption. They focus on the best equilibrium of the game without commitment. The ruler's preferred deviation expropriates all the economy's resources for its own private consumption; thus, higher capital increases the attractiveness of this deviation. As a result, the best equilibrium discourages accumulation, implying a positive marginal tax on capital. In contrast to our main result regarding progressivity, in their setting all agents face the same positive tax rate. In addition, unless the ruler is impatient, these distortions disappear in the long run because promised consumption transfers to the ruler are backloaded in a way that makes the credibility constraint eventually not bind.[5] Bisin and Rampini (2006) study optimal policy in an economy with unobservable endowments, when the planner has

[5] Acemoglu, Golosov and Tsyvinski (2007) considers an extension where the ruler's objective is a weighted average of utilitarian welfare and the utility from its private consumption. This model is closer to ours, although they do not consider the ruler's weight on private consumption to be zero. We conjecture that our main result on progressive capital taxation may obtain for this extension. However, this is not addressed because they only study the aggregate distortions to capital accumulation, not individual ones.

limited commitment. In their model, in contrast to ours, reforms are not associated with any cost. As a result, full redistribution always occurs in the second period. They show that in this context, it is welfare improving to give agents access to anonymous markets that the government cannot monitor.

The normative literature on capital taxation provides an important benchmark for our results. Many normative optimal taxation models prescribe zero capital taxation. One the one hand, Chamley (1986) and Judd (1985) have shown that in Ramsey models, capital taxes should not be used in steady-state to finance government expenditures. On the other hand, in a Mirrlees context, the uniform-taxation result by Atkinson and Stiglitz (1976) shows that optimal capital taxes are zero when is no uncertainty and preferences are separable. A few theoretical papers analyze non-linear capital taxes under commitment. Saez (2002) considers a model where the only source of heterogeneity is initial wealth. In this setting, an initial capital levy that fully redistributes capital is optimal. Saez assumes an exogenous upper bound on the marginal tax rate and characterizes the optimal sequence of piecewise linear capital tax schedules. Benabou (2002) constructs a model with human capital, instead of physical capital, and studies nonlinear taxation of income, within a one-dimensional parametric class. As mentioned above, Farhi and Werning (2008), study the related issue of nonlinear estate taxation.

Two branches of the political economy literature have touched upon the issue of capital taxation. Both strands of literature have rationalized positive tax rates on capital, but have largely ignored the nonlinear taxation of capital.

The first branch revolves around the idea of time inconsistency first introduced by Kydland and Prescott (1977). The typical setup is a Ramsey model with a representative agent and a government which finances a public good using linear taxes. The central idea is that once sunk, capital is inelastic, so that capital taxation is equivalent ex-post to lump sum taxation. See Fischer (1980) and Klein and Rios-Rull (2003) for a more recent treatment. Several papers analyze how reputation mechanisms can alleviate the time inconsistency problem and result in intermediate levels of capital taxation. See for example Kotlikoff, Persson and Svensson (1988), Chari and Kehoe (1990) and Phelan and Stacchetti (2001).

The second branch is closest to our paper. It studies the linkage between income distribution, redistribution and growth, but mostly abstracts from time inconsistency problems.[6] The typical setup features heterogenous agents and linear taxation combined with

[6]One notable exception is Persson and Tabellini (1994b), who reintroduce a time inconsistency problem in an otherwise similar model. They emphasize strategic delegation, whereby voters might elect a policymaker that has a disproportionate stake in capital income.

lump sum rebates. If the median voter is less productive than the mean voter (Persson and Tabellini, 1994a), or if the median voter derives a lesser fraction of its total income from capital than the mean voter (Alesina and Rodrick, 1994; Bertola, 1993), strictly positive and higher than optimal capital tax rates will be chosen in the political equilibrium.

Our model combines elements of both literatures. Time inconsistency arises in our setup because of the interaction between dynamic incentive provision and redistribution. Incentives require inequality in consumption and savings. However, because of a concern for equality, it is tempting to expropriate capital holdings and fully redistribute. The main result of the paper that capital taxes are progressive is a new insight.

The rest of the paper is organized as follows. Sections II and III contain the two period model when policies are determined respectively under commitment or through successive elections with an exogenous cost of reforms. Section IV describes the policy game and the solution concept in the infinite-horizon model. Section V derives the results under the simplifying assumptions that shocks are i.i.d. Section VI tackles the general case of persistent shocks.

2 A Two Period Economy

We begin with a two period version of the model. The analysis in this case is simple and helps bring out the essential mechanism underlying our results. The drawback is that one cannot capture a concern for reputation, where the government's desire to deviate is kept in check to avoid a negative shift in the expectations of the private sector. In an infinite horizon setting this is possible in a trigger-strategy equilibrium. However, in our finite horizon model, the unique equilibrium can be solved by backward induction, implying that the outcome is independent of past policy choices, precluding any concerns for reputation. To avoid this, we assume that current governments can make policy choices that partially bind future governments. In particular, we assume that reforming a previously enacted law requires paying a fixed cost in terms of output. This introduces a state variable which makes the game depend on past policy choices. In the limit where the exogeneous cost is zero, no commitment is possible; if the cost is arbitrarily large, then full commitment is possible. Intermediate values of the cost allow us to capture intermediate levels of commitment.

We later study an infinite horizon model where this fixed cost is unnecessary. Indeed, in a way, the infinite horizon setting endogenizes the cost of reform. We study credible equilibria, closely related to the game theoretic notion of subgame perfect equilibria, and focus on the best equilibrium.

2.1 Preferences and Technology

The economy lasts for two periods $t = 0, 1$ and is populated by a continuum of agents. Agents work, save and consume in the first period, and simply consume in the second. We introduce heterogeneity in the productivity, θ_0, with which they convert effort e_0 into effective units of labor $n_0 = e_0 \cdot \theta_0$. Their lifetime utility is given by

$$u(c_0) - h\left(\tfrac{n_0}{\theta_0}\right) + \beta u(c_1), \tag{1}$$

where c is consumption. We assume that the utility functions u and h are twice differentiable, respectively concave and convex, and satisfy the Inada conditions $u'(0) = \infty$, $u'(\infty) = 0$, $h(0) = 0$, $h'(0) = 0$. Moreover, we assume that $n_0 \leq \bar{n}$, where $\bar{n} \in \bar{\mathbb{R}}^+$ is the maximal amount of effective units of labor individuals can supply. [7] The latter constraint on allocations will be imposed on all the allocations discussed in the rest of the paper. To save on notation, we will make this constraint implicit and avoid repeating it.

An allocation specifies the assignment of consumption and labor for each agent as a function of productivity $(c_0(\theta_0), c_1(\theta_0), n_0(\theta_0))$. We assume technology is linear in labor and capital, so that the resource constraints is given by

$$\int c_0(\theta_0)\, dF(\theta_0) + K_1 \leq \int n(\theta_0)\, dF(\theta_0) + RK_0,$$
$$\int c_1(\theta_0)\, dF(\theta_0) \leq RK_1,$$

where K_t denotes the aggregate capital, with rate gross rate of return R. Combining these two leads to the intertemporal resource constraint

$$\int c_0(\theta_0)\, dF(\theta_0) + \frac{1}{R} \int c_1(\theta_0)\, dF(\theta_0) \leq \int n(\theta_0)\, dF(\theta_0) + Rk_0. \tag{2}$$

2.2 Information, Incentives and Taxes

Following Mirrlees (1971), we assume that individual productivity θ_0 and work effort e_0 are privately observed. Only the product of the two, the effective units of labor $n_0 = e_0 \cdot \theta_0$ and consumption are publicly observable. Thus, type specific lump-sum taxes that ensure full efficiency are unavailable. Instead, we study constrained efficient allocations and the distorting tax systems that implement them.

[7] At times, we will need to restrict $\bar{n}$ to be finite. While this would slightly complicate the analysis, we could also assume that the upper bound applies to the units of labor n_0/θ_0 rather than to the effective units of labor n_0.

Consider for a moment the economy where the government has perfect commitment. By the revelation principle, any allocation that is attainable by some mechanism or tax system must satisfy the incentive compatibility constraints

$$u(c_0(\theta_0)) - h\left(\tfrac{n_0(\theta_0)}{\theta_0}\right) + \beta u(c_1(\theta_0)) \geq u(c_0(\theta_0')) - h\left(\tfrac{n_0(\theta_0')}{\theta_0}\right) + \beta u(c_1(\theta_0')) \quad \text{for all } \theta_0, \theta_0'. \tag{3}$$

In words, under a direct mechanism, the agent is asked to report productivity and is assigned consumption and labor as a function of this report. The incentive constraint ensures that reporting the truth is optimal. Of course, there are other ways of implementing allocations that are incentive compatible and we are interested in those that resemble tax systems.

Our first result provides a simple tax system that implements incentive compatible allocations. It features two separate nonlinear tax schedules, one for labor income and another for capital income. Agents are subject to the following budget constraint

$$\begin{aligned} c_0 + k_1 &\leq n_0 - T^n(n_0) + Rk_0, \\ c_1 &\leq Rk_1 - T^k(Rk_1). \end{aligned} \tag{4}$$

After observing their productivity θ_0 agents make consumption, saving and labor choices. Given tax schedules T^n and T^a, a competitive equilibrium is an allocation $c_0(\theta_0)$, $c_1(\theta_0)$, $n_0(\theta_0)$ and $k_1(\theta_0)$ such that (i) agents optimize: each agent θ_0 maximizes their utility (1) subject to (4); and (ii) markets clear: the intertemporal resource constraint (2) is satisfied with equality. We say that tax schedules (T^n, T^k) implements an incentive compatible allocation $(c_0(\theta_0), c_1(\theta_0), n_0(\theta_0))$ if the latter is a competitive equilibrium for some $k_1(\theta_0)$. Our implementation result can now be simply stated.

Proposition 1. *For any allocation $(c_0(\theta_0), c_1(\theta_0), n_0(\theta_0))$ that is nondecreasing in θ_0 satisfying (3) there exists tax schedules (T^n, T^k) that implement this allocation as a competitive equilibrium.*

One can show that incentive compatibility requires that $n_0(\theta_0)$ and $u(c_0(\theta_0)) + \beta u(c_1(\theta_0))$ be nondecreasing in θ_0. That is, higher skill workers produce and consume more (on average). Thus, as long as $c_0(\theta_0)$ and $c_1(\theta_0)$ are positively related it is possible to implement an incentive allocation with two separate tax schedules. Moreover, efficient allocations, with or without commitment, have nondecreasing $c_0(\theta_0)$ and $c_1(\theta_0)$, so that restricting attention to allocations that are implemented by separable tax schedules is not restrictive.

Note that Proposition 1 allows for bunching (i.e. allocations that are constant over a non-trivial interval of types). When bunching does not occur, so that the allocation $(c_0(\theta_0), c_1(\theta_0), n_0(\theta_0))$ is strictly increasing in θ_0, the first order conditions for the agent can be rearranged to obtain familiar expressions relating allocations to marginal tax rates on labor and capital allocations with

$$\frac{h'\left(\frac{n_0(\theta_0)}{\theta_0}\right)}{u'(c_0(\theta_0))} = \theta_0 \left(1 - T^{n\prime}(n_0(\theta_0))\right), \tag{5}$$

$$u'(c_0(\theta_0)) = \beta R(1 - T^{k\prime}(Rk_1(\theta_0)))u'(c_1(\theta_0)). \tag{6}$$

The first condition equates the marginal rate of substitution between effort and consumption to the after tax marginal wage; the second is the usual intertemporal Euler equation once one identifies $R(1 - T^{k\prime})$ as the appropriate gross after-tax marginal rate of return.

2.3 Perfect Commitment: Zero Capital Taxation

As a benchmark, consider the case with full commitment. An allocation is constrained efficient if it is incentive compatible and resource feasible and there is no other allocation with these properties that delivers the same or greater utility to all agents.

Our first result shows that it is optimal to set the tax on capital to zero $T^k(Rk_1) = 0$ for all k_1. This result follows as a simple corollary of the celebrated uniform taxation result by Atkinson and Stiglitz (1976). They showed that if preferences over a group of consumption goods are weakly separable from work effort, then these consumption goods should be uniformly taxed, so that no distortions are introduced in their relative consumption. In our case, the consumption in both periods (c_0, c_1) is weakly separable from work effort in the first period, so the result applies.

Proposition 2. *Any efficient allocation $(c_0^*(\theta_0), c_1^*(\theta_0), n_0^*(\theta_0))$ can be implemented with a nonlinear income tax T^n and a zero tax on capital $T^k(Rk_1) = 0$.*

This result is important because it establishes a benchmark for the results that follow. With commitment capital taxation should be zero. Thus, any distortion on capital that arises in the sequel is due to the lack of commitment.

3 No Commitment and Politics

In this section we depart from the assumption of full commitment to study an economy where policy choices are taken sequentially through democratic elections.

We adopt a probabilistic voting framework. At the beginning of each period two candidates $i = A, B$ face off in an election. The winner of the contest is determined by simple majority. Candidates attempt to maximize the probability of winning the electoral race. Before voting takes place, candidates present their platforms to the electorate, stating the policies they will pursue if elected. The winner then holds office for one period and is committed to implementing the platform previously proposed. Voters make choices weighing a self interested welfare calculation against an idiosyncratic inclination for one candidate or another. As is well known, under suitable assumptions, the end result of this political process is that candidates seek to maximize a utilitarian objective.

Before moving forward, it is worth remarking that all the analysis in our paper applies immediately if, instead of developing the elections with probabilistic voting, we simply assumed that policy is chosen in each period by a "Utiliarian planner" without commitment, or by a sequence of such planners. The only role that probabilistic voting plays is to provide a foundation for this assumption.

Some readers may find it useful to adopt the Utilitarian planner perspective directly, and consider the probabilistic structure as a motivating example. Although we do not object to this position, in our view, it is also important to provide a fully explicit political economy structure leading to the Utilitarian representation. Although the political economy model we use is certainly highly stylized, it captures elements of a representative democracy which is the main focus of our theory and the motivation for our analysis.

3.1 Probabilistic Voting

We now describe this setup in more detail. At the beginning of period $t = 0$ each candidate $j = A, B$ announces its platform, which consists of a proposed tax system $(T_{0,i}^{n,j}, T_0^{k,j})$. Voting then takes place and a winner j^* is determined. We describe voting behavior below. The winner j^* takes office and enacts $(T_0^n, T_0^k) = (T_0^{n,j^*}, T_0^{k,j^*})$. In the second period, two new candidates – which we denote again by $j = A, B$ – take this inherited tax system T_0^k as given. They also take as given the distribution of asset holdings in the population summarized by $k(\theta_0)$. At this point they can propose a platform to reform this system or not; if they do, they must specify a new tax schedule $T_1^{k,j}$. We assume that reforming requires paying a fixed cost $\kappa \geq 0$ in terms of good, otherwise the default policy from the first period is implemented. In particular, if a reform takes place then the resource constraint in the second period becomes

$$\int c(\theta_0)\, dF(\theta_0) \leq RK_1 - \kappa. \tag{7}$$

While one may interpret the fixed cost literally, perhaps as the opportunity cost of timely legislative procedures, its purpose here is to introduce a form of limited commitment. At one extreme, the case with $\kappa = \infty$ effectively delivers full commitment to the first period policy maker, as in the previous section. Indeed, the same outcome obtains for finite but high enough values of κ. On the other side of the spectrum, when $\kappa = 0$ there is no commitment and reform is imminent, no matter what tax schedules or asset distributions are inherited. Intermediate values of κ capture intermediate levels of commitment in the sense that reform may occur for some inherited policies and capital distributions, but not for others. Later, in the infinite-horizon setting, we will dispense with the fixed cost and study reputational equilibria that are sustained by trigger strategies, which is one way to endogenize the cost of reform.

In deciding which candidate to cast their vote for, agents care about the sum of two variables: the welfare the platform will imply for them and an idiosyncratic candidate-specific taste shock. The latter captures ideological preferences, fondness based on a candidates personality, or any other consideration that make individuals not vote entirely based on their self interest. It implies that for each productivity type θ_0, voters take different sides in the election. As a result, candidates choose their platform with an eye for pleasing agents across the productivity spectrum. Indeed, we will obtain the standard result that, in equilibrium, both candidates pick the platform that maximizes a utilitarian average of utility. This is in sharp contrast with the median voter setup, where there is a single type θ_0 that is the marginal voter and candidates cater their platform to this single agent. Specifically, in the first period agents consider the welfare implications of platforms $j = A, B$ and compute

$$v_0^j(\theta_0) = u(c_0^{0,j}(\theta_0)) + \beta u(c_1^{0,j}(\theta_0)) - h\left(\frac{n_0^{0,j}(\theta_0)}{\theta_0}\right) \qquad j = A, B,$$

where $(c_0^{0,j}(\theta_0), c_1^{0,j}(\theta_0), n_0^{0,j}(\theta_0))$ denotes the allocation that would result if in period $t = 0$ platform j won the election. Likewise, in period $t = 1$ agents compute the implications of platform $j = A, B$ into

$$v_1^j(\theta_0) = u(c_1^{1,j}(\theta_0)) \qquad j = A, B,$$

which captures their remaining lifetime utility. In period t, an agent i with productivity θ_0 votes for A over B if and only if

$$v_t^A(\theta_0) + \varepsilon^{i,A} > v_t^B(\theta_0) + \varepsilon^{i,B}; \tag{8}$$

ties are broken by voting with equal probability for each candidate. We assume that $\Delta_\varepsilon^i = \varepsilon^{i,B} - \varepsilon^{i,A}$ is independent with respect to θ_0 and is distributed uniformly on a symmetric interval around zero, $[-m_\varepsilon, m_\varepsilon]$ where $m_\varepsilon > 0$ is a measure of the dispersion of political biases in the population. [8] The probability that platform A wins the election is given by

$$\int G\left(v_t^A(\theta_0) - v_t^B(\theta_0)\right) dF(\theta_0) \tag{9}$$

where G is the distribution for Δ_ε. The political equilibrium takes a very simple form when the dispersion of beliefs is large enough relative to the range of utility payoffs resulting from different policy proposals.

Assumption 1. *The upper bound on effective labor is finite $\bar{n} < \infty$. There exists a bound $M > 0$ such that for all c, $n < \bar{n}$ and θ, $|u(c) - h(n/\theta)| < M$. Moreover, $2M(1+\beta) < m_\varepsilon$.*

The first part of Assumption 1 requires u and h to be bounded functions. The second part ensure that for every θ, and for any two given allocations of consumption and work effort $(c_0(\theta), c_1(\theta), n(\theta))$ and $(c_0'(\theta), c_1'(\theta), n'(\theta))$ for type θ, the dispersion of political biases is large enough that some agents of type θ will prefer the first allocation and some will prefer the second allocation. This ensures that the cumulative distribution function in (9) is linear $G\left(v_t^A(\theta_0) - v_t^B(\theta_0)\right) = \left(v_t^A(\theta_0) - v_t^B(\theta_0)\right)/(2m_\varepsilon)$. Since G is linear, each candidate positions their platform to maximize the utilitarian welfare criterion

$$\int v_t(\theta_0)\, dF(\theta_0). \tag{10}$$

It also follows that both candidates choose the same platform and get elected with equal probability.

Thus, in the first period politicians choose their platform to maximize

$$\int \left(u(c_0(\theta_0)) + \beta u(c_1(\theta_0)) - h\left(\frac{n_0(\theta_0)}{\theta_0}\right)\right) dF(\theta_0), \tag{11}$$

while in the second period they maximize

$$\int u\left(c_1(\theta_0)\right) dF(\theta_0). \tag{12}$$

3.2 Planning Problem with Political Constraints

We now show that the model can be solved backwards, starting from $t = 2$, and leads to a simple planning problem with a commitment constraint. We then study this problem

[8] Assuming the density is uniform simplifies the analysis but is not critical. As is well known, the same results would obtain for a larger class of non-uniform distributions that ensures that the candidates platform problem is sufficiently convex.

and derive the results for the equilibrium tax schedules (T^n, T^k).

The first thing to note is that, from the point of view at $t = 2$, the default tax schedule T_0^{k,j^*} and the distribution of assets in the population $k_1(\theta_0)$ combine to give a default allocation for consumption

$$c_1(\theta_0) = Rk_1(\theta_0) - T_0^{k,j^*}(Rk_1(\theta_0)).$$

Obviously, this is what is relevant for voters and candidates. If a candidate decides to reform then the platform will maximize (12) subject to the resource constraint (7) that applies in case of reform. Thus, if a reform takes place then consumption will equal

$$c_1(\theta_0) = RK_1 - \kappa$$

. Comparing the two alternatives, it follows that a reform can be avoided if and only if

$$\int u\big(Rk_1(\theta_0) - T_0^{k,j^*}(Rk_1(\theta_0))\big)\, dF(\theta_0) \geq u(RK_1 - \kappa). \tag{13}$$

Note that if a reform takes place then $T_1^k(Rk_1) = Rk_1 - RK_1 + \kappa$, so that capital is completely expropriated.

Turning to period $t = 0$, it always in the interest of candidates to propose platforms that will not be reformed. For suppose otherwise, that they propose a policy that leads to a reform in the second period. Then they could have done better by offering tax schedules that obtain the same constant allocation for consumption in $t = 1$. This saves them the fixed cost κ, allowing lower savings and higher consumption at $t = 0$. This shows that assuming candidates at $t = 1$ anticipate a reform leads to a contradiction. Hence, at $t = 0$ candidates will propose platforms that satisfy (13).

So far, we have argued in terms of the tax schedules. We now translate the argument directly to allocations. In period $t = 0$ candidates will propose platforms that solve the following planning problem

$$\max_{c_0, c_1, n_0} \int \Big(u\big(c_0(\theta_0)\big) + \beta u\big(c_1(\theta_0)\big) - h\big(n_0(\theta_0)/\theta_0\big)\Big)\, dF(\theta_0) \tag{14}$$

subject to the resource constraints (2) the incentive compatibility constraint (3) and the credibility constraint

$$\int u\big(c_1(\theta_0)\big)\, dF(\theta_0) \geq u\left(\int c_1(\theta_0)\, dF(\theta_0) - \kappa\right). \tag{15}$$

This program differs from the full commitment problem that maximizes a utilitarian objective only by the last constraint. Although reforms do not take place, the threat of one shape policy in the efforts to avoid it.

The dual of the above planning problem will prove more convenient. It is defined by minimizing the present value cost of delivering a certain average utility level V in an incentive compatible way while avoiding a reform:

$$\min_{c_0, c_1, n_0} \int \left(c_0(\theta_0) + \frac{1}{R} c_1(\theta_0) - n(\theta_0) \right) dF(\theta_0) \tag{16}$$

subject to the promise keeping constraint

$$\int \left(u\big(c_0(\theta_0)\big) + \beta u\big(c_1(\theta_0)\big) - h\big(n_0(\theta_0)/\theta_0\big) \right) dF(\theta_0) \geq V, \tag{17}$$

the incentive compatibility constraints (3), and the credibility constraint (15).

The political equilibrium in our two-period economy is entirely isomorphic to an environment with a utilitarian planner with limited commitment as in Chari and Kehoe (1990). To be more precise, consider an economy where in the first period, a utilitarian planner chooses an allocation to maximize (11). This planner, however, lacks commitment. Plans made in period 0 about period 1 consumption are not binding once the economy moves into period 1: the allocation can be reformed to maximize the period 1 utilitarian criterion (12) at cost κ. We could have adopted this alternative framework from the outset at a minimal cost in terms of economic substance. Moreover, this would have allowed us to dispense with Assumption 1. We chose instead to focus on an explicit equilibrium in a political game where politicians propose competing tax policies in regular in elections. In our view, this reinforces our implementation since the concept of a utilitarian planner can only be thought of a useful abstraction. By contrast, we provide an explicit example of an economy where decisions are taken by politicians chosen through elections and where political platforms consist of explicit tax schedules.

3.3 Optimal Progressive Capital Taxation

To derive the first-order condition, let $\nu \geq 0$ be the multiplier on the credibility constraint (15) and consider minimizing the Lagrangian

$$\int \left(c_0(\theta_0) + \frac{1}{R} c_1(\theta_0) - \nu u(c_1(\theta_0)) - n(\theta_0) \right) dF(\theta_0) + \nu u \left(\int c_1(\theta_0)\, dF(\theta_0) - \kappa \right) \tag{18}$$

subject to the incentive compatibility constraints (3) and the promise keeping constraint (17).

Note that in both constraints $c_0(\theta_0)$ and $c_1(\theta_0)$ enter through the expression $U(\theta_0) \equiv u(c_0(\theta_0)) + \beta u(c_1(\theta_0))$. It follows that any solution must solve the subproblem of minimizing (18) subject to $u(c_0(\theta_0)) + \beta u(c_1(\theta_0)) = U(\theta_0)$. The first-order conditions are

$$1 = \lambda(\theta_0) u'(c_0(\theta_0))$$
$$\frac{1}{R} + \nu \left(u'(RK_1 - \kappa) - u'(c_1(\theta_0)) \right) = \lambda(\theta_0) \beta u'(c_1(\theta_0))$$

Combining gives

$$u'(c_0(\theta_0)) = \beta R \frac{1}{1 + R\nu(u'(RK_1 - \kappa) - u'(c_1(\theta_0)))} u'(c_1(\theta_0)) \tag{19}$$

so that the marginal tax on capital is given by

$$T^{k\prime}(Rk_1(\theta_0)) = 1 - \frac{1}{1 + R\nu(u'(RK_1 - \kappa) - u'(c_1(\theta_0)))}. \tag{20}$$

Several implications follow from this simple formula. First, the tax schedule is progressive in the sense that it is increasing in consumption $c_1(\theta_0)$. Second, the sign of $T^{k\prime}$ is determined by the sign of $u'(RK_1 - \kappa) - u'(c_1(\theta_0))$, which may depend on θ_0. Indeed, for the agent consuming the most we have $RK_1 - \kappa = \int c_1(\theta_0)\, dF(\theta_0) - \kappa < \max_{\theta_0} c_1(\theta_0)$ ensuring that $T^{k\prime} > 0$ at the top. Similarly, for the credibility constraint to be binding, it must be the case that consumption at the bottom is lower than consumption if a reform were to take place: $\min_{\theta_0} c_1(\theta_0) < RK_1 - \kappa$. This implies that the marginal tax rate is always negative at the bottom.

Proposition 3. *Suppose that κ is low enough so that the full commitment solution is not feasible, so that the credibility constraint is strictly binding. The equilibrium tax function $T^k(Rk_1)$ is convex, with the marginal tax rate $T^{k\prime}(Rk_1)$ strictly increasing in capital income. The marginal tax rate is positive at the top, $T^{k\prime}(R\bar{k}_1) > 0$ for $\bar{k}_1 \equiv \max_{\theta_0} k_1(\theta_0)$, and negative at the bottom $T^{k\prime}(R\underline{k}_1) < 0$ for $\underline{k}_1 = \min_{\theta_0} k(\theta_0)$.*

The sign of the marginal tax rate is driven by the sign of $u'(RK_1 - \kappa) - u'(c_1(\theta_0))$ because this determines whether an additional unit of capital saved in the hands of an agent with productivity θ_0 tightens or loosens the credibility constraint. An extra unit of capital raises individual consumption, and thus raises the left hand side of the credibility constraint. At the same time, more capital raises the right hand side of the credibility constraint, the value of reforming. Indeed, for rich enough agents the constraint surely becomes tighter, since individual marginal utility is low enough. Conversely, an extra unit of capital in the hands of the poorest agent must increase average utility by more than it does in the reform state, given that the credibility constraint is binding reforming

4 Sustainability in an Infinite Horizon

We now turn to an infinite horizon version of the economy and study the policy problem as a dynamic game. Unlike the two-period economy, there is no longer a need to introduce an exogenous cost of reform to avoid the trivial outcome. Instead, reputational equilibria, sustained by trigger strategies, can deter policy makers from deviating away from expected policies for fear that such a deviation would spur a change in the private sector's expectations about future policy, leading to a worse equilibrium outcome.

As in the two-period version of the model, we rely on the Probabilisting voting model of electoral competition. All the analysis that follows applies immediately if, instead, we simply assumed that policy is chosen in each period by a sequence of Utiliarian planners, without commitment. The only role that probabilistic voting plays is to provide an explicit foundation for this assumption.

Model Setup

We now describe the general dynamic model and introduce the concept of sustainability.

Preferences. There is a continuum private agents with preferences given by

$$V_t = \sum_{s=0}^{\infty} \beta^s \, \mathbb{E}_{t-1} \left[u(c_{t+s}) - h(n_{t+s}/\theta_{t+s}) \right] \tag{21}$$

where $\beta < 1$ is the discount rate. We assume that the utility functions u and h are twice differentiable, u is concave, h convex, and impose the Inada conditions $u'(0) = \infty$, $u'(\infty) = 0$, $h(0) = 0$, $h'(0) = 0$. We assume that there is some maximal amount of effective labor that individuals can supply $n_{t+s} \leq \bar{n}$, $\bar{n} \in \bar{\mathbb{R}}^+$.[9] The latter constraint on allocations will be imposed on all the allocations discussed in the rest of the paper. To save on notation, we will make this constraint implicit and avoid repeating it.

We identify agents by their initial utility entitlement v_0 with distribution ψ in the population. Let V be the set of possible values for v.

[9] Note that all agents then have the same feasible set of consumption-labor bundles, regardless of their current shock θ_t. An alternative assumption is that all agents have some maximal amount of labor time n_t/θ_t. We conjecture that this would complicate the analysis, but not affect the main results in any way.

Information and incentives. We assume that types θ_t are independently and identically distributed across agents. Let π^∞ be the probability distribution over infinite paths $\theta^\infty \in \Theta^\infty$. Note that we do not, at this stage, we not necessarily assume that shocks are independent over time. Productivity shocks are assumed to be private information to each agent, requiring allocations to satisfy incentive compatibility constraints, which we describe below.

Technology. The resource constraints are given by:

$$C_t + K_{t+1} \leq F(K_t, N_t) \qquad t = 0, 1, \ldots \tag{22}$$

where C_t, N_t and K_t are aggregate consumption, aggregate labor and aggregate capital, respectively, in period t. The production function F is assumed to have constant returns to scale and is strictly increasing and continuously differentiable in both arguments.

Political Economy: The Policy Game. We set up the policy problem in terms of a dynamic game with actions and communication. Actions are dictated by technology, preferences and the political system. The communication requires a language or set of possible messages that can be used by agent i at time t: a general message space M_t. We now describe the game for any given choice of $\{M_t\}_{t=0}^\infty$. For every t, we denote by $M^t = \Pi_{s=0}^t M_s$ the space of message histories at date t.

We adapt the concept of sustainable equilibrium introduced by Chari and Kehoe (1990). This is essentially a refinement that focuses on symmetric perfect bayesian equilibria of anonymous games. It captures the notion that private agents assume that their individual decisions can neither affect other agent's decisions nor policy choices made by governments: in this sense, agents are behaving competitively.

We enter a period t with some aggregate history H_t. This aggregate history consists of the sequence of past implemented policies $\{\{c_s^v, x_s^v\}, K_{s+1}\}_{s\leq t-1}$. For every agent i, we denote by $\underline{h}_t^i$ his full individual history $\left\{\{\theta_s^i, m_s^i, n_s^i, c_s^i\}_{s\leq t-1}, \theta_t\right\}$. We denote by $\hat{h}_t^i$ his interim public individual history $\left\{\{m_s^i, n_s^i, c_s^i\}_{s\leq t-1}, m_t^i, n_t^i\right\}$ and by $\underline{\hat{h}}_t^i$ his interim full individual history $\left\{\{\theta_s^i, m_s^i, n_s^i, c_s^i\}_{s\leq t-1}, \theta_t^i, m_t^i, n_t^i\right\}$.

The timing of events within the period is then as follows:

1. Agent types θ_t^i are realized and observed privately by each agent. Each agent i sends a message $m_t^i \in M_t^i$ to the planner and choose how much labor to supply n_t^i; output is produced: $F(K_t, \int n_t^i di)$.

2. Two new candidates $j = A, B$ make proposals regarding the current distribution of consumption $\{c_t^{v,j}, x_t^{v,j}\}$. For every v, $c_t^{v,j}$ is a function that maps an interim individual public history $\hat{h}_t$ into a current consumption level; $x_t^{v,j}$ is a function that maps an interim individual public history $\hat{h}_t$ into 0,1. When $x_t^{v,j}(\hat{h}_t) = 1$, the platform proposed by candidate j imposes a $-\infty$ payoff in the period to an agent with initial promised utility v and individual public history $\hat{h}_t$. We are implicitly making the assumption that governments can inflict arbitrary large punishments—such as jail or death—on agents.[10] They also make a proposal for capital investment K_{t+1}^j, where $j \in \{A, B\}$. These proposals must be resource feasible in period t. Let $\hat{H}_t$ be the interim history generated by adding to H_t the two proposed platforms $\{c_t^{v,j}, x_t^{v,j}\}$ for $j = A, B$.

3. Agents then observe these proposals. In deciding how to vote, they consider their own personal utility from each proposal, denoted $v_t^{i,j} \equiv v_t^j(\hat{\underline{h}}_t^i)$ for $j = A, B$. In addition, each agent i receives an idiosyncratic "political bias shock" denoted by $\varepsilon_t^{i,j}$ for each candidate $j = A, B$. Agents then vote for candidates A or B, by choosing $z_t^i \in A, B$, to maximize the sum $v_t^{i,j} + \varepsilon_t^{i,j}$ for $j = A, B$. We assume that the political bias shocks are independent across agents and identically distributed and that $\varepsilon_t^{i,B} - \varepsilon_t^{i,A}$ is uniformly distributed over the interval $[-m_\varepsilon, m_\varepsilon]$.

4. The candidate with the most votes wins and enforces the platform they ran on. If both have half the votes then each is elected with $1/2$ probability.

5. We move to the next period.

We now define strategies. In period t, each agent conditions his action on H_t and his interim full individual history $\underline{h}_t$ for the choice of labor and messages made by agents and on $\hat{H}_t$ and his interim full individual history $\hat{\underline{h}}_t$ for the choice of platforms by candidates. Voting is mechanical: in period t, each agent votes for the candidate which, if elected would assure him the highest continuation utility v_t under the considered strategy σ. Hence we do note incorporate $z_t(\hat{H}_t, \hat{\underline{h}}_t)$ in the definition of an individual strategy. Thus, we write

$$\sigma \equiv \{n_t^v(H_t, \underline{h}_t), m_t^v(H_t, \underline{h}_t)\}_{t=0}^{\infty}$$

for the agents' strategies.

[10]This assumption facilitates the analysis and guarantees that we can restrict restrict our attention to certain individual deviations: deviations where an agent adopts the actions adopted another agent in equilibrium. Since agents are atomistic, all the other individual deviations are deterred at no cost by setting $x^t = 1$ when the corresponding histories are observed.

A candidate $j = A, B$ from period t has strategy

$$\tau_t^j \equiv (\{c_t^{v,j}(H_t), x_t^{v,j}(H_t)\}, K_{t+1}^j(H_t)),$$

and let $\tau \equiv \{(\tau_t^A, \tau_t^B)\}_{t=0}^{\infty}$ denote policy strategies. Given a history H_t, the continuation of the plan (σ, τ) determines an elected platform in period t and hence a history H_{t+1}.

Given a message space M^{∞}, and an initial distribution of utility entitlements ψ, a sustainable equilibrium is a a pair of strategy profiles (σ, τ) that satisfies the following conditions: (i) Given the policy plan τ, σ_t^v maximizes utility given H_t, v and $\underline{h}_t$ at stage 1 of period t; (ii) Given σ and τ_t^{-j}, τ_t^j maximizes the number of votes given H_t; (iii) the utility attained by an agent with initial promised utility v is equal to v.

In order for the political equilibrium to be well behaved, we assume that utility is bounded and that the support for political bias shocks is wide enough.

Assumption 2. *The upper bound on effective labor is finite $\bar{n} < \infty$. There exists a bound $M > 0$ such that for all c, $n < \bar{n}$ and θ, $|u(c) - h(n/\theta)| < M$. Moreover, $2M/(1-\beta) < m_\varepsilon$.*

Assumption 2 is an adaptation of Assumption 1 to an infinite-horizon setting. It guarantees that, regardless of the two candidates proposals, politicians receive some positive votes from every agent type. In this sense, the electoral game is at an interior. This implies that the equilibrium maximizes a utilitarian welfare criterion.

To each sustainable equilibrium, we can associate an initial resource cost K_0. In what follows, we are interested in characterizing, for an given initial distribution of utility entitlements ψ, the best sustainable equilibrium: the one with the minimal initial resource cost.

Following Chari and Kehoe (1990), we develop a set of necessary and sufficient conditions for an allocation to be the outcome of a sustainable equilibrium. We term these allocations sustainable. Studying the best sustainable equilibrium boils down to characterizing the optimal sustainable allocation using a constrained programming problem.

5 Independent Shocks

In this section, we study the case where productivity shocks are independent and identically distributed (i.i.d.) over time. In Section 6 we fully relax this assumption and allow general stochastic processes for productivity.

5.1 Sustainable Allocations and the Credibility Constraint

In the previous section we developed notation and definitions for the policy game and sustainable equilibria that were general enough to handle productivity shocks that may not be dependent over time. In particular, we did not assume that the revelation principle held and allowed for more general message spaces. However, when shocks are i.i.d. using a direct mechanism where $M_t = \Theta$ and we impose that agent's reveal truthfully is without loss of generality. There are two reasons for this. First, at the interim stage within a period, when politicians propose their platforms, current output is already fixed, so that knowledge of current shocks is irrelevant. Second, because shocks are independent over time, knowledge of past shocks do not create opportunities for exploiting this information.[11]

Any allocation that can be implemented as a sustainable equilibrium can also be derived as a sustainable equilibrium where agents report their true type θ_t in every period. The i.i.d. assumption is crucial to this result: the information revealed in equilibrium is irrelevant for the set of payoffs achieved if a deviation occurs.

We start our characterization of the best sustainable equilibrium by laying down some notations for the i.i.d. case, using the fact that we can focus on direct mechanisms.

Feasible allocations. An *allocation* $\left(\{c_t^v, n_t^v\}_{t\geq 0, v\in V}, \{K_t\}_{t\geq 0}\right)$ consists of sequences of consumption functions $c_t^v : \Theta^{t+1} \to \mathbb{R}^+$, labor supply functions $n_t^v : \Theta^{t+1} \to \mathbb{R}^+$, capital stocks $K_t \in \mathbb{R}$.

An agent's reporting strategy $\sigma \equiv \{\sigma_t\}$ is a sequence of functions $\sigma_t : \Theta^{t+1} \to \Theta$ that maps histories of shocks θ^t into a current report $\hat{\theta}_t$. Any strategy σ induces a history of reports $\sigma^t : \Theta^{t+1} \to \Theta^{t+1}$. We use σ^* to denote the truth-telling strategy with $\sigma_t^*(\theta^t) = \theta_t$ for all $\theta^t \in \Theta^{t+1}$.

Given an allocation $\left(\{c_t^v, n_t^v\}_{t\geq 0, v\in V}, \{K_t\}_{t\geq 0}\right)$, the utility obtained by an agent with initial utility entitlement v from any reporting strategy σ is

$$U\left(\{c_t^v, n_t^v\}, \sigma; \beta\right) \equiv \sum_{t=0}^{\infty} \sum_{\theta^t \in \Theta^{t+1}} \beta^t \left[u\left(c_t^v\left(\sigma^t(\theta^t)\right)\right) - h\left(n_t^v(\sigma^t(\theta^t))/\theta_t\right)\right] \Pr(\theta^t).$$

The allocation delivers utility v to all agents entitled to v if

$$U\left(\{c_t^v, n_t^v\}, \sigma^*; \beta\right) = v$$

[11] Similar results regarding the use of direct mechanisms arise in other settings such as Acemoglu, Golosov and Tsyvinski (2007) and by Sleet and Yeltekin (2008b).

The allocation is *incentive compatible* if truth-telling is optimal, so that

$$U\big(\{c_t^v, n_t^v\}, \sigma^*; \beta\big) \geq U\big(\{c_t^v, n_t^v\}, \sigma; \beta\big) \tag{23}$$

for all strategies σ and initial utility entitlement v.

For a given initial distribution of entitlements ψ, we say that an allocation $\big(\{c_t^v, n_t^v\}, K_t\big)$ is *feasible* if: (i) it is incentive compatible; (ii) it delivers expected utility of v to all agents initially entitled to v; and (iii) it satisfies the following sequence of resource constraints:

$$\int \sum_{\theta^t} c_t^v(\theta^t) \Pr(\theta^t)\, d\psi(v) \leq F\big(K_t, \int \sum_{\theta^t} n_t^v(\theta^t) \Pr(\theta^t)\, d\psi(v)\big) - K_{t+1} \qquad t = 0, 1, \dots \tag{24}$$

Sustainable allocations. We say that an allocation $\big(\{c_t^v, n_t^v\}, K_t\big)$ is *sustainable* if it is the outcome of a sustainable equilibrium. Following Chari and Kehoe (1990), we can derive a simple set of necessary and sufficient conditions for a feasible allocation to be the outcome of a sustainable equilibrium.

$$\int U\big(\{c_{t+s}^v, n_{t+s}^v\}_{s\geq 0}, \sigma^*; \beta\big) d\psi(v) \geq \hat{W}\big(K_t, \{n_t^v(\theta^t)\}\big) \qquad t = 0, 1 \dots \tag{25}$$

Thus, an allocation is sustainable if: (i) it is feasible; and (ii) it satisfies the sequence of credibility constraints (25).

The value of deviating from the equilibrium outcome path is given by $\hat{W}$. This endogenous object represents the payoff corresponding the most profitable deviation by the government followed by the worst equilibrium payoff:

$$\hat{W}(K, \{n_\theta\}) \equiv \max_{K'} \Big\{ u(F(K, \sum_\theta \int n_\theta \Pr(\theta)) - K') - \\ \sum_\theta \int h(n_\theta/\theta) \Pr(\theta) + \beta W(K') \Big\} \tag{26}$$

where $W(K)$ represents the welfare associated with the worst sustainable equilibrium of the policy game, which we characterize below.

Given a distribution of welfare entitlements ψ, we term optimal sustainable allocations the allocations that are the outcome of the best sustainable equilibria. More precisely, given ψ optimal sustainable allocation solves the following planning problem:

$$\min K_0 \tag{27}$$

subject to $\left(\{c_t^v, n_t^v\}, \{K_t\}\right)$ being a sustainable allocation.

In our case the worst equilibrium payoff W can be characterized quite easily.

Lemma 1. *The worst payoff function W can be represented as the fixed point in a simple functional equation:*

$$W(K) = \min_{n \in [0, \bar{n}]} \max_{K'} \left\{ u(F(K, n) - K') - \mathbb{E}\left[h(n/\theta)\right] + \beta W(K') \right\} \tag{28}$$

Moreover: (i) $W(K)$ is nondecreasing and concave, and (ii) $\hat{W}(K, \{n_\theta\})$ is increasing, concave, and differentiable.

Note that by the envelope theorem, we have

$$\hat{W}_K(K, \{n^v\}) = F_K\left(K, \int n^v\right) u'\left(\hat{c}(K, \{n^v\})\right) \tag{29}$$

where $\hat{c}(K, \{n^v\})$ is the optimal equalized consumption level chosen in (26).

5.2 A Modified Inverse Euler Equation

Putting multipliers $\beta^t \mu_t$ and $\beta^t \nu_t$ on the resource constraints (24) and the credibility constraints (25), we can derive two key necessary first order condition:

$$\frac{\mu_{t+1}}{\mu_t} \beta F_K(K_{t+1}, N_{t+1}) - \frac{\nu_{t+1}}{\mu_t} \beta \hat{W}_K(K_{t+1}, \{n_{t+1}^v\}) = 1 \tag{30}$$

and

$$\frac{1}{u'(c^v(\theta^t))} = \frac{\mu_{t+1}}{\mu_t} \mathbb{E}_{\theta^t}\left[\frac{1}{u'(c^v(\theta^{t+1}))}\right] - \frac{\nu_{t+1}}{\mu_t}, \tag{31}$$

where $\mathbb{E}_{\theta^t}$ is the expectations operator, conditional on the event $\hat{\theta}^t = \theta^t$.

Equation (30) states that the social intertemporal rate of return $\frac{\mu_t}{\beta \mu_{t+1}}$ is given by

$$F_K(K_{t+1}, N_{t+1}) - \frac{\nu_{t+1}}{\mu_{t+1}} \hat{W}_K(K_{t+1}, \{n_{t+1}^v\}).$$

Accumulating capital from t to $t+1$ has the dual effect of relaxing the resource constraint and tightening the credibility constraint at date $t+1$. The social intertemporal rate of return on capital incorporates these two effects. This is a crucial difference with Farhi and Werning (2008): political economy considerations introduce a wedge between the social rate of return on capital and the marginal product of capital.

The left-hand side of (31) together with the first term on the right-hand side is the standard inverse Euler equation.

The second term on the right-hand side is novel, since it is zero when the credibility constraint at time $t+1$ is ignored. The negative constant $-\frac{\nu_{t+1}}{\mu_t}$ on the right hand side of (31) shows that as in our two period example, and similarly to Farhi and Werning (2007) and Farhi and Werning (2008), the transmission of consumption inequality from one period to the next is less than one for one when $\frac{\mu_{t+1}}{\mu_t} > 1$ – as is the case, for example, in steady states. Consumption then mean-reverts towards $\frac{\nu_{t+1}}{\mu_{t+1}-\mu_t}$ from one period to the next. Reducing inequality allows the planner to improve the credibility of the allocation by reducing the risk for the desired social arrangement to be overturned in a future election.

5.3 Tax Implementation

We now explore two tax implementations and focus on their implications for capital taxation. Our first implementation is along the lines of Kocherlakota (2005) and features linear taxes on capital. These tax rates generally depend on the entire history of reports, including the current shock's report. The dependence on the current report makes the net-of-tax return on capital risky. As shown by Kocherlakota, this feature is sufficient to discourage double deviations in reports and savings. Our second implementation features two-stages and is closer to the one used in the two period version of the model. In the first stage, before the period's productivity is realized, savings are taxed according to a nonlinear schedule that is decreasing and convex. This tax on capital reflects the features of our main results, both in terms of the progressivity and the sign of taxation. In the second stage, a linear wealth tax is implemented exactly as in Kocherlakota (2005). In particular, tax rates are zero on average, but vary to deter double deviations in reports and savings. If there were no uncertainty in the next period's skill, then the second-stage linear wealth tax would be identically zero, just as in the two period version of the model.

We proceed as follows. We first describe the first implementation, which is both simpler and closer to existing implementations. We then characterize some of its important implications of capital taxation. Finally, we describe the second implementation.

Linear capital taxes. Any allocation that is incentive compatible and feasible, and has strictly positive consumption, can be implemented by a combination of taxes on labor income and taxes on capital income. Here we first describe this implementation, and explore some features of the optimal capital tax in the next subsection.

For any incentive-compatible and feasible allocation $\{c_t^v(\theta^t), n_t^v(\theta^t)\}$ we propose an implementation along the lines of Kocherlakota (2005). In each period, conditional on the history of their dynasty's reports $\hat{\theta}^{t-1}$ and wealth b_t, individuals report their current shock $\hat{\theta}_t$, produce, consume, pay taxes and save wealth subject to the following set of budget constraints

$$c_t + b_{t+1} \leq n_t(\hat{\theta}^t) W_t - T_t^v(\hat{\theta}^t) + \big(1 - \tau_t^v(\hat{\theta}^t)\big) R_{t-1,\,t} b_t \qquad t = 0, 1, \ldots \tag{32}$$

In this equation, $W_t = F_N\big(K_t, \int \sum_{\theta^t} n_t^v(\theta^t) \Pr(\theta^t)\, d\psi(v)\big)$ is the before-tax wage, $R_{t-1,t} = F_K\big(K_t, \int \sum_{\theta^t} n_t^v(\theta^t) \Pr(\theta^t)\, d\psi(v)\big)$ is the before-tax interest rate across periods, and initially $b_0 = K_0$. Individuals are subject to two forms of taxation: a labor income tax $T_t^v(\hat{\theta}^t)$, and a proportional tax on wealth $R_{t-1,\,t} b_{t-1}$ at rate $\tau_t^v(\hat{\theta}^t)$.[12]

Given a tax policy $\{T_t^v(\theta^t),\ \tau_t^v(\theta^t)\}$, an equilibrium is a sequence of wages and interest rates $\{W_t, R_{t,\,t+1}\}$, an allocation for consumption, labor and bequests $\{c_t^v(\theta^t), n_t^v(\theta^t), b_t^v(\theta^t)\}$; and a reporting strategy $\{\sigma_t^v(\theta^t)\}$ such that: (i) $\{c_t^v,\ b_t^v, \sigma_t^v\}$ maximize utility subject to (32), taking wages and interest rates $\{W_t, R_{t-1,t}\}$ and tax policy $\{T_t,\ \tau_t\}$ as given; (ii) in each period t, aggregate capital K_t and labor N_t maximize profits $F(k,n) - R_{t-1,t}k - W_t n$ taking the wage and interest rate as given, or equivalently $W_t = F_N(K_t, N_t)$ and $R_{t-1,t} = F_K(K_t, N_t)$; (iii) markets clear: the resource constraints (22) are satisfied with equality. We seek a tax policy that implements efficient allocations as a competitive equilibrium with truth-telling.

For any feasible, incentive-compatible allocation $\{c_t^v,\ n_t^v\}$ with strictly positive consumption we construct a tax policy that induces an equilibrium where all agents bequeath $b_t = K_t$. First, using the budget constraint with equality, let

$$T_t^v(\theta^t) = W_t n_t^v(\theta^t) + (1 - \tau_t^v(\theta^t)) R_{t-1,t} K_t - c_t^v(\theta^t) - K_{t+1}.$$

Second, following Kocherlakota (2005), set the linear tax on inherited wealth to

$$\tau_t^v(\theta^t) = 1 - \frac{1}{\beta R_{t-1,\,t}} \frac{u'(c_{t-1}^v(\theta^{t-1}))}{u'(c_t^v(\theta^t))}. \tag{33}$$

These choices work because for any reporting strategy σ, the agent's consumption Euler

[12] In this formulation, taxes are a function of the entire history of reports, and labor income n_t is mandated given this history. However, if the labor income histories $n^t \colon \Theta^t \to \mathbb{R}^t$ being implemented are invertible, then by the taxation principle we can rewrite T and τ as functions of this history of labor income and avoid having to mandate labor income. Under this arrangement, individuals do not make reports on their shocks, but instead simply choose a budget-feasible allocation of consumption and labor income, taking as given prices and the tax system.

equation

$$u'\left(c_t^v\left(\sigma^t(\theta^t)\right)\right) = \beta R_{t,t+1} \sum_{\theta_{t+1}} u'\left(c_{t+1}^v\left(\sigma^{t+1}(\theta^t, \theta_{t+1})\right)\right)\left(1 - \tau_{t+1}^v\left(\sigma^{t+1}(\theta^t, \theta_{t+1})\right)\right) \Pr(\theta_{t+1})$$

holds. Since the budget constraints hold with equality, savings choice is optimal regardless of the reporting strategy σ. The allocation is incentive compatible by hypothesis, so it follows that truth telling σ^* is optimal. Resource feasibility ensures that the markets clear.[13]

Optimal Progressive Capital Taxation. For efficient allocations, the assignment of consumption and labor at in any period depends on the history of reports in a way that can be summarized by the continuation utility $v_t(\theta^{t-1})$. Therefore, the capital tax $\tau^v(\theta^{t-1}, \theta_t)$ can be expressed as a function of $v_t(\theta^{t-1})$ and θ_t; abusing notation we denote this by $\tau_t(v_t, \theta_t)$. Similarly write $c_{t-1}(v_t)$ for $c_{t-1}^v(\theta^{t-1})$. The average capital tax rate $\bar{\tau}_t(v_t)$ is then defined by

$$\bar{\tau}_t(v_t) \equiv \sum_{\theta} \tau(v_t, \theta) \Pr(\theta).$$

Combining equations (30)and (31), we can derive the following formula:

$$\bar{\tau}_t\left(v_t\right) = \frac{\beta \hat{W}_K(K_{t+1}, \{n_{t+1}^v\}) - u'\left(c_t(v_t)\right)}{\beta F_K\left(K_{t+1}, N_{t+1}\right)} \frac{\nu_{t+1}}{\mu_{t+1}} \tag{34}$$

This formula has crisp implications for both the level of capital taxes and for their progressivity.

The severity of the commitment problem arising from voting in the next period depends on how much capital is accumulated. The higher K_{t+1}, the higher the utility achieved under the worst credible allocation, and the tighter the credibility constraint. The optimal credible allocation takes this into account and mitigates future commitment problems by lowering the rate of capital accumulation. This is reflected in the implementation by the positive term $\hat{W}_K(K_{t+1}, \{n_{t+1}^v\})$ in (34).

An opposing force pushes average capital taxes in the opposite direction. Agents do not internalize the effects of their economic decisions on future political outcomes. In particular, agents do not internalize that by delaying consumption, they contribute to increasing average future welfare and thereby to loosening future credibility constraints

[13] A version of Ricardian equivalence holds, so that the same allocation can be implemented with the same capital taxes, but adjusting the income taxes and savings. In particular, it is possible to have agents with higher v_t saving more.

or in other words to lowering the likelihood of a political renegociation. This can be interpreted as a form of externality from future consumption. Taxes can then be seen as a way of countering these externalities as prescribed by Pigou. This is the origin of the negative term $-u'\left(c_t^v(v_t)\right)$ in (34).

Hence the sign of the average capital tax $\mathbb{E}_{\theta^t}\left[\tau(\theta^{t+1})\right]$ depends on the balancing act between the alleviation of the Pigovian externality and the mitigation of future commitment problems. The former pushes in the direction of negative capital taxes, while the latter introduces a force in the direction of positive capital taxes.

More precisely the sign of the expected capital tax burden $\mathbb{E}_{\theta^t}\left[\tau(\theta^{t+1})\right]$ after history θ^t is determined by the sign of

$$\beta \hat{W}_K(K_{t+1}, \{n_{t+1}^v\}) - u'(c_t(v_t))$$

Equation (29) demonstrates that $\hat{W}_K(K_{t+1}, \{n_{t+1}^v\})$ is the product of R_{t+1} times the marginal utility of consumption $u'\left(\hat{c}(K_{t+1}, \{n_{t+1}^v\})\right)$ that would prevail if a deviation from equilibrium occurred at $t+1$ where the elected government implements a platform that reaps the benefits of equalizing consumption at $t+1$ at the cost of triggering a reversion to the worst equilibrium from date $t+2$ on.

This equation can be interpreted as a fictitious Euler equation, determining wether an agent after history θ^t would save at the margin if he anticipated consumption and work effort to be distributed according to the worst credible allocation from next period on. We can therefore develop the following heuristic. The expected capital burden after history θ^t is positive if and only the corresponding agent would save at the margin if he anticipated consumption and work effort to be distributed according to the worst credible allocation from next period on.

Equation (34) can be rewritten in the following useful way

$$\bar{\tau}_{t+1}\left(v_{t+1}\right) = \frac{\hat{W}_K(K_{t+1}, \{n_{t+1}^v\}) - R_{t+1}\left(\mathbb{E}_{\theta^t}\left[u'\left(c^v(\theta^{t+1})\right)\right]\right)^{-1}}{\beta^{-1}\mu_t \nu_{t+1}^{-1} + \hat{W}_K(K_{t+1}, \{n_{t+1}^v\}) - R_{t+1}\left(\mathbb{E}_{\theta^t}\left[u'\left(c^v(\theta^{t+1})\right)\right]\right)^{-1}}.$$

Hence the sign of the average capital tax $\bar{\tau}_{t+1}\left(v_{t+1}\right)$ is determined by the sign of

$$u'\left(\hat{c}(K_{t+1}, \{n_{t+1}^v\})\right) - \left(\mathbb{E}_{\theta^t}\left[u'\left(c^v(\theta^{t+1})\right)\right]\right)^{-1}.$$

Capital taxes are positive if the harmonic average of the marginal utility of consumption on the equilibrium path is lower than the marginal utility of consumption that would occur under a political deviation where consumption is equalized next period.

Independently of the predictions for the level of capital taxes, our model has sharp implications for the optimal dependence of the average capital tax with respect to the history of past shocks encoded in the promised continuation utility v_t. The average capital tax is an increasing function of consumption, which, in turn, is an increasing function of v_t. Thus, capital taxation is progressive.

Proposition 4. *An optimal credible allocation with strictly positive consumption can be implemented by a combination of income and capital taxes. The optimal average capital tax $\bar{\tau}_t(v_t)$ defined by (34) is increasing in promised continuation utility v_t.*

Thus, the structure of capital taxes that arise from a concern for reputation in an infinite horizon model has close parallels with the two period model featuring the fixed cost of reform.

A Nonlinear Capital Tax Implementation. Building on the discussion above, we now describe our second implementation. Each period, individuals report their current shock $\hat{\theta}_t$, produce, consume, pay taxes and save subject to the budget constraints

$$c_t(\theta^t) + b_t(\theta^t) \leq W_t n_t^v(\hat{\theta}^t) - T_t^{n,v}(\hat{\theta}^t) - T_t^b(b_t) + (1 - \tau_t^{b,v}(\hat{\theta}^t))R_{t-1,t}b_{t-1}(\theta^{t-1}) \tag{35}$$

Individuals are subject to three forms of taxation: an income tax $T_t^{n,v}(\hat{\theta}^t)$, a capital tax $T_t^b(b_t)$ and a proportional tax on saved wealth $R_{t-1,t}b_{t-1}$ with rate $\tau_t^{b,v}(\hat{\theta}^t)$. Given a tax policy $\{T^{n,v}, T^b, \tau^{b,v}\}$, a competitive equilibrium is defined exactly as before, but replacing the budget constraint (32) with (35). Once again, we will construct a tax policy that implements efficient allocations as a competitive equilibrium with truth-telling.

We have already argued that continuation utility v_t is a sufficient state variable for efficient allocations. The continuation utility v_t depends on the history of a reports θ^{t-1} and the initial welfare entitlement v, so we write $v_t = v_t(\theta^{t-1}, v)$ to emphasize this dependence. In our implementation, there will be a one to one mapping between bequests and continuation utility, so that we can keep track of the latter using the former.

First, select any sequence of strictly increasing functions $B_t(v_{t+1})$, normalized so that

$$\int \sum_{\theta^t} B_t(v_{t+1}(\theta^t, v)) \Pr(\theta^t) d\psi(v) = K_{t+1}.$$

Next, let the capital tax schedule $T_{t-1}^b(\cdot)$ for any $t = 1, 2, \ldots$ solve

$$T_{t-1}^{b\prime}(B_{t-1}(v_t)) = \frac{\bar{\tau}_t(v_t)}{1 - \bar{\tau}_t(v_t)}$$

with $T^b_{t-1}(K_t) = 0$. Since $B_{t-1}(v_t)$ is increasing in v_t and $\bar{\tau}_t(v_t)$ is negative and increasing in v_t, the capital tax schedule $T^b_t(\cdot)$ is decreasing and convex. Set the wealth tax rate to

$$\tau_t^{b,v}(\theta^t) \equiv 1 - \frac{u'(c^v_{t-1}(\theta^{t-1}))}{\beta R_{t-1,t} u'(c^v_t(\theta^t))(1 - \bar{\tau}_t(v_t(\theta^{t-1}, v)))}.$$

Finally, the income tax schedule $T^n_t(\theta^t)$ is defined so that the budget constraint holds with equality at the proposed allocation and savings are given by $b^v_{t-1}(\theta^{t-1}) = B_{t-1}(v_t(\theta^{t-1}, v))$.

Proposition 5. *Efficient interior allocations can be implemented by a combination of an income tax T^n_t, a capital tax T^b_t and a wealth tax $\tau^{b,v}_t$. The wealth tax is linear and its average is equal to zero:*

$$\sum_{\theta_t} \tau_t^{b,v}(\theta^{t-1}, \theta_t) \Pr(\theta_t) = 0.$$

The capital tax schedule $T^b_t(\cdot)$ is convex: the associated marginal tax rate on capital $T^{b\prime}_t(\cdot)$ is increasing in the amount of savings.

The proof of the proposition is similar to that of the previous implementation. In particular, by construction, the agent's consumption Euler equation

$$\begin{aligned} &u'(c^v_t(\sigma^t(\theta^t)))(1 + T^{b\prime}(B_t(v_{t+1}))) \\ &\qquad = \beta R_{t,t+1} \sum_{\theta_{t+1}} u'(c^v_{t+1}(\sigma^{t+1}(\theta^t, \theta_{t+1})))(1 - \tau^{b,v}_{t+1}(\sigma^{t+1}(\theta^t, \theta_{t+1}))) \Pr(\theta_{t+1}) \end{aligned}$$

holds for any reporting strategy σ. In addition, note that given any reporting strategy σ, the budget set is convex since the capital tax T^b_t is convex and the wealth tax is linear. Thus, first order conditions are sufficient for optimality of consumption and savings decisions, given σ. Hence, given a reporting strategy σ, the resulting consumption and labor allocation $\{c^v_t(\sigma^t(\theta^t)), n^v_t(\sigma^t(\theta^t))\}$ with savings given by $\{B_t(v_{t+1}(\sigma^t(\theta^t), v))\}$ is optimal from the perspective of the agents. Since the original allocation is incentive compatible, it follows that truth-telling is optimal. The resource constraint together with the budget constraints then ensure that the asset market clears.

This implementation is appealing because it decouples a nonlinear capital tax schedule $T^b_t(\cdot)$, that parallels the analysis of the two period model, from the linear wealth tax associated with the standard inverse Euler equation as studied in Kocherlakota (2005).

6 Non-independent Shocks

So far, we have assumed that the productivity shocks are independent over time. In this section we extend the analysis to the case where shocks are not independent.

Assuming that shocks are independent over time simplified the analysis of the optimal credible equilibrium for two related reasons. First, it made the worst equilibrium particularly simple, with allocations that were independent of the history of reports. This justified, without loss of generality, having agents reveal all their information. That is, we could use the revelation principle, employing a direct mechanism and imposing truth telling in equilibrium. This is the second property we made use of.

In contrast, when shocks are not independent over time, the payoff after a deviation may potentially depend on the information obtained from past reports about future productivities. For example, if the planner has very good information regarding agent's productivities it could exploit that information to obtain a better allocation. Ex ante, this implies that it may be better to play a game where some information is withheld. In general, the revelation principal fails and the best equilibrium may not use a direct mechanism, with truth-telling, between the private sector and the government.

Given these challenges, we make progress along two different routes. First, we find additional assumptions that guarantee that the revelation principle applies even when shocks are persistent. Under certain conditions the worst equilibrium calls for ignoring any past information and setting labor to zero for all agents. The previous results then immediately generalize.

Second, we tackle the case where the revelation principle fails. We show that the variations behind the our main results go through even if the government uses a mechanism other than direct reports on productivity as long as one assumes that the worst equilibrium is differentiable with respect to capital. We show that the average marginal distortion on capital imposed on some group, conditional on all information known to the government is increasing in current consumption.

6.1 Sustainable Allocations and the Credibility Constraint

In this subsection, we revisit the notation of a sustainable allocation for the non i.i.d. case and develop some needed notation.

Sustainable allocations. An *allocation* $\left(\{c_t^v, n_t^v\}_{t\geq 0, v\in V}, \{K_t\}_{t\geq 0}\right)$ consists of sequences of consumption functions $c_t^v : M^t \to \mathbb{R}^+$, labor supply functions $n_t^v : M^t \to \mathbb{R}^+$, capital stocks $K_t \in \mathbb{R}$.

A agent's reporting strategy $\sigma \equiv \{\sigma_t\}$ is a sequence of functions $\sigma_t : \Theta^{t+1} \to M_t$ that maps histories of shocks θ^t into a current report $m(\theta_t)$. Any strategy σ induces a history of reports $\sigma^t : \Theta^{t+1} \to M^t$.

Given an allocation $\left(\{c_t^v, n_t^v\}_{t\geq 0, v\in V}, \{K_t\}_{t\geq 0}\right)$, the utility obtained by an agent with initial utility entitlement v from any reporting strategy σ is

$$U\left(\{c_t^v, n_t^v\}, \sigma; \beta\right) \equiv \sum_{t=0}^{\infty} \sum_{\theta^t \in \Theta^{t+1}} \beta^t [u\left(c_t^v\left(\sigma^t(\theta^t)\right)\right) - h(n_t^v(\sigma^t(\theta^t))/\theta_t)] \Pr(\theta^t).$$

We say that allocation $\left(\{c_t^v, n_t^v\}_{t\geq 0, v\in V}, \{K_t\}_{t\geq 0}\right)$ and the set of strategies $\{\sigma^{*,v}\}$ is *incentive compatible* if

$$U\left(\{c_t^v, n_t^v\}, \sigma^{*,v}; \beta\right) \geq U\left(\{c_t^v, n_t^v\}, \sigma; \beta\right) \tag{36}$$

for all strategies σ and initial welfare entitlement v. It delivers utility v to all agents entitled to v, if in addition

$$U\left(\{c_t^v, n_t^v\}, \sigma^{*,v}; \beta\right) = v.$$

For a given initial distribution of entitlements ψ, we say that the allocation $\left(\{c_t^v, n_t^v\}, K_t\right)$ and the set of strategies $\{\sigma^{*,v}\}$ is *feasible* if: (i) it is incentive compatible; (ii)it delivers v to agents entitled to v; and (iii) the following sequence of resource constraints are satisfied:

$$\begin{aligned} &\int \sum_{\theta^t} c_t^v(\sigma^{*,v}(\theta^t)) \Pr(\theta^t)\, d\psi(v) \\ &\qquad \leq F\left(K_t, \int \sum_{\theta^t} n_t^v(\sigma^{*,v}(\theta^t)) \Pr(\theta^t)\, d\psi(v)\right) - K_{t+1} \qquad t = 0, 1, \ldots \end{aligned} \tag{37}$$

As in the i.i.d. case, a simple set of necessary and sufficient conditions for a feasible allocation and set of strategies to be the outcome of a sustainable equilibrium is that:

$$\int U\left(\{c_{t+s}^v, n_{t+s}^v\}_{s\geq 0}, \sigma^{*,v}; \beta\right) d\psi(v) \geq \hat{W}_t(K_t, \{n_t^v\}, \{\pi^{v,t}(.|m^t)\}) \qquad t = 0, 1 \ldots \tag{38}$$

where $\hat{W}_t(K_t, \{n_t^v\}, \{\pi^{v,t}(.|m^t)\})$ and $\pi^{v,t}(.|m^t)$ are defined below. The conditional probability $\pi^{v,t}(\theta^t|m^t)$ is defined by

$$\pi^{v,t}(\theta^t|m^t) \equiv \frac{\Pr\left((\hat{\theta}^t, \sigma^{*,v}(\hat{\theta}^t)) = (\theta^t, m^t)\right)}{\Pr\left(\sigma^{*,v}(\hat{\theta}^t) = m^t\right)}.$$

It is the probability distribution over Θ^t conditional on observing the sequence of messages m^t. When we consider a deviation at date t, each agent is characterized by an observable type which is composed of the initial promised value v and past sent messages m^t. The conditional probabilities $\pi^{v,t}(\theta^t|m^t)$ encode all the information about the agent's sequence of shocks θ^t up to date t that is revealed in public interim histories $\hat{h}_t$ in equilibrium. The agent's future probability distribution can then be inferred by the Chain Rule $\pi^{v,\infty}(\hat{\theta}^\infty \in \Theta^\infty|m^t) = \pi^{v,t}(\theta^t|m^t) \Pr\left(\hat{\theta}^\infty \in \Theta^\infty|\hat{\theta}^t = \theta^t\right)$.

The function $\hat{W}_t(K_t, \{n_t^v\}, \{\pi^{v,t}(.|m^t)\})$ is defined as follows:

$$\begin{aligned}\hat{W}_t(K_t, \{n_t^v\}, \{\pi^{v,t}(.|m^t)\}) \equiv \max_{K_{t+1}} \Big\{ & u(F(K_t, N_t) - K_{t+1}) \\ & - \sum_{\theta^t} \int h(n_t^v(\sigma^{*,v}(\theta^t))/\theta_t) \Pr(\theta^t), d\psi(v) + \beta W_{t+1}(K_{t+1}, \{\pi^{v,t}(.|m^t)\}) \Big\}\end{aligned} \tag{39}$$

where $W_{t+1}(K_{t+1}, \{\pi^{v,t}(.|m^t)\})$ represents the welfare associated with the worst sustainable equilibrium of the policy game at date $t+1$ given the information about agents types revealed in equilibrium and encoded in the conditional probability distribution $\pi^{v,t}(\theta^t|m^t)$.

We say that an allocation $\left(\{c_t^v, n_t^v\}, K_t\right)$ and a set of strategies $\{\sigma^{*,v}\}$ is *sustainable* if it is the outcome of a sustainable equilibrium. Equivalently, an allocation is sustainable if: (i) it is feasible; and (ii) it satisfies the sequence of credibility constraints (38).

Given ψ optimal sustainable allocation solves the following planning problem:

$$\min K_0 \tag{40}$$

subject to $\left(\{c_t^v, n_t^v\}, \{K_t\}\right)$ and $\{\sigma^{*,v}\}$ being a sustainable allocation.

6.2 The Revelation Principle Again

In this section, we make assumptions that guarantee that in the worst equilibrium, no information is used and all agents receive the same allocation. The idea is quite simple: we assume that labor is sufficiently beneficial that the worst equilibrium features no labor being provided.

It will prove useful to introduce the following definitions. We say that the worst payoff function W_{t+1} is *information-independent* if

$$W_{t+1}(K_{t+1}, \{\pi^{v,t}(.|m^t)\}) = W_{t+1}(K_{t+1}, \{\pi^t(.)\}).$$

When this condition holds, we use the simplified notation $W_{t+1}(K_{t+1})$. This is justified by the fact that $\{\pi^t(\cdot)\}$ is entirely exogenous. We use similar definitions and notations for $\hat{W}_t$.

Lemma 2. *Suppose that: (i) there is exists a maximal capital stock level $\bar{K} < \infty$; (ii) there exists $w > 0$ such that $\min_{N\in[0,\bar{n}],K\in[0,\bar{K}]} F_N(K,N) > w$; (iii) $u'(F(\bar{K},\bar{\theta}\bar{n}))w > h(\bar{n}/\underline{\theta})/\bar{n}$.. Then both $\hat{W}_t$ and W_{t+1} are information-independent. Moreover, the worst payoff function can be represented as the fixed point in a simple functional equation:*

$$W_t(K_t) = \max_{K_{t+1}} \{u(F(K_t,0) - K_{t+1}) - h(0) + \beta W_{t+1}(K_{t+1})\} \tag{41}$$

Moreover, $W_t(K_t)$ is nondecreasing and concave, and $\hat{W}_t(K_t,\{n_t^v\})$ is increasing, concave,and differentiable in K_t.

Lemma 2 proves that the worst equilibrium entails zero labor supply. As a result, no use is made of revealed information after a deviation has occurred. Moreover, no new information is revealed after a deviation and the continuation of the worst is the worst. Therefore, we can rely on the revelation principle on the equilibrium path just as we did for the i.i.d. All the the results we proved for the i.i.d. case can be extended to this case.

6.3 Failure of the Revelation Principle

In this section we tackle the general case where the revelation principle might fail. We show that a version of our main optimality result actually relies on the availability of a perturbation that does not affect the revelation of information over time. Thus, it can be performed for general mechanisms, regardless of the information being revealed in equilibrium.

A Class of Perturbations. For any sustainable allocation $(\{c_t^v, n_t^v\}, \{K_t\})$ and strategies $\{\sigma^{*,v}\}$, we associate a class of allocations $\Omega(\{\sigma^{*,v}\}, \{c_t^v, n_t^v\}, \{K_t\})$ as follows. It is the set of allocations $(\{c_t'^v, n_t'^v\}, \{K_t'\})$ that verify two properties: (i) they have the same labor allocation $n_t'^v = n_t^v$; (ii) they are sustainable using the set of strategies $\{\sigma^{*,v}\}$. This class of allocation preserves the same allocation for labor and the same strategy, so that it preserves the information revealed by individual histories. Hence, the conditional probabilities $\{\pi^{v,t}(.|m^t)\}$ are the same for all the allocations in $\Omega(\{\sigma^{*,v}\}, \{c_t^v, n_t^v\}, \{K_t\})$.

Suppose that the minimum in the planning problem (40) is attained for a sustainable allocation $(\{c_t^v, n_t^v\}, \{K_t\})$ and set of reporting strategies $\{\sigma^{*,v}\}$. Then $(\{c_t^v, n_t^v\}, \{K_t\})$ is also the optimum in the restricted planning problem

$$\min K_0 \tag{42}$$

subject to $\left(\{c_t^{\prime v}, n_t^{\prime v}\}, \{K_t^{\prime}\}\right) \in \Omega\left(\{\sigma^{*,v}\}, \{c_t^v, n_t^v\}, \{K_t\}\right)$.

First order conditions and implementation. In what follows, we make the extra assumption that $\hat{W}_t(K_t, \{n_t^v\}, \{\pi^{v,t}(.|m^t)\})$ is differentiable with respect to K_t. Then, the following first order conditions will hold:

$$\frac{\mu_{t+1}}{\mu_t}\beta F_K\left(K_{t+1}, N_{t+1}\right) - \frac{\nu_{t+1}}{\mu_t}\beta \hat{W}_{t+1,K}(K_{t+1}, \{n_{t+1}^v\}, \left\{\pi^{v,t+1}(.|m^{t+1})\right\}) = 1 \tag{43}$$

and
$$\frac{1}{u'(c^v(\sigma^{*v}(\theta^t)))} = \frac{\mu_{t+1}}{\mu_t}\mathbb{E}_{\sigma^{*,v}(\theta^t)}\left[\frac{1}{u'(c^v(\sigma^{*,v}(\theta^{t+1})))}\right] - \frac{\nu_{t+1}}{\mu_t}, \tag{44}$$

where $\mathbb{E}_{\sigma^{*,v}(\theta^t)}$ denotes the expectations operator, conditional on the event $\sigma^{*,v}(\hat{\theta}^t) = \sigma^{*,v}(\theta^t)$.

These first order conditions are the exact analogues of (30) and (31). Note however that the Inverse Euler Equation (44) is conditional on the information revealed along the equilibrium path – $\sigma^{*,v}(\theta^t)$ – and not on the information received by the agent θ^t. The former contains weakly less information than the latter.

We can then develop an implementation exactly as in subsection 5.3. The remarkable feature of our implementation is that it continues to work even when the agents and the planners have different information sets. This is because the implementation makes sure that the agent's Euler equation holds no matter what his reporting strategy his. The difference is that the average capital tax rate $\mathbb{E}_{\sigma^{*,v}(\theta^t)}\left[\tau(\sigma^{*,v}(\theta^{t+1}))\right]$ that we can characterize sharply is potentially different from the average capital tax $\mathbb{E}_{\theta^t}\left[\tau(\sigma^{*,v}(\theta^{t+1}))\right]$ faced by the agent under the implementation.

Just as in subsection 5.3, we can derive a formula for $\mathbb{E}_{\sigma^{*,v}(\theta^t)}\left[\tau(\sigma^{*,v}(\theta^{t+1}))\right]$:

$$\frac{\beta \hat{W}_{t+1,K}(K_{t+1}, \{n_{t+1}^v\}, \{\pi^{v,t+1}(.|m^{t+1})\}) - u'\left(c_t^v(\sigma^{*,v}(\theta^t))\right)}{\beta F_K\left(K_{t+1}, N_{t+1}\right)}\frac{\nu_{t+1}}{\mu_{t+1}}. \tag{45}$$

so that this average distortion is progressive, in the sense of being increasing in the level of current consumption $c_t^v(\sigma^{*,v}(\theta^t))$ assigned to this group.

It is important to understand more precisely how this result resembles and differs from the previous cases, where all information was revealed along the equilibrium path.

The key difference is that the conditional expectation is now taken with respect to all the information that the government possesses, which is the history of messages $\sigma^{*,v}(\theta^t)$. This is what determines current consumption. However, agents may possess more information, since they know their full history of shocks θ^t. This may allow agents to predict their future messages, which determine their future marginal tax rates. This means that the average tax on capital expected by the agent, computed as $\mathbb{E}_{\theta^t}\left[\tau(\sigma^{*,v}(\theta^{t+1}))\right]$, may vary with the true history θ^t in addition to the reported sequence of reports $\sigma^{*,v}(\theta^t)$. Since this implies that $\mathbb{E}_{\theta^t}\left[\tau(\sigma^{*,v}(\theta^{t+1}))\right]$ is not solely a function of current consumption, it makes little sense to ask whether it is progressive.

On the other hand, as equation (45) shows, progressivity continues to hold on average across groups of agents that are distinguishable from the government's point of view in period t.The government is less informed, but it still imparts progressivity on the dimension of average taxes that they control $\mathbb{E}_{\sigma^{*,v}(\theta^t)}\left[\tau(\sigma^{*,v}(\theta^{t+1}))\right]$.

7 Conclusion

The basic idea behind our result can be stated as follows: in settings where: (a) the credibility of future policies is of concern, and (b) credibility depends on keeping inequality in check, policies will be put into place to avoid the accumulation of inequality. A progressive tax on capital is one such policy.

Our stylized model delivered this sharp result in a transparent way. But the main mechanism behind the model, however, does not appear to be dependent on the particular simplifying assumptions we made.We conjecture that the progressivity of capital taxation is likely to be robust to a number of elements of the model and that perhaps the general logic extends in interesting directions to other policy instruments.

Appendix

A Proof of Proposition 1

Consider an allocation $c_0(\theta_0), c_1(\theta_0), n_0(\theta_0)$ that satisfies the assumptions of Proposition 1. The implicit capital tax for agent θ_0 is given by

$$\tau(\theta_0) = 1 - \frac{u'(c_0(\theta_0))}{\beta R u'(c_1(\theta_0))}.$$

We define the capital tax schedule T^k as the solution of the following ODE:

$$T^{k\prime}(Rk) = \tau\big(c_1^{-1}(Rk - T^k(Rk))\big)$$

where c_1^{-1} is the inverse of the function c_1, and $T^k(0) = 0$[14]. Note that since $\tau(\theta_0) < 1$, $Rk - T^k(Rk)$ is strictly increasing in k. Hence we can define the function $k(\theta_0)$, increasing in θ_0, as follows:

$$Rk(\theta_0) - T^k(Rk(\theta_0)) = c_1(\theta_0).$$

Define the labor income tax T^n so that

$$n_0(\theta_0) - c_0(\theta_0) - T^n(n_0(\theta_0)) = k(\theta_0).$$

Let $y(\theta_0) = n_0(\theta_0) - T^n(n_0(\theta_0))$.

Consider, for a given level of income y net of labor income tax, the following problem. Maximize

$$U(k, c_1; y) \equiv u(y - k) + \beta u(c_1) \tag{46}$$

subject to

$$c_1 = Rk - T^k(Rk).$$

By construction, $\big(k(\theta_0), c_1(\theta_0)\big)$ satisfies the first order conditions in (46) when $y = y(\theta_0)$.

Note that we have the following single crossing property:

$$\frac{\partial\big(-\frac{U_{c_1}}{U_k}\big)}{\partial y} > 0$$

Together with the fact that $y(\theta_0)$ and $c_1(\theta_0)$ are increasing in θ_0, this is enough to en-

[14] possible flat portions of $c_1(\theta_0)$ define discontinuous jumps

sure that for $y = y(\theta_0)$, $\big(k(\theta_0), c_1(\theta_0)\big)$ attains the maximum in (46). Hence an agent of type θ_0 who supplies $n(\theta_0')$ units of effective labor will optimally choose to consume $\big(c_0(\theta_0'), c_1(\theta_0')\big)$ when confronted with the taxes T^n and T^k. Since the original allocation is incentive compatible, working $n(\theta_0)$ and consuming $\big(c_0(\theta_0), c_1(\theta_0)\big)$ is the optimal choice for an agent of type θ_0. Therefore, the taxes T^n and T^k implement the allocation.

B Proof of Lemma 1

Let us denote by $\Gamma(K)$ the set of static income allocations $\{n_\theta\}$ such that there exists an incentive compatible and resource feasible allocation with initial capital given by K and first period income for type θ given by n_θ. Sargent and Ljungqvist (2004) provide a two step algorithm to determine the worst.

Consider the following program:

$$\tilde{W}(K) = \min_{\{n_\theta\}\in\Gamma(K)} \max_{K', c_\theta} \left\{ u\left(\sum c_\theta \Pr(\theta)\right) - \sum h(n_\theta/\theta) \Pr(\theta) + \beta \tilde{W}(K') \right\} \tag{47}$$

subject to

$$\sum c_\theta \Pr(\theta) \le F\left(K, \sum n_\theta \Pr(\theta)\right) - K'$$

This Bellman equation admits a unique bounded fixed point W^*. The worst equilibrium payoff, which is bounded, is necessarily greater than than the value of the unique bounded fixed point of (47): $W \ge W^*$.

We proceed in two stages: a characterization stage and a verification stage. In the characterization stage, characterize W^*. In the verification stage, we use our characterization to construct an equilibrium with a payoff given by W^*. Since $W \ge W^*$, this proves that $W = W^*$.

Characterization. Let us characterize the unique bounded fixed point of (47). The maximization over $\{c_\theta\}$ is very simple. The planner just equalizes consumption across agents. We are then led to study, for a given aggregate output N, the following problem:

$$\min_{\{n_\theta\}\in\Gamma(K)} \mathbb{E}[h(n_\theta/\theta)] \tag{48}$$

subject to $\mathbb{E}[n_\theta] = N$. A necessary condition for n_θ to be in $\Gamma(K)$ is that n_θ be increasing in θ. We can therefore study the relaxed planning problem

$$\min_{\{n_\theta\}} \mathbb{E}[h(n_\theta/\theta)]$$

subject to $\mathbb{E}[n_\theta] = N$ and n_θ increasing. This problem is concave. Moreover, we can verify that ignoring the monotonicity condition at θ leads to its violation. Hence the monotonicity constraint is binding for all θ: and $n_\theta = N$ for all θ. Clearly this static income allocation is in $\Gamma(K)$. Hence it also solves (48). Therefore, the unique bounded fixed point W^* of (47) also solves the following simplified program

$$W^*(K) = \min_{n\in[0,\bar{n}]} \max_{K'} \left\{u(F(K,n) - K') - \mathbb{E}\left[h(n/\theta)\right] + \beta W^*(K')\right\} \tag{49}$$

We denote by $n(K)$, $c(K) = F(K, n(K)) - K'(K)$ and $K'(K)$ the optimal policies in this program. Since W^* is bounded, the allocation that is constructed by iterating these policy functions delivers the payoff $W^*(K)$.

Verification. Consider the following strategies for the agents and the governments respectively, where all agents initially have the same v (we therefore drop the v superscripts): $n_t(H_t, \underline{h}_t) = n(K_t)$, $m_t(H_t, \underline{h}_t) = \theta_t$, $x_t^j(H_t)(\hat{h}_t) = 1$ if and only if $n_t \neq n(K_t)$, $c_t^j(H_t)(\hat{h}_t) = c(K_t)$, $K_{t+1}^j(H_t) = K'(K_t)$. These strategies define a sustainable equilibrium. Therefore, the unique bounded fixed point W^* of (47) is an equilibrium payoff.

Therefore, the worst payoff function W is equal to the unique bounded fixed point W^* of (47) and is characterized by (49). In words, all types are asked to work the same so that no incentives have to be provided: deviating from the worst leads to exactly the same allocation and payoffs as those that occur on the equilibrium path of the worst equilibrium. The continuation of the worst is the worst.

The Bellman operator in (49) maps concave and nondecreasing functions into concave and nondecreasing functions. It follows from this that W is nondecreasing and concave. An simple application of the Benveniste-Scheinkman theorem (see theorem 4.10 in Stokey et al., 1989) then proves that $\hat{W}$ is increasing, concave and differentiable.

C Proof of Lemma 2

The proof proceeds along similar lines as that of Lemma 1. We first develop a candidate representation for the worst and then proceed to a verification.

Consider the following dynamic program:

$$\tilde{W}_t(K_t, \left\{\pi^{v,t-1}(.|m^{t-1})\right\}) = \min_{(n_t^{v,m^{t-1}}(.),m_t^v(.))\in\Gamma(K_t,\left\{\pi^{v,t-1}(.|m^{t-1})\right\})} \max_{c_t^{v,m^{t-1}}(.),K_t+1}$$

$$\sum_{m^{t-1},\theta^t} \int \left[u(c_t^{v,m^{t-1}}(m_t^v(\theta^t))) - h(n_t^{v,m^{t-1}}(m_t^v(\theta^t)))/\theta_t\right] \pi^{v,t-1}(\theta^{t-1}|m^{t-1}) \Pr(\theta^t|\theta^{t-1}) d\psi(v)$$
$$+ \beta\tilde{W}_{t+1}(K_{t+1}, \{\pi^{v,t}(.|m^t)\}) \quad (50)$$

subject to

$$\sum_{m^{t-1},\theta^t} \int c_t^{v,m^{t-1}}(m_t^v(\theta^t))\pi^{v,t-1}(\theta^{t-1}|m^{t-1}) \Pr(\theta^t|\theta^{t-1})d\psi(v) + K_{t+1}$$
$$\leq F\left(K_t, \sum_{m^{t-1},\theta^t} \int n_t^{v,m^{t-1}}(m_t^v(\theta^t))\pi^{v,t-1}(\theta^{t-1}|m^{t-1}) \Pr(\theta^t|\theta^{t-1})d\psi(v)\right)$$

where $\Gamma(K_t, \{\pi^{v,t-1}(.|m^{t-1})\})$ is the set of labor allocations $n_t^{v,m^{t-1}}(.)$ and date t messages m_t^v that are a component of incentive compatible allocations given capital K_t and information $\{\pi^{v,t-1}(.|m^{t-1})\}$.

The Bellman equation can be used to map any sequence of functions $\{\tilde{W}_t\}$ into a new sequence, call it $\{\tilde{\tilde{W}}_t\}$, where $\tilde{\tilde{W}}_t$ is given by the right hand side of equation (50) using $\tilde{W}_t$. Denote by T the operator that maps the space of sequences of functions $\{\tilde{W}_t\}$ into itself in this way. The Bellman equation (50) represents a fixed point of this operator.

We first define the space of uniformly bounded sequences of functions. These are the sequences of functions $\{\tilde{W}_t\}$ that satisfy the following property:

$$\max_{t,K_t,\left\{\pi^{v,t-1}(.|m^{t-1})\right\}} |\tilde{W}_t(K_t, \left\{\pi^{v,t-1}(.|m^{t-1})\right\})| < \infty.$$

Because by Assumption 2, u and h are bounded, T preserve the space of uniformly bounded sequences of functions. This space is a Banach space when endowed with the supremum norm. We therefore know that T has a unique fixed point $\{W_t^*\}$ in the space of uniformly bounded sequences of functions. Moreover, given any uniformly bounded sequence of functions $\{\tilde{W}_t\}$, this unique fixed point is the limit of $T^n(\{\tilde{W}_t\})$ when n goes to ∞.

The worst payoff $\{W_t\}$ is uniformly bounded by $M/(1-\beta)$. Moreover, for all t, we necessarily have $W_t \geq W_t^*$.

We break-up the proof into two parts: a characterization stage, and a verification stage.

In the characterization stage, we will characterize $\{W_t^*\}$. We will use the properties of (50), Assumption 2, and the assumptions in Lemma 2. In the verification stage, we use our characterization to construct an equilibrium with a payoff given by W_t^*. Since $W_t \geq W_t^*$ for all t, this proves that $W_t = W_t^*$ for all t.

Characterization. We first prove that the Bellman operator in (50) maps sequences the space of information independent functions into itself. Suppose that $\tilde{W}_{t+1}(K_{t+1}, \{\pi^{v,t}(.|m^t)\})$ is information independent. Then maximizing over $c_t^{v,m^{t-1}}(.)$ clearly implies that $c_t^{v,m^{t-1}}(m_t^v(\theta^t))$ is equalized for all agents. We are then left with a constrained minimization over $n_t^{v,m^{t-1}}(.)$ and $m_t^v(.)$ of a concave objective function. The solution is of the bang-bang type and involves setting, for each agent, $n_t^{v,m^{t-1}}(m_t(\theta^t))$ to either 0 or $\bar{n}$.

As an intermediate step, we will prove that for all $n > 0$, the min-max value of the per period objective function that is reached when

$$\sum_{m^{t-1},\theta^t} \int n_t^{v,m^{t-1}}(m_t(\theta^t)) d\psi(v) \pi^{v,t-1}(\theta^{t-1}|m^{t-1}) \Pr(\theta^t|\theta^{t-1}) = N > 0$$

and K_{t+1} is invested is higher than when $N = 0$. Indeed, for all $N \geq 0$, this value is greater than $u(F(K_t, N) - K_{t+1}) - h(\bar{n}/\underline{\theta})N/\bar{n}$ with an equality for $N = 0$. The derivative of this function is greater than $u'(F(\bar{K}, \bar{n}\bar{\theta}))w - h(\bar{n}/\underline{\theta})/\bar{n}$ which is positive under assumptions (i), (ii), (iii). This concludes the proof of the intermediate step.

Hence the solution involves setting $n_t^{v,m^{t-1}}(m_t(\theta^t)) = 0$ and $c_t^{v,m^{t-1}}(m_t^v(\theta^t)) = F(K_t, 0) - K_{t+1}$. This implies that the iterated sequence $T\left(\{\tilde{W}_t\}\right)$ is information independent and concludes the proof that the set space of sequences of information independent functions is preserved by the Bellman operator T underlying (50).

Next we observe that there necessarily exists an information independent, uniformly bounded solution to (50). In fact, consider any uniformly bounded sequence of information independent functions $\{\tilde{W}_t\}$. For example, one can consider the sequence of functions that is uniformly equal to 0. We know that $T^n\left(\{\tilde{W}_t\}\right)$ converges to the unique fixed point of T in the space of uniformly bounded sequences of functions when n goes to ∞. Since the space of uniformly bounded sequence of information independent functions is closed and preserved by T, this shows that the unique fixed point of T in the space of uniformly bounded sequences of functions, given by the limit when n goes ∞ of $T^n\left(\{\tilde{W}_t\}\right)$, is information independent.

We conclude that the unique uniformly bounded fixed point of (50) $\{W_t^*\}$ is information independent. Moreover, the optimal policies in (50) involve setting $n_t^{v,m^{t-1}}(m_t(\theta^t)) = 0$ and $c_t^{v,m^{t-1}}(m_t^v(\theta^t)) = F(K_t, 0) - K_{t+1}$. And since $\{W_t^*\}$ is uniformly bounded, the allocation that is constructed by iterating these policy functions delivers the payoff $W_t^*(K)$.

Verification. The verification stage follows naturally from the characterization stage, exactly as in Lemma 1. Since the same allocation is given to every agent, no incentives have to be provided, deviating from the worst leads to exactly the same allocation and payoffs as those that occur on the equilibrium path of the worst equilibrium and the continuation of the worst is the worst.

References

Acemoglu, Daron, Michael Golosov, and Aleh Tsyvinski, "Dynamic Mirrlees Taxation under Political Economy Constraints," 2007. mimeo.

_, _, and _, "Political Economy of Mechanisms," *Econometrica*, 05 2008, *76* (3), 619–641.

Alesina, Alberto and Danny Rodrick, "Distributive Politics and Economic Growth," *Quarterly Journal of Economics*, 1994, *109*, 465–490.

Atkeson, Andrew and Robert E., Jr. Lucas, "On Efficient Distribution with Private Information," *Review of Economic Studies*, 1992, *v59* (3), 427–453.

Atkinson, Anthony B. and Joseph E. Stiglitz, "The Design of Tax Structure: Direct vs. Indirect Taxation," *Journal of Public Economics*, 1976, *6*, 55–75.

Benabou, Roland, "Tax and Education Policy in a Heterogeneous-Agent Economy: What Levels of Redistribution Maximize Growth and Efficiency?," *Econometrica*, March 2002, *70* (2), 481–517.

Benhabib, Jess and Aldo Rustichini, "Social Conflict and Growth," *Journal of Economic Growth*, 1996, *1*, 125–139.

Bertola, Guiseppe, "Factor Shares and Savings in Endogenous Growth," *American Economic Review*, 1993, *83*, 1184–1199.

Besley, Timothy and Stephen Coate, "An Economic Model of Representative Democracy," *The Quarterly Journal of Economics*, February 1997, *112* (1), 85–114.

Bisin, Alberto and Adriano A. Rampini, "Markets as beneficial constraints on the government," *Journal of Public Economics*, May 2006, *90* (4-5), 601–629.

Chamley, Christophe, "Optimal Taxation of Capital in General Equilibrium," *Econometrica*, 1986, *54*, 607–22.

Chari, V.V. and Patrick J. Kehoe, "Sustainable Plans," *Journal of Political Economy*, 1990, *98*, 616–636.

Coughlin, Peter J., "Elections and income redistribution," *Public Choice*, 1986, *50* (1–2), 27–91.

Farhi, Emmanuel and Iván Werning, "Inequality and Social Discounting," *Journal of Political Economy*, 2007, *115(3)*, 365–402.

— **and** — , "Progressive Estate Taxation," 2008. mimeo.

Fischer, Stanley, "Dynamic Inconsistency, Cooperation, and the Benevolent Dissembling Government," *Journal of Economic Dynamics and Control*, 1980, *2*, 93–107.

Judd, Kenneth L., "Redistributive Taxation in a Perfect Foresight Model," *Journal of Public Economics*, 1985, *28*, 59–83.

Klein, Paul and Jose-Victor Rios-Rull, "Time-Consistent Optimal Fiscal Policy," *International Economi Review*, 2003, pp. 1217–1245.

Kocherlakota, Narayana, "Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation," *Econometrica*, 2005, *73*, 1587–1621.

Kotlikoff, Larry, Torsten Persson, and Lars Svensson, "Social Contracts as Assets: a Possible Solution to the Time-Inconsistency Problem," *American Economic Review*, 1988, *78*, 662–677.

Kydland, Finn and Edward C. Prescott, "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 1977, *85*, 473–491.

Lindbeck, Assar and Joerg Weibull, "Balanced-budget Redistribution as the Outcome of Political Competition," *Public Choice*, 1987, pp. 273–97.

Mirrlees, James A., "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 1971, *38* (2), 175–208.

Osborne, Martin J. and Al Slivinski, "A Model of Political Competition with Citizen-Candidates," *The Quarterly Journal of Economics*, February 1996, *111* (1), 65–96.

Persson, Torsten and Guido Tabellini, "Is Inequality Harmful for Growth," *American Economic Review*, 1994, *84*, 600–621.

— **and** — , "Representative Democracy and Capital Taxation," *Journal of Public Economics*, 1994, *55*, 53–70.

Phelan, Christopher, "Opportunity and Social Mobility," *Review of Economic Studies*, 2006, *73* (2), 487–505.

— **and Ennio Stacchetti**, "Sequential Equilibria in a Ramsey Tax Model," *Econometrica*, November 2001, *69* (6), 1491–1518.

Saez, Emmanuel., "Optimal Progressive Capital Income Taxes in the Infinite Horizon," 2002. NBER WP9046.

Sargent, Thomas J. and Lars. Ljungqvist, *Recursive Macroeconomic Theory*, MIT Press, 2004.

Sleet, Christopher and Sevin Yeltekin, "Credibility and Endogenous Social Discounting," *Review of Economic Dynamics*, 2006, *9*, 410–437.

_ and _ , "Optimal Growth and Inequality in Uncommitted Societies," 2008. mimeo.

_ and _ , "Politically Credible Social Insurance," *Journal of Monetary Economics*, 2008, *55*, 129–151.

Stokey, Nancy L., Robert E., Jr. Lucas, and Edward C. Prescott, *Recursive Methods in Economic Dynamics*, Harvard University Press, 1989.