

Federal Reserve Bank of Minneapolis
Research Department Staff Report 122

June 1989

SUSTAINABLE PLANS

V. V. Chari

Patrick J. Kehoe*

Federal Reserve Bank of Minneapolis
and University of Minnesota

ABSTRACT

We propose a definition of time consistent policy for infinite horizon economies with competitive private agents. Allocations and policies are defined as functions of the history of past policies. A sustainable equilibrium is a sequence of history-contingent policies and allocations that satisfy certain sequential rationality conditions for the government and for private agents. We provide a complete characterization of the sustainable equilibrium outcomes for a variant of Fischer's (1980) model of capital taxation. We also relate our work to recent developments in the theory of repeated games.

*Kehoe acknowledges financial support from the National Science Foundation and the Sloan Foundation.

The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

I. Introduction

This paper develops an equilibrium concept and a technique for analyzing equilibria in a representative agent economy where the government chooses policy sequentially and private allocations are determined in a competitive equilibrium. This work builds on three distinct literatures. First, we extend the analysis of policy design in dynamic general equilibrium models without commitment technologies (see Kydland and Prescott 1977, 1980; Calvo 1978; Fischer 1980; and Lucas and Stokey 1983). Second, we build on ideas developed in the theory of repeated games, particularly in the oligopoly literature (see Friedman 1971, Fudenberg and Maskin 1986, and Abreu 1988). Third, we build on the recent literature on time consistency in macroeconomic models (see Barro and Gordon 1983 and Rogoff 1987).

The analysis of policy design in models without commitment in turn builds on the standard framework for policy design in general equilibrium models with commitment--a framework stemming from Ramsey (1927). Ramsey studied a static, representative consumer economy with many goods. In that economy, a government requires fixed amounts of these goods, which are purchased at market prices and financed by proportional excise taxes. Given the excise taxes, prices and quantities are determined in a competitive equilibrium. The government's problem is to choose tax rates to maximize the welfare of the representative consumer over the resulting competitive equilibrium allocations. An optimal policy together with the resulting competitive equilibrium is called a Ramsey equilibrium. It is straightforward to extend Ramsey's formulation to dynamic models by reinterpreting the goods as date-contingent commodities. In this dynamic context the Ramsey equilibrium can be interpreted as a one-time choice of policy with consumers making decisions sequentially. This formulation of the policy design problem is appropriate in

environments where societies have a commitment technology to bind the actions of future governments.

In many situations, however, it is more appropriate to think of policies as being chosen at each date and society as having no ability to commit to future policies. For an environment without commitment, a solution to the policy design problem requires that policies be sequentially rational, in the sense that they maximize welfare at each date. In turn, rationality on the part of private agents requires that they forecast future policies as being optimal for society. In a finite horizon economy, we can use backward induction (as in Kydland and Prescott 1977) to solve the forecasting problem. We compute an optimal policy for the last period for every state. This yields a policy function describing how policies will be chosen in every state. Taking this policy function and the resulting competitive equilibrium function as given, we compute optimal policies for the next-to-last period, and so on.¹

For infinite horizon problems, however, this procedure is no longer available. So we borrow from the game theory tradition to solve the forecasting problem imposed by sequential rationality. We allow policies and allocations to depend on the history of past decisions by the government. Thus, policies, allocations, and prices are defined as history-contingent functions. We need to break from the general equilibrium approach of defining equilibria as functions contingent only on exogenous events because the government is not a price taker. Since the government must predict how private agents will respond to its decisions, private allocations and prices are defined as functions of the history of the government's policies. To satisfy the requirement of sequential rationality, various definitions of equilibrium have been offered in the literature (see, for example, Lucas and Stokey

1983). To distinguish our definition from others, we call a sequence of history-contingent policy rules, allocations, and prices that together satisfy sequential rationality a sustainable equilibrium.

To illustrate the notion of sustainable equilibrium, we use an infinitely repeated version of Fischer's (1980) model of capital taxation. In the model, each period is divided into two distinct stages. In the first stage, consumers have a given endowment and must decide how much of it to save and to consume. In the second stage, they decide how much labor to supply and how much to consume. At the second stage, the government sets tax rates on capital and labor supply to finance an exogenously specified amount of government consumption.

Our results show that there is a unique sustainable equilibrium when the horizon is finite; but when the horizon is infinite, the set of sustainable equilibria is large and difficult to characterize. Fortunately, the set of allocations and policies resulting from a sustainable equilibrium is easy to characterize. We show that an arbitrary policy and allocation sequence is sustainable if and only if two conditions are met. First, the sequence must be a date-zero competitive equilibrium. Second, it must be the outcome of a particular sustainable equilibrium called the revert-to-autarky equilibrium. This equilibrium specifies that both private agents and government continue with the given sequence if there has been no past deviation from it. If, however, there has been a deviation, then the equilibrium specifies that all agents revert to the finite horizon sustainable equilibrium forever. This result allows us to characterize the set of sustainable policies and allocations by a simple set of inequalities. We use these inequalities to show that with sufficiently little discounting, even the Ramsey allocations can be supported by a sustainable equilibrium.

One question that arises in the analysis is, How does the notion of time consistency considered here relate to the notion of perfection in game theory? We analyze this matter by mapping our economy, with budget constraints and competitive private agents, into an anonymous game. We then show that the symmetric perfect Bayesian equilibrium outcomes of the game coincide with the sustainable outcomes of our economy.²

Before turning to the model, we briefly summarize how our work relates to the three distinct literatures on which it builds: In the literature on policy design in equilibrium models, the decisions of private agents and the government depend, at most, on exogenous events; in our formulation these decisions can depend on the complete history of past policy. This difference explains why in our formulation the Ramsey plan may be dynamically consistent, whereas in other formulations it typically is not. From game theory we borrow the idea of history-contingent decisions and from the literature on repeated games, Abreu's (1986, 1988) technique of using the worst equilibrium to characterize the entire set of equilibrium outcomes. In that literature, however, the games consist of a finite number of players, each of whom has strategic power. In ours, there is one large agent, the government, and a continuum of competitive private agents. For this reason, standard results from the theory of repeated games cannot be applied. Finally, in the literature on macroeconomic policy games it often appears as if the government plays a game against a coalition of private agents, who may have different objectives than the government. In our model the government maximizes the welfare of private agents, who behave competitively.

II. A One-Period Economy

Consider a one-period economy along the lines of Fischer (1980). The economy contains a large number of identical consumers and a government. There is a

linear production technology, for which the marginal product of capital is a constant $R > 1$ and for which the marginal product of labor is 1. Consumers make decisions at two distinct points in time: the first stage and the second stage. At the first stage, consumers are endowed with ω units of the consumption good out of which they consume c_1 and save k . At the second stage, they consume c_2 and work ℓ units. Second-stage income, net of taxes, is $(1-\delta)Rk + (1-\tau)\ell$, where δ and τ denote the tax rates on capital and labor. For simplicity we assume that first-stage consumption and second-stage consumption are perfect substitutes. A consumer, confronted with tax rates δ and τ , chooses $(c_1, k; c_2, \ell)$ to solve

$$\max U(c_1 + c_2, \ell) \tag{1}$$

subject to

$$c_1 + k \leq \omega$$

$$c_2 \leq (1-\delta)Rk + (1-\tau)\ell.$$

Notice that when $(1-\delta)R = 1$, the consumer is indifferent between saving the entire endowment ω and not saving at all.

The government sets proportional tax rates on capital and labor income to finance an exogenously given amount of second-stage per capita government spending g . The government's budget constraint is

$$g \leq \delta Rk + \tau \ell. \tag{2}$$

We assume, throughout, that

$$g > (R-1)\omega. \tag{3}$$

This assumption implies that in any equilibrium, the government must tax labor. We also assume that it is feasible to finance government spending with only a tax on labor.

In what follows, we let $x = (x_1, x_2)$ denote each consumer's first- and second-stage allocations, where $x_1 = (c_1, k)$ and $x_2 = (c_2, \ell)$.

A. Commitment

Consider, first, the case where an institution or commitment technology is available through which the government can bind itself to a tax policy at the beginning of the period. This technology is formalized by the following sequence of events: the government sets tax rates, then consumers make their first- and second-stage decisions.

In this setup, a policy for the government is a pair of numbers $\pi = (\delta, \tau)$ with $0 \leq \delta, \tau \leq 1$. Since the government needs to predict how consumers will respond to its policies, consumer behavior is described by rules that associate government policies with allocations. Formally, an allocation rule is a function f that maps government policies into allocations. For any allocation rule f , the problem of the government is to choose some policy π that maximizes consumer utility³

$$V(\pi, f) = U(c_1(\pi) + c_2(\pi), \ell(\pi)) \quad (4)$$

subject to the budget constraint

$$g \leq \delta Rk(\pi) + \tau \ell(\pi), \quad (5)$$

where $f(\pi) = (c_1(\pi), k(\pi), c_2(\pi), \ell(\pi))$. We then have the following definition: A Ramsey equilibrium is a policy π and an allocation rule f that satisfy

- Government maximization. The policy π solves the government's problem (4).
- Rationality by consumers. For every policy π' , the allocation $f(\pi')$ solves the consumer's problem (1).

In the Ramsey equilibrium (π, f) , some particular allocation will be realized, namely, $x = f(\pi)$. We call the Ramsey policy together with this allocation the Ramsey outcome and denote it by (π^f, x^f) .

Proposition 1. The Ramsey Outcome. The Ramsey outcome (π^f, x^f) has first-stage allocations $c_1^f = 0$ and $k^f = \omega$, and a capital tax rate $\delta^f = (R-1)/R$. Second-stage allocations c_2^f and l^f and the labor tax rate τ^f solve

$$U^f = \max U(c_2, l) \tag{6}$$

subject to

$$c_2 \leq \omega + (1-\tau)l$$

$$-U_l / U_c = (1-\tau)$$

$$g \leq (R-1)\omega + \tau l.$$

Proof. Consider the allocation rule for capital $k(\pi)$. If the tax rate on capital is strictly greater than δ^f , then consumers save zero; if $\delta = \delta^f$, consumers are indifferent among all levels of savings; and if $\delta < \delta^f$, consumers save their entire endowments. For now, assume that when $\delta = \delta^f$, the allocation rule specifies $k(\pi) = \omega$. The tax on capital acts like a lump-sum tax when it is selected at any level less than or equal to δ^f . Clearly, it is optimal to raise as much revenue as possible from this tax. Since $g > (R-1)\omega$, government spending is greater than the maximal possible revenue from this capital tax, namely $\delta^f R\omega$; therefore, it is optimal to set δ so that

$(1-\delta)R = 1$. Faced with this tax, consumers save their entire endowments. Given these facts, the optimal tax problem reduces to choosing c_2 , λ , and τ to solve (6).

Now suppose that when $\delta = \delta^f$, $k(\pi)$ equals some number α with $0 < \alpha < \omega$. We claim that such an allocation rule is not consistent with an equilibrium. For such an allocation rule, the government's objective function is discontinuous at $\delta = \delta^f$, and no optimal policy exists. With such a rule, the government can increase its utility by setting δ arbitrarily close to but smaller than δ^f and by setting τ close enough to τ^f so that the labor tax can raise the rest of the needed revenue. Consumers now choose to save their entire endowments, and the government is strictly better off. Thus, such a specification of $k(\pi)$ is inconsistent with equilibrium. Q.E.D.

B. No Commitment

For the equilibrium just described to be viable, the government needs to have some special commitment technology by which it can bind itself to specific policies. When no such technology is available, the notion of what constitutes an optimal policy is quite different. Formally, the lack of commitment is modeled by assuming that the government does not set policy until after consumers have made their first-stage decisions. That is, events are timed as follows: consumers make their first-stage decisions, then the government sets tax policies, and then consumers make their second-stage decisions.

In this setup, consumers' second-stage decisions are allowed to depend on the tax policy. Thus, a consumer's second-stage decisions are described by a pair of functions $f_2(\pi) = (c_2(\pi), \lambda(\pi))$. We call f_2 a second-stage allocation rule to distinguish it from a particular second-stage allocation x_2 . Consumers' allocations over the two stages are described by a pair (x_1, f_2) . A policy for the government is again a pair of tax rates

$\pi = (\delta, \tau)$. An equilibrium is a policy plan and a set of allocation rules that satisfy certain conditions that are developed recursively. First, we consider the second-stage problem faced by consumers for a given tax policy. We then consider the problem faced by the government. Finally, we consider the first-stage problem faced by consumers. For the second-stage problem, we say the pair (x_1, f_2) is rational for consumers at the second stage if for every history π' of tax policy, the second-stage decisions $f_2(\pi')$ solve

$$\max_{c_2, \ell} U(c_1 + c_2, \ell) \quad (7)$$

subject to

$$c_2 \leq R(1 - \delta')k + (1 - \tau')\ell.$$

Next, consider the government's problem. Given the decisions x_1 and knowing that second-stage decisions are selected according to the rule f_2 , the government selects a policy π that maximizes consumer welfare subject to its budget constraint. Thus, given x_1 and f_2 , a policy π is rational for the government if it maximizes

$$V(\pi, x_1, f_2) = U(c_1 + c_2(\pi), \ell(\pi)) \quad (8)$$

subject to

$$g \leq \delta Rk + \tau \ell(\pi).$$

Finally, consider the consumer's problem at the first stage. Each consumer picks an individual allocation for the first stage, $x_1 = (c_1, k)$, together with an allocation rule f_2 for taking actions at the second stage. Each consumer takes as given that the future policy is set according to π . Then, for a given policy π , an allocation plan (x_1, f_2) is rational for con-

sumers at the first stage if x_1 and $f_2(\pi)$ solve

$$\max U(c_1 + c_2(\pi), l(\pi))$$

subject to

$$c_1 \leq \omega - k$$

$$c_2(\pi) \leq (1-\delta)Rk + (1-\tau)l(\pi).$$

We have now defined the consumer's and the government's problems recursively. Combining these gives an equilibrium with sequential rationality built in for both the private agents and the government. As a result, we say the equilibrium is sustainable. Formally, a sustainable equilibrium is a triple (π, x_1, f_2) that satisfies⁴

- Sequential rationality by consumers. The pair (x_1, f_2) is rational for consumers at the first stage, and for every history π' the allocation rule $f_2(\pi')$ is rational for consumers at the second stage.
- Sequential rationality by the government. Given x_1 and f_2 , the policy π is rational for the government.

In Section IV we use the sustainable equilibrium from this one-period model to help characterize the set of sustainable policies and allocations for the infinite horizon model. To help distinguish these two equilibria, we call the sustainable equilibrium for the one-period model the autarky equilibrium and denote it by superscript a. We then have

Proposition 2. Autarky Equilibrium. In any autarky equilibrium (π^a, x_1^a, f_2^a) , the first-stage allocations are $c_1^a = \omega$ and $k^a = 0$. The capital tax δ^a satisfies $(1-\delta^a)R \leq 1$. The labor tax τ^a is given in the solution to

the problem

$$\max_{\tau, \ell} U(\omega + (1-\tau)\ell, \ell) \quad (9)$$

subject to

$$-U_{\ell} / U_c = (1-\tau)$$

$$g \leq \tau \ell.$$

The allocation rule f_2^a is defined as follows: for any π' , the allocations $f_2^a(\pi')$ solve

$$\max_{c_2, \ell} U(\omega + c_2, \ell) \quad (10)$$

subject to

$$c_2 \leq (1-\tau')\ell.$$

Proof. We first verify that any such triple (π^a, x_1^a, f_2^a) is an equilibrium. Consider the second-stage allocation rule. By definition this rule must satisfy (7). With $k^a = 0$ and $c_1^a = \omega$, the problem (7) reduces to (10), and f_2^a is rational for consumers. Then consider the first-stage allocations. If $(1-\delta)R < 1$, it is optimal for consumers to save nothing and consume their endowments. If $(1-\delta)R = 1$, consumers are indifferent among all savings levels, including saving nothing. Thus, x_1^a is rational for consumers for any δ satisfying $(1-\delta)R \leq 1$.

Next, consider the problem of the government. With $k^a = 0$, the government is indifferent among all capital tax rates, including those rates satisfying $(1-\delta)R \leq 1$. Thus, any such δ is rational for the government. Notice that any δ satisfying $(1-\delta)R > 1$ would be inconsistent with consumers' saving decisions. Finally, consider the labor tax τ^a . Putting together the

fact that savings are zero and that f_2^a solves (10), we can reduce the optimal tax problem of the government to (9). Thus, the specified (π^a, x_1^a, f_2^a) is a sustainable equilibrium.

It should be clear that no other equilibria are sustainable. Suppose, for example, there were an equilibrium with $k > 0$. An optimal tax policy would raise as much revenue as possible from this capital; thus, δ would be set equal to $\min(g/Rk, 1)$. Since we have assumed $g > (R-1)\omega$, this capital tax would satisfy $(1-\delta)R < 1$. For such a tax rate, however, it would be optimal for consumers to set $k = 0$. Thus, there can be no equilibria with positive savings. Q.E.D.

In an autarky equilibrium, the realized second-stage allocations are $f_2^a(\pi^a)$. We let the realized value of utility be denoted U^a , where $U^a = U(\omega + (1-\tau^a)k^a(\pi^a), l^a(\pi^a))$. Notice that while there is a unique value of utility realized in an autarky equilibrium, there is a range of capital tax rates consistent with equilibrium. For concreteness, we define $\delta^a = 1$.

It is immediate to show that for this economy, the value of an institution that allows the government to commit to its policy is strictly positive:

Lemma 1. The level of utility in the Ramsey equilibrium is strictly greater than the level of utility in a sustainable equilibrium; that is, $U^f > U^a$.

Proof. Comparing the Ramsey tax problem (6) to the autarky tax problem (9), it is clear that the autarky tax problem is simply the Ramsey problem with a larger level of g , namely, $\hat{g} = g + (R-1)\omega$. Since the utility function is strictly increasing in both arguments, the maximized value of the Ramsey problem is strictly greater than that of the autarky problem. Q.E.D.

III. An Infinite Horizon Economy

Consider an infinite repetition of the one-period economy described in Section II. Each period is indexed by t , for $t = 0, 1, \dots$, and each consumer's utility is the sum of each period's utility discounted at rate β . We assume throughout that government spending is constant and satisfies $g > (R-1)\omega$. We consider two versions of this economy, one with commitment and one without. In the commitment version, the government sets a sequence of tax rates once and for all at the beginning of time. Consumers then choose a sequence of allocations for all time. In contrast, in the no-commitment version the government and the consumers make decisions sequentially. For each period the timing of decisions is the same as in the one-period problem. We then compare the optimal policies for these two versions.

A. Commitment

Consider, first, the commitment economy in which the government and the consumers make their decisions at the beginning of time. In particular, let $\pi = (\pi_0, \pi_1, \dots)$ denote an infinite sequence of tax rates starting at time zero. For each period t , let $x_t = (x_{1t}, x_{2t})$ be the allocations for the first and second stages of t , and let $x = (x_0, x_1, \dots)$ denote the infinite sequence of such allocations. For this environment, a policy for the government is an infinite sequence of tax rates π . An allocation rule is a sequence of functions $f = (f_0, f_1, \dots)$ that maps government policies into sequences of allocations. We define a Ramsey equilibrium as a policy π and an allocation rule f that satisfy

- Government maximization. The policy π maximizes

$$V(\pi, f) = \sum_{t=0}^{\infty} \beta^t U(c_{1t}(\pi) + c_{2t}(\pi), k_t(\pi))$$

subject to $g \leq \delta_t R k_t(\pi) + \tau_t k_t(\pi)$ for $t = 0, \dots, \infty$.

- Rationality by consumers. For every policy π' , the allocation rule $f(\pi')$ solves

$$\max \sum_{t=0}^{\infty} \beta^t U(c_{1t} + c_{2t}, \lambda_t)$$

subject to $c_{1t} \leq \omega - k_t$ and $c_{2t} \leq (1-\delta'_t)Rk_t + (1-\tau'_t)\lambda_t$.

We denote a Ramsey equilibrium by the pair (π, f) . It immediately follows that the Ramsey equilibrium for this infinite horizon model is simply the Ramsey equilibrium for the one-period model repeated infinitely many times. Notice that any Ramsey equilibrium (π, f) induces some particular outcome (π, x) , where $x = f(\pi)$. By construction, this outcome satisfies consumer maximization and the sequence of government budget constraints. Hence, this outcome is some specific date-0 competitive equilibrium. More generally, we say a pair of sequences (π, x) is a date-0 competitive equilibrium if it satisfies consumer maximization and the sequence of government budget constraints (but not necessarily government maximization). (In Section IV we characterize conditions under which a date-0 competitive equilibrium can be supported by a sustainable equilibrium.)

B. No Commitment

As in the one-period economy, the lack of a commitment technology is formally modeled by assuming that the timing scheme is as follows: for each period t , consumers make their first-stage decisions, then the government sets current tax rates, and then consumers make their second-stage decisions.

In each period the consumers and the government can vary their decisions, depending on the history of government policies up to the point when the decision is made. At the first stage of period t , each consumer chooses a first-stage allocation as a function of the history $h_{t-1} = (\pi_s | s=0, \dots, t-1)$ and a contingency plan for setting future actions for all

possible histories. The first-stage consumption and savings decisions are denoted by $f_{1t}(h_{t-1})$. After the first-stage decisions have been made, the government sets time- t tax rates as a function of the current history h_{t-1} and chooses a contingency plan for setting future tax rates as a function of any future aggregate history it may face. Let $\sigma_t(h_{t-1})$ denote the time- t tax rates chosen by the government when faced with history h_{t-1} . At the second stage of t , an individual's history is $h_t = h_{t-1} \cup (\pi_t)$. Faced with h_t , consumers choose a second-stage allocation $f_{2t}(h_t)$ and a contingency plan for setting future actions for all possible future histories.

To define a sustainable equilibrium, we now need to explain how policy plans induce future histories. We let σ^t denote a policy plan for the government from time t onward; that is, $\sigma^t = (\sigma_t, \sigma_{t+1}, \dots)$ is a sequence of policy rules from time t onward. For any policy plan $\sigma = (\sigma_0, \sigma_1, \dots)$ and any time t , we call σ^t the continuation of σ . Let f^t denote the corresponding objects for the allocation rules. Given a history h_{t-1} , the policy plan σ^t induces future histories by $h_t = (h_{t-1}, \sigma_t(h_{t-1}))$, and so on.

Consider the second stage of period t . Given some symmetric history h_{2t} , an individual consumer chooses a contingency plan (f_{2t}, f^{t+1}) . Each consumer takes as given that future policies will evolve according to the histories induced by σ^{t+1} . Formally, given a policy plan σ , the allocation rule f is rational for consumers at the second stage after the history h_t if the continuation of the rule maximizes

$$\sum_{s=t}^{\infty} \beta^{s-t} U(c_{1s}(h_{s-1}) + c_{2s}(h_s), l_s(h_s))$$

subject to

$$c_{1s}(h_{s-1}) \leq \omega - k_s(h_{s-1})$$

$$c_{2s}(h_s) \leq (1-\delta_s(h_{s-1}))Rk_s(h_{s-1}) + (1-\tau_s(h_{s-1}))l_s(h_s),$$

where for all $s \geq t$, the future histories are induced by σ^{t+1} . For any history h_{t-1} , rationality at the first stage of t is defined in a similar fashion.

Next, consider the situation of the government in period t . Given some history h_{t-1} and the fact that allocations evolve according to f , the government chooses a policy plan σ^t that maximizes the welfare of the consumers subject to its budget constraints. Given an allocation rule f , the policy plan σ is rational for the government after the history h_{t-1} if the continuation of σ maximizes

$$V_t(\sigma^t, f^t, h_{t-1}) = \sum_{s=t}^{\infty} \beta^{s-t} U(c_{1s}(h_{s-1}), c_{2s}(h_s), l_s(h_s)) \quad (11)$$

subject to

$$g \leq \delta_s(h_{s-1})Rk_s(h_{s-1}) + \tau_s(h_{s-1})l_s(h_s)$$

for all histories induced by σ^t , for all $s \geq t$. Let $\Sigma^t(f^t, h_{t-1})$ denote the set of policy plans σ^t that satisfy these budget constraints.

Combining these requirements gives us a type of equilibrium that will not break down as time evolves, since by construction the various contingency plans will be carried out for any possible set of histories. We then define a sustainable equilibrium as a pair (σ, f) that satisfies

- Sequential rationality by consumers. Given a policy plan σ , the allocation rule f is rational for consumers at the first stage for every history h_{t-1} , and this allocation rule is rational for consumers at the second stage for every history h_t .
- Sequential rationality by governments. Given an allocation rule f , the plan σ is rational for the government for every history h_{t-1} .

For later use, let $V_0(\sigma, f)$ denote the value of utility in a sustainable equilibrium.

IV. Characterization of Sustainable Equilibria

In this section we characterize the allocations and policies that result from sustainable equilibria. Recall that a sustainable equilibrium (σ, f) is a sequence of functions that specify policies and allocations for all possible histories. Starting from the null history at date 0, a sustainable equilibrium induces a particular sequence of policies and allocations, say (π, x) . We call this the outcome induced by the sustainable equilibrium. The technique for characterizing the set of such outcomes builds on Abreu's (1988) seminal work on repeated games. In our models, however, agents behave competitively rather than strategically; thus, we need to reformulate Abreu's arguments.

We first show that the autarky allocations of the one-period economy (defined in Section II) are sustainable allocations in the infinite horizon economy. We then prove that a sequence of policies and allocations can be induced by some sustainable equilibrium if and only if it can be induced by a particular sustainable equilibrium called the revert-to-autarky equilibrium. We use this result to show that an arbitrary sequence of policies and allocations is an outcome of a sustainable equilibrium if and only if it satisfies two conditions: first, the sequence is a competitive equilibrium at date 0; second, the sequence satisfies some simple inequalities.

Consider, first, the autarky equilibrium (σ^a, f^a) , which is defined as follows. For any history h_{t-1} , the autarky policy plan sets $\sigma_t^a(h_{t-1})$ equal to the autarky policy π^a , as defined in the one-period problem in Proposition 2. The autarky allocation rule f^a is defined as follows: For the first stage, for every history h_{t-1} this rule specifies that consumers save nothing and consume all their endowment; that is, set $k^a(h_{t-1}) = 0$ and

$c_1^a(h_{t-1}) = \omega$. For the second stage, given any history h_t this rule sets the allocations equal to those defined in the one-period autarky problem, namely, $f_2^a(\pi_t)$, where the function f_2^a is defined in Proposition 2. We then have

Lemma 2. The autarky equilibrium is sustainable.

Proof. We first verify sequential rationality by consumers. For the first stage it is clear, given σ^a specifies a capital tax rate of 1, that it is optimal for consumers never to save. For the second stage, the consumer's problem reduces to the static problem of Proposition 2, so the second-stage allocations are optimal by construction. We next verify sequential rationality by the government. We need to show that for every history h_{t-1} ,

$$V_t(\sigma^a, f^a; h_{t-1}) \geq V_t(\hat{\sigma}^t, f^a; h_{t-1})$$

for every $\hat{\sigma}^t \in \Sigma^t(f^a; h_{t-1})$. Given that the consumer's contingency plan specifies zero saving for all future histories regardless of the past policies of the government, the tax problem of the government reduces to a series of static problems. For each such problem the solution is, by construction, the autarky policy of Proposition 2. Q.E.D.

In the next lemma we show that the autarky equilibrium is the worst sustainable equilibrium. Proving this is the key to our method of characterizing the set of sustainable allocations. We have

Lemma 3. The autarky equilibrium is the worst sustainable equilibrium. That is, for any sustainable equilibrium (σ, f) ,

$$V_0(\sigma, f) \geq V_0(\sigma^a, f^a). \tag{12}$$

Proof. For a given sustainable equilibrium (σ, f) , we construct a plan $\bar{\sigma}$ that satisfies

$$V_0(\sigma, f) \geq V_0(\bar{\sigma}, f) \geq V_0(\sigma^a, f^a). \quad (13)$$

Note that $V_0(\bar{\sigma}, f)$ can be well defined even when $(\bar{\sigma}, f)$ is not an equilibrium. In the construction, we exploit a fact about sustainable equilibria: second-stage decisions solve a simple static problem. To see this, notice that at the second stage of any period t , given the first-stage decisions and the current tax rates τ_t and δ_t , the contingency plans σ^{t+1} and f^{t+1} have no effect on the second-stage problem of consumers. More precisely, for any history $h_t = (h_{t-1}, \pi_t)$, the function $f_{2t}(h_{t-1}, \pi_t)$ can be written as some function $f_2(k_t, \pi_t)$ given in the solution to the following problem: choose c_2 and ℓ to solve

$$\max_{c_2, \ell} U(\omega - k_t + c_2, \ell) \quad (14)$$

subject to

$$c_2 \leq (1 - \delta_t)Rk_t + (1 - \tau_t)\ell.$$

Let us define $\bar{\sigma}$ as follows: For any h_{t-1} , let $\bar{\sigma}_t(h_{t-1})$ be the optimal tax policy in the problem

$$U^d(k_t) = \max_{\delta, \tau, c_2, \ell} U(\omega - k_t + c_2, \ell) \quad (15)$$

subject to

$$c_2 \leq (1 - \delta)Rk_t + (1 - \tau)\ell$$

$$-U_\ell / U_c = (1 - \tau)$$

$$g \leq \delta Rk_t + \tau\ell,$$

where k_t is given by $f_{1t}(h_{t-1})$. Now by construction of $\bar{\sigma}$ and the fact that f_{2t} solves (14), it is clear that $\bar{\sigma}$ is feasible for any such allocation rule. Thus, the first inequality in (13) follows from the definition of sequential rationality of the government.

Consider the second inequality in (13). The utility realized under the autarky plans σ^a and f^a is simply the discounted value of the one-period autarky utility repeated forever. We argue that the utility realized under the plans $\bar{\sigma}$ and f is at least as high. Let \bar{h}_t denote the history induced by $\bar{\sigma}$. For any t such that $f_{1t}(\bar{h}_{t-1})$ specifies zero saving, the time- t utility coincides with that of autarky. For any t such that $f_{1t}(\bar{h}_{t-1})$ specifies positive savings, the time- t utility exceeds that of autarky. In any such period the government will collect a strictly positive amount of revenue using what is essentially a lump-sum tax on capital. Thus, in period t , the amount of revenue the government raises using the distortionary tax on labor is smaller, so welfare is higher. Since this argument holds for any such period t , welfare under $(\bar{\sigma}, f)$ must be at least as high as it is under (σ^a, f^a) , where all revenue is raised through the distortionary labor tax. Q.E.D.

In the next proposition--which is this paper's main result--we characterize the conditions under which an arbitrary sequence of allocations and policies is sustainable. The proof uses a modified version of the autarky plans, which we call the revert-to-autarky plans. (For those familiar with Abreu 1988, it will be clear that these plans are related to Abreu's optimal simple penal codes.) For an arbitrary sequence (π', X') , the revert-to-autarky policy plans and allocation rules are denoted σ^r and f^r . The policy plan σ^r is defined as follows. If $h_{t-1} = (\pi'_0, \pi'_1, \dots, \pi'_{t-1})$, then $\sigma^r_t(h_{t-1}) = \pi'_t$. For all other histories, revert to the autarky tax rule of Lemma 2. Define the revert-to-autarky allocation rules f^r as follows. If $h_{t-1} =$

$(\pi'_0, \pi'_1, \dots, \pi'_{t-1})$, let $f_{1t}^r(h_{t-1}) = x_{1t}$. For all other histories let $f_{1t}^r(h_{t-1})$ be the autarky allocations x_1^a as specified in Lemma 2. The second-stage allocation functions are defined similarly. We then have

Proposition 3. Sustainable Outcomes. An arbitrary pair of sequences (π, x) is the outcome of a sustainable equilibrium if and only if (i) the pair (π, x) is a competitive equilibrium at date 0 and (ii) for every t , the following inequality holds:

$$\sum_{s=t}^{\infty} \beta^{s-t} U(c_{1s} + c_{2s}, k_s) \geq U^d(k_t) + \frac{\beta}{1-\beta} U^a, \quad (16)$$

where $U^d(k_t)$ is defined in (15).

Proof. Suppose, first, that (π, x) is the outcome of a sustainable equilibrium (σ, f) . Sequential rationality by consumers requires that (π, x) maximize consumer welfare at date 0. Sequential rationality by the governments implies that (π, x) satisfies the government's budget constraint at date 0. Thus, (π, x) is a competitive equilibrium under commitment. Next, at time t , given a history h_{t-1} , a deviation to the plan $\bar{\sigma}$ defined in Lemma 3 is feasible. Under this deviation, the time- t utility is $U^d(k_t)$, as defined in Lemma 3, and for any $s > t$, Lemma 3 guarantees that the time- s utility is at least U^a . Clearly, then, the utility of the government must be at least as large as the right side of (16) for every period t . Thus conditions (i) and (ii) hold.

Next, suppose some arbitrary pair of sequences (π, x) satisfies (i) and (ii). We show that the associated revert-to-autarky plans (σ^r, f^r) constitute a sustainable equilibrium. Consider histories under which there have been no deviations from π up until time t . Since (π, x) is a competitive equilibrium at date 0, it is obvious that its continuation from time t is

rational for consumers. Thus, sequential rationality for consumers holds for such histories. Consider the situation of the government. For any deviation at time t , the discounted value of utility from time $t + 1$ onward is given by the second term on the right side of (16). Recall that the policy plan $\bar{\sigma}$ was constructed to maximize time- t utility for any k_t . Thus, faced with such a history, the maximal utility attainable under any deviation by the policymaker at t is simply the right side of (16). Hence, given that the assumed inequality holds, then sticking with the specified plan is always optimal.

Now consider histories for which there has been a deviation before time t . The revert-to-autarky rules (σ^r, f^r) specify autarky from then onward. Clearly, the autarky allocations are rational for consumers. Finally, faced with the autarky allocation rule, it is optimal for the government to choose the autarky policy. Thus, (σ^r, f^r) is a sustainable equilibrium. Q.E.D.

Proposition 3 completely characterizes the conditions under which an arbitrary sequence of policies and allocations is sustainable. In particular, the proposition gives necessary and sufficient conditions for a date-0 competitive equilibrium to be the outcome of a sustainable equilibrium. It is worth noting that some competitive equilibria cannot be the outcome of any sustainable equilibrium. Indeed, any equilibrium that generates lower utility than autarky cannot be a sustainable outcome. For instance, consider an equilibrium with the tax on capital identically equal to one and with the tax on labor inefficiently high (for example, let the tax on labor be on the far side of the Laffer curve). Clearly, this equilibrium generates lower utility than the autarky equilibrium and, thus, is not sustainable. Notice that this equilibrium cannot be sustained for any discount factor in the unit interval.

It follows from Proposition 3 that if an outcome (π, x) is sustainable for some discount factor, then it is sustainable for a larger discount

factor. A more interesting result also follows: namely, if the discount factor is sufficiently high, the Ramsey outcome is sustainable. Formally, we have

Proposition 4. Sustainability of Ramsey Allocations. There is some discount factor $\underline{\beta} \in (0,1)$ such that for all $\beta \in (\underline{\beta},1)$ the Ramsey allocations are sustainable.

Proof. From Proposition 3 it suffices to show that the inequality (16) holds for the Ramsey allocations. Recall from Proposition 1 that in the one-period model, the level of utility under commitment is U^f . For the infinite horizon model, the utility under the Ramsey equilibrium is the discounted sum of these numbers U^f . Recall from Proposition 2 that in the one-period model, the level of utility under autarky is U^a . For the infinite horizon model, the utility under autarky is the discounted sum of utilities. Thus, to prove the result it suffices to verify the inequality

$$\frac{U^f}{1-\beta} \geq U^d(k^f) + \frac{\beta}{1-\beta} U^a.$$

Rearranging terms gives

$$\frac{\beta}{1-\beta} [U^f - U^a] \geq [U^d(k^f) - U^f]. \quad (17)$$

From Lemma 1, the left side of (17) is strictly positive. Thus, there is some $\underline{\beta}$ strictly less than one such that this inequality holds for all $\beta \geq \underline{\beta}$. Q.E.D.

Two remarks about Propositions 3 and 4 are warranted. First, it is important to emphasize that in the revert-to-autarky equilibrium, consumers do not "punish" the government when it deviates; rather, they choose the autarky allocations because, taking the future aggregate allocations and policies as

beyond their control, it is optimal to choose those allocations. Second, in these propositions, we develop conditions under which an infinite sequence of specified outcomes can be sustained by an equilibrium. A separate question is whether or not some specified discounted value of utility can arise in an equilibrium. For instance, from Lemma 2 it follows that for any discount factor β , the autarky utility $U^a/(1-\beta)$ is sustainable. From Proposition 3 it follows that for a high enough discount factor, the Ramsey utility $U^f/(1-\beta)$ is sustainable. One might wonder if any utility between the two is sustainable. Or, to rephrase the question more precisely: Given any number U satisfying $U^a < U < U^f$, is there some discount factor such that $U/(1-\beta)$ is the date-0 utility level of some sustainable equilibrium? Clearly, by considering an equilibrium that alternates in an appropriate fashion between the autarky and the Ramsey allocations and by choosing the discount factor to be high enough, any such utility level can be sustained.

V. An Example

We present an example to illustrate four features of sustainable outcomes and their associated utility levels that follow from Propositions 3 and 4. First, for low enough values of the discount factor, the only sustainable outcome is autarky. Second, if a certain outcome is sustainable for some discount factor, then it is sustainable for a larger discount factor. Third, for large enough values of the discount factor, the Ramsey outcome is sustainable. Fourth, for large enough values of the discount factor, all utilities between autarky and Ramsey are sustainable.

We focus on stationary outcomes--namely, outcomes (π, x) for which π_t and x_t are independent of t . For such outcomes the inequalities in (16) reduce to the single inequality

$$U(c_1+c_2, \ell) \geq (1-\beta)U^d(k) + \beta U^a. \quad (18)$$

In terms of characterizing the set of utilities that satisfy (18), it suffices to consider outcomes in which $\delta = (R-1)/R$, in which the tax on labor is set optimally, and in which k takes on all values in $[0, \omega]$. For any k , let $U(k)$ denote the maximized value of utility under such an outcome. For any discount factor β in $[0, 1]$, let $E(\beta)$ be the set of stationary sustainable utility levels; that is,

$$E(\beta) = \{U(k) \mid U(k) \geq (1-\beta)U^d(k) + \beta U^a, k \in [0, \omega]\}. \quad (19)$$

Let the utility function be

$$U(c_1+c_2, \ell) = [(c_1+c_2)^\alpha + \gamma(\bar{x}-\ell)^\alpha]^{1/\alpha} \quad (20)$$

and let $\alpha = -0.3$, $\gamma = 1.2$, $\bar{x} = 100$, $\omega = 10$, $g = 25$, and $R = 2$. For this example, the set of stationary sustainable utility levels illustrates the four features (see Figure 1). First, for $\beta < 0.1$, $E(\beta) = U^a$. Second, $E(\beta) \subset E(\beta')$ for $\beta < \beta'$. Third, for $\beta \geq 0.1$, the $U^f \in E(\beta)$. Fourth, the $\beta \geq 0.48$, $E(\beta) = [U^a, U^f]$.

Finally, a rather special feature of the example is that for some values of β --namely, $\beta \in [0.1, 0.48]$, the Ramsey utility is sustainable, but some utilities between the Ramsey utility and the autarky utility are not sustainable (at least with stationary outcomes).

VI. Anonymous Games

In this section we provide one rationalization of the equilibria considered in the previous sections, but in a game-theoretic context. In particular, we make precise the relationship between the equilibria of the economies with and without commitment and the perfect equilibrium of certain games. We first

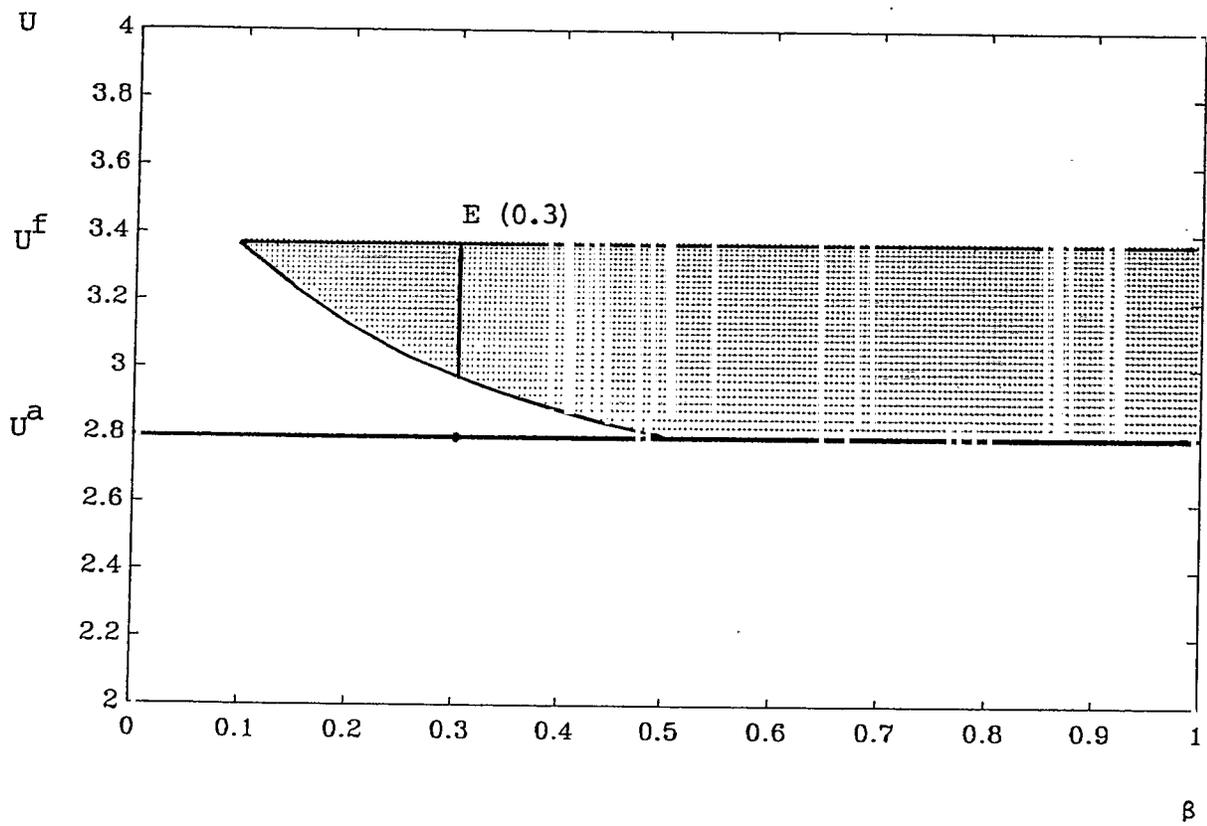


Fig. 1 -- The set of stationary sustainable utility levels

show that the Ramsey equilibrium is the unique subgame perfect equilibrium of a game with commitment. More important, we then show that the set of sustainable equilibria correspond to the set of symmetric perfect Bayesian equilibria of a game with no commitment. (For related work, see Atkeson 1988.)

In the economies considered earlier, we modeled private agents as behaving competitively, in the sense that they assume policies are unaffected by their individual decisions. We capture this feature in a game by using two assumptions. First, we assume there is a continuum of agents. Second, we assume individuals observe only their own decisions and aggregate outcomes. A game with this feature is called an anonymous game (see Green 1980, 1984).

A. General Setup

There is a continuum of private agents, represented by Lebesgue measure λ on the interval $[0,1]$, and a player called the government. A policy for the government is a pair of tax rates $\pi = (\delta, \tau)$ with $0 \leq \delta, \tau \leq 1$. An action profile for private agents is a pair of measurable functions $x = (k, \ell): [0,1] \rightarrow [0, \omega] \times \mathbb{R}_+$. We denote the implied action of an individual agent i by $x(i) = (k(i), \ell(i))$. The single-period payoffs of agent i are

$$V_i(\pi, x(i), x) = U(\omega - k(i) + (1 - \delta)Rk(i) + (1 - \tau)\ell(i), \ell(i)) + W(\delta, \tau, K, L), \quad (21)$$

where $K = \int k(i)\lambda(di)$ and $L = \int \ell(i)\lambda(di)$ and where the function W equals zero if its arguments satisfy $g \leq \delta RK + \tau L$ but equals some large negative number, say $-M$, if otherwise. The government's payoff is

$$V(\pi, x) = \int V_i(\pi, x(i), x)\lambda(di). \quad (22)$$

Recall that in the usual definition of a game, there are no budget constraints. The function W incorporates the budget constraint of the government into its preferences in such a way that the government will seek to balance the budget.

In what follows, we consider infinite horizon games in which agents maximize the discounted present value of the stream of single-period payoffs.

B. Commitment Game

In a commitment game, the government first chooses an infinite sequence of policies $\pi = (\pi_t)_0^\infty$. A strategy for the government is thus just an infinite sequence of policies. Private agents, having seen π , then make their decisions. A strategy profile for private agents is a sequence of functions $f = (f_t)_0^\infty$ that maps policies π into action profiles x . A strategy profile f naturally induces strategies for each agent, and these strategies take the form $f_t(i, \pi)$ for every period. For private agents, the payoffs over strategies are defined as

$$\sum_{t=0}^{\infty} \beta^t V_i(\pi_t, f_t(i, \pi), f_t(\pi)). \quad (23)$$

Likewise, the payoffs for the government are

$$\sum_{t=0}^{\infty} \beta^t V(\pi_t, f_t(\pi)). \quad (24)$$

We can now define an equilibrium: A subgame perfect equilibrium for the commitment game is a strategy π for the government and a strategy profile f for private agents that together satisfy two conditions: (i) for each agent i , given the strategies of other agents as specified by f and any policy π' for the government, the strategy $f(i, \pi')$ maximizes the agent's payoff, and (ii) given the strategy profile f , the strategy π maximizes the government's payoff. Comparing this definition with the Ramsey equilibrium of Section III, we immediately have

Proposition 5. Equilibrium Outcomes of the Commitment Game. The subgame perfect equilibrium policies and allocations $(\pi, f(\pi))$ of the commitment game are identical to the Ramsey policies and allocations.

The proof of Proposition 5 is given in the Appendix. The requirement of subgame perfection is crucial in demonstrating this proposition. Suppose we drop the requirement of perfection and consider Nash equilibria instead. Recall that a Nash equilibrium is defined as above, except we require the strategy profile, say f^* , to be an equilibrium for private agents only at the equilibrium policy of the government, say π^* . Thus, for policies other than π^* the strategy profile f^* is unrestricted. It is easy to see that the set of Nash equilibria is considerably larger than the set of subgame perfect equilibria. Indeed, any competitive equilibrium (π, x) is the outcome of a Nash equilibrium. To see this, let the strategy profile f^* specify x if the policy π is chosen and specify zero savings and zero labor supply if any other policy is chosen. By construction of W , the government's payoff is some large negative number for any policy other than π . Hence, it is optimal for the government to choose π . Then, since (π, x) is a competitive equilibrium, x is a best response to π . Thus, (π, f^*) is a Nash equilibrium with outcome (π, x) .

C. The No-Commitment Game

Next, consider a game without a commitment technology. Let the timing of the moves be the same as in the no-commitment infinite horizon economy. In defining this game, we must be careful about what the players have observed when they make their decisions. We formalize this by defining histories both of the game and for the players. The history of the game is a complete description of all the actions chosen in the past by all players. In particular, at the first stage of period t , the history of the game is $h_{1t} = (x_s, \pi_s | s < t)$. At the second stage of t , the history of the game is $h_{2t} = h_{1t} \cup (x_{1t}, \pi_t)$. In contrast, the history for a player i consists only of observed outcomes. Each individual observes only aggregate outcomes and, of course, that individual's

own past decisions. Thus, a player i 's history at the first stage of period t is $h_{1t}(i) = (x_s(i), X_s, \pi_s | s < t)$, where $X_s = \int x_s(i) \lambda(di)$. The history for player i at the second stage is similarly defined. The other player, the government, observes only aggregate outcomes. A history for the government at time t is $H_t = (X_s, \pi_s | s < t) \cup X_{1t}$.

Players' histories correspond to information sets in the obvious way. For example, player i 's history $h_{1t}(i)$ at time t corresponds to the information set consisting of all histories of the game h_{1t} consistent with $h_{1t}(i)$. More precisely, an individual history $h_{1t}(i) = (x_s(i), X_s, \pi_s | s < t)$ corresponds to an information set consisting of all histories $h'_{1t} = (x'_s, \pi'_s | s < t)$ that satisfy $x_s(i) = x'_s(i)$, $X_s = \int x'_s(i) \lambda(di)$, and $\pi_s = \pi'_s$ for all $s < t$. From now on we identify information sets with histories in this way.

Consider, next, the strategies for the players in the game. A strategy for the government is a sequence of functions $\sigma = (\sigma_t)_{t=0}^{\infty}$ which, for each t , maps government histories H_t into policies π_t . A strategy profile for private agents is a sequence of functions $f = (f_{1t}, f_{2t})_{t=0}^{\infty}$ which, for each stage, maps histories of the game into action profiles. A strategy profile naturally induces strategies of the form $f_{1t}(i, h_{1t})$ and $f_{2t}(i, h_{2t})$ for each agent. To be consistent with our informational restrictions, we require that for each player i , the strategies $f_{1t}(i, \cdot)$ and $f_{2t}(i, \cdot)$ depend only on individual histories. [Technically, we require that $f_t(i, \cdot)$ be measurable with respect to the σ -algebra generated by the individual histories.] Such profiles will be called anonymous strategy profiles.

Payoffs for the players are naturally defined from the outcomes that the strategies induce. For example, the payoff for player i at time t , given a history of the game h_{1t} , is

$$W_{it}(\sigma, f(i), f; h_{1t}) = \sum_{s=t}^{\infty} \beta^{s-t} v_i(\pi_s, x_s(i), x_s), \quad (25)$$

where the future actions are induced from h_{1t} by f and σ . The payoff for the government at time t , given a history of the game h_t , is similarly defined:

$$W_t(\sigma, f; h_t) = \sum_{s=t}^{\infty} \beta^{s-t} V(\pi_s, f_s), \quad (26)$$

where the future actions are induced from h_t by f and σ .

Now we want to define some type of perfect equilibrium for this game. One approach would be to consider subgame perfect equilibrium. Given the informational restrictions, however, the only proper subgame is the original game itself; hence, any Nash equilibrium is subgame perfect.⁵ An alternative is to consider a type of Bayesian equilibrium. (See, for example, Fudenberg and Tirole 1988.)

A Bayesian equilibrium consists of strategy profiles together with a sequence of probability distributions. For every information set there is a probability distribution over histories of the game consistent with that information set. Let $\mu(h_{1t} | h_{1t}(i))$ denote a probability distribution over the histories of the game h_{1t} that are consistent with the information set associated with player i 's first-stage history $h_{1t}(i)$. Likewise, let $\mu(h_t | H_t)$ and $\mu(h_{2t} | h_{2t}(i))$ denote probability distributions over a government information set and over a player i 's second-stage information set. Let μ denote the collection of these probability distributions. Given some collection of probability distributions μ and strategies σ and f , we can use (25) to write the expected utility of player i at the information set associated with history $h_{1t}(i)$ as

$$\int W_{it}(\sigma, f(i), f; h_{1t}) d\mu(h_{1t} | h_{1t}(i)).$$

We use (26) to define the expected utility for the government at the information set associated with a history H_t as

$$\int W_t(\sigma, f; h_t) d\mu(h_t | H_t).$$

The payoffs for players at the second-stage are similarly defined.

A perfect Bayesian equilibrium is an anonymous strategy profile f , a government strategy σ , and a collection of probability distributions μ such that (i) for each player i , period t , and history $h_{jt}(i)$, for $j = 1, 2$, the continuation of the strategy $f(i)$ maximizes player i 's expected payoff; (ii) for each period t and history H_t , the continuation of σ maximizes the government's expected payoff; and (iii) μ assigns probability one to histories of the game along the equilibrium path.

To understand condition (iii), consider an example of a probability distribution over an information set of the government at date 1. The history of the game along the equilibrium path is $h_1 = x_1 = f_1(h_{10})$, where h_{10} is a null history. This history of the game is a member of the government's information set corresponding to the history $H_1 = X_1 = \int f_1(i, h_{10}) \lambda(di)$. Condition (iii) requires that μ assign probability one to this history, namely, $\mu(h_1 | H_1) = 1$, and probability zero to any other history, $h'_1 \neq f_1(h_{10})$, in this information set. Notice that condition (iii) places no restrictions on μ for histories of the game off the equilibrium path.

In the equilibria of Sections III and IV, we used a representative agent to model the private agents. The standard interpretation of this construct is that the representative agent stands in for a large number of competitive private agents who act identically in equilibrium. To keep the analysis of the game model parallel with a representative agent model, the equilibria must be symmetric. In the commitment game it is easy to see that all the equilibria are (almost everywhere) symmetric, so we did not need to impose symmetry. But in the no-commitment game there typically are asymmetric equilibria; hence, for that game we require symmetry.

We say (σ, f) is a symmetric perfect Bayesian equilibrium if, in addition to satisfying conditions (i)-(iii), it satisfies two other conditions: (iv) the strategies of consumers are symmetric, and (v) μ assigns probability one to symmetric histories (both on and off the equilibrium path). To understand condition (v) consider, for example, the information set of player i corresponding to the history $h_{1t}(i) = (x_s(i), X_s, \pi_s | s < t)$. Condition (v) requires that $\mu(\cdot | h_{1t}(i))$ assign probability one to the symmetric history of the game associated with $h_{1t}(i)$. That is, probability one is assigned to the history of the game $h'_{1t} = (x'_s, \pi'_s | s < t)$ that for each $s < t$ satisfies $\pi'_s = \pi_s$, $x'_s(j) = X_s$ for each $j \neq i$, and $x'_s(i) = x_s(i)$. In other words, condition (v) requires that at any information set, player i believes with probability one that all the other private agents have behaved symmetrically in the past. Similarly, condition (v) requires the government to believe with probability one that all private agents have behaved symmetrically in the past.

Proposition 6. Equilibrium Outcomes of the No-Commitment Game. The set of symmetric perfect Bayesian equilibrium policies and allocations of the no-commitment game is the same as the set of sustainable equilibrium policies and allocations.

Here we provide an intuitive explanation of the proposition and relegate the formal proof to the Appendix. The essential difference between the definitions of a sustainable equilibrium and a symmetric perfect Bayesian equilibrium is that the latter requires rationality after histories with private deviations, whereas the sustainable equilibrium does not even consider such histories. The main point of Proposition 6 is that in the symmetric perfect Bayesian equilibrium, the extra conditions imposed after such his-

tories are irrelevant to the set of outcome paths. Part of the proof relies on a straightforward result from game theory: Consider two equilibria whose strategy profiles coincide for all histories without simultaneous deviations in the past. The result is that these two equilibria generate identical outcome paths. The reason is simply that in checking whether a deviation from a candidate equilibrium is profitable, we need, by the definition of equilibrium, to check only deviations by a single player. Intuitively, after any history without simultaneous deviations, no player acting alone can induce future histories with simultaneous deviations. Thus, regardless of how we specify the continuation equilibrium after simultaneous deviations (as long as it is some equilibrium), we get the same outcome path.

In our anonymous game, the only deviations by private agents that can influence the future behavior of other private agents or the government (by affecting their information sets) are ones in which a positive measure of agents deviate simultaneously and change the aggregate outcomes. By this result we can ignore such deviations, in the sense that no matter how we fill in the continuation equilibrium after such histories, we get the same outcome path. Moreover, given that no single private agent's deviation can affect the payoff or the information sets reached by other agents, we can ignore single deviations by private agents. Combining these results and using the definitions of sustainable outcomes and symmetric perfect Bayesian outcomes, the result follows.

Four factors--perfection, symmetry, anonymity, and continuum of players--all play crucial roles in the proof of Proposition 6. First, for essentially the same reasons as in the commitment game, the set of Nash equilibria is much larger than the set of perfect equilibria. Second, the set of perfect Bayesian equilibria is larger than the set of symmetric perfect Baye-

sian equilibria. For example, since consumers are indifferent among all saving levels when $(1-\delta)R = 1$, we can have asymmetric equilibria where some consumers save all of their endowment and others save none. Furthermore, the assumption that the probability distributions assign probability one to symmetric histories is important. Without it, the government's strategies off the equilibrium path are affected, and consequently the set of equilibria can be larger.

The third factor, anonymity, is somewhat more subtle. Suppose, for example, that private agents can observe each other's actions. We can show that for sufficiently little discounting, it is possible to support the equilibrium allocations obtained with lump-sum taxation. We support these equilibrium allocations, the so-called command optimum allocations, using certain trigger strategies. These strategies specify that in every period, agents save their entire endowment and supply the amount of labor specified by the command optimum--as long as all agents have chosen these actions in the past. If any player deviates, the strategies specify that each agent chooses the worst sustainable equilibrium allocations. With sufficiently little discounting, the gains from deviating are outweighed by the future losses, so no agent will deviate. Notice that while no single private agent has any effect on current aggregate outcomes, the fact that each agent's actions are observable means that a deviation by a single agent can trigger a move to a "bad" equilibrium. Our restriction that actions are unobservable and our assumption that there is a continuum of agents together imply that a single agent can deviate without being detected by any other player in the game. In our game, these types of trigger strategies are inconsistent with the information structure.

Fourth, notice that the type of game set up here is quite different from the standard repeated oligopoly game of Friedman (1971), as well as the more general class of repeated games analyzed by Fudenberg and Maskin (1986) and Abreu (1988). Their games have a finite number of players with standard information structures. In contrast, our game has one large player and a continuum of small anonymous players. Such a structure does not fall into the class of games analyzed by those authors, and their results do not directly apply. Of course, they also do not directly apply to sustainable equilibrium either. The essential difference is that in the game, our private agents are anonymous (or, in the sustainable equilibrium, they are competitive).⁶

VII. Conclusion

We wrote this paper to address four related questions:

- Is it possible to build a simple general equilibrium model in which private agents are competitive, in which the government maximizes the welfare of these agents, and which exhibits trigger-type equilibria?
- If so, precisely what is the equilibrium concept; in particular, what are the decision problems of private agents?
- Is it possible to characterize all the equilibria?
- How is the notion of time consistency related to standard notions of perfection in game theory?

We analyzed these questions in a variant of Fischer's taxation model. We developed an equilibrium concept in which private agents are competitive and in which trigger-type equilibria are possible. We characterized the equilibrium outcomes by a pair of simple conditions. And we showed the equivalence between sustainable outcomes and the symmetric perfect Bayesian equilibrium outcomes of an appropriately defined anonymous game.

We believe that our results will readily generalize to a wide variety of models. For some preliminary work in that direction, see papers by Chari and Kehoe (1989) and Chari, Kehoe, and Prescott (1989).

Appendix

Proof of Proposition 5

Actually, to be technically precise the proposition should say that the realized allocations in the subgame perfect equilibrium, say \hat{x} , coincide with the Ramsey outcomes, say x^* , almost everywhere [in the sense that for every t , the measurable functions $\hat{x}_t(i)$ are constant and equal to the scalars x_t^* up to some set A_t that has a Lebesgue measure of zero].

Proof. We show that the equilibrium action profiles of consumers in the game coincide (up to sets of measure zero) with the Ramsey equilibrium allocation function. It is then immediate that the government's choice is the same in both environments. Note that for any policy π such that $\delta_t \neq \delta^f$, all consumers make the same decisions in the game as in the economic environment. Suppose, for some t , $\delta_t = \delta^f$ and a measurable set of consumers choose to save less than ω . Using the same argument as in Proposition 1, the government can increase its utility by choosing a slightly lower capital tax rate and by raising the rest of the needed revenues from labor taxation. Therefore, the equilibrium action profiles coincide with the Ramsey equilibrium allocation function. Q.E.D.

Proof of Proposition 6

To establish Proposition 6, we use a simple lemma which establishes that the set of strategies and beliefs (the natural analogues of the autarky policy plans and allocation rules) is a symmetric perfect Bayesian equilibrium. In the proof of Proposition 6, we only use this lemma in cases where we construct a continuation equilibrium after histories with simultaneous deviations. Throughout this section we identify histories with information sets in the obvious manner.

The autarky strategies and beliefs (σ^a, f^a, μ^a) are defined as follows. For any history $H_t = (H_{t-1}, X_{2t-1}, X_{1t})$ with $X_{1t} = K_t$, let $\sigma^a(H_t)$ specify the tax policies that solve (15), with k_t replaced by K_t . For any history $h_{1t}(i)$, let $f_{1t}^a(i, h_{1t}(i))$ specify zero savings. For any history $h_{2t}(i)$ with $x_{1t}(i) = k_t(i)$, let $f_{2t}^a(i, h_{2t}(i))$ specify the value of labor that solves (14) with k_t replaced by $k_t(i)$. For any information set, let μ^a assign probability one to the appropriate symmetric history of the game.

Lemma. The autarky strategies and beliefs (σ^a, f^a, μ^a) constitute a symmetric perfect Bayesian equilibrium.

Proof. We first verify condition (i): optimality by consumers. Consider the consumer's first-stage problem after some history $h_{1t}(i)$. The consumer expects the government to set the capital taxes according to $\delta^a(H_t)$. Since $\delta^a(H_t)$ solves (15), we know it equals $\min(g/RK_t, 1)$. Since, by assumption, $g > (R-1)\omega$, it then follows that $(1-\delta^a(H_t))R \leq 1$. Thus, it is optimal for the consumer not to save. Next consider the consumer's second-stage problem after some history $h_{2t}(i)$. It is clear that this consumer's second-stage action has no effect on the individual's future payoffs. Thus, the consumer's problem reduces to the static problem of maximizing current utility given $k_{1t}(i)$, τ_t , and δ_t . By construction, the solution is $f_{2t}^a(i, h_{2t}(i))$.

Now consider condition (ii): optimality by government. Faced with a strategy profile f^a , nothing the government does at t influences its payoffs from $t + 1$ onward. Thus, the government's problem reduces to a static problem of maximizing current utility given the current history. By construction, μ^a assigns probability one to the symmetric history of the game in this information set. Thus, the problem of the government reduces to (15) which, by construction, is solved by $\sigma^a(H_t)$.

Next, by construction, conditions (iii)-(v) are satisfied. Hence, (σ^a, f^a, μ^a) is a symmetric perfect Bayesian equilibrium. Q.E.D.

Using this lemma, we can now prove Proposition 6.

Proof. Suppose (σ, f) is a sustainable equilibrium with outcome (π, x) . We construct a symmetric perfect Bayesian equilibrium $(\hat{\sigma}, \hat{f}, \hat{\mu})$ that has the same outcome. We consider two kinds of histories in the game. First, we consider histories of the government, $H_t = (\pi'_s, X_s \mid s < t)$, that satisfy $X_{1s} = f_{1s}(h'_{s-1})$, $X_{2s} = f_{2s}(h'_s)$ for all $s < t$ and for all $h'_s = (\pi'_0, \dots, \pi'_{s-1})$. Such histories are histories with no simultaneous deviations by consumers relative to the sustainable equilibrium (σ, f) . For such histories, let $\hat{\sigma}_t(H_t) = \sigma_t(h'_{t-1})$. For all other histories, let $\hat{\sigma}_t(H_t) = \sigma_t^a(H_t)$, as defined in the above lemma. The consumer's strategies \hat{f} are constructed from the allocation rule f in the sustainable equilibrium, and the autarky strategies \hat{f} are constructed analogously. [A minor detail is that we set $\hat{f}_{2t}(i, h_{2t}(i))$ equal to the decision that solves (14), with k_t replaced by $x_{1t}(i)$ for all histories $h_{2t}(i)$.] The construction of $\hat{\mu}$ is obvious.

For histories with no simultaneous deviations, no player can profitably deviate from $(\hat{\sigma}, \hat{f})$, since (σ, f) is a sustainable equilibrium. For all other histories, no player can profitably deviate from $(\hat{\sigma}, \hat{f})$ since from the above lemma, (σ^a, f^a, μ^a) is a symmetric perfect Bayesian equilibrium. By construction, $(\hat{\sigma}, \hat{f}, \hat{\mu})$ satisfies conditions (iii)-(v) of the definition of such an equilibrium. Therefore $(\hat{\sigma}, \hat{f}, \hat{\mu})$ is a symmetric perfect Bayesian equilibrium.

The converse is immediate, since a symmetric perfect Bayesian equilibrium requires optimality for all histories, whereas a sustainable equilibrium requires decisions to be optimal only for histories with no deviations by consumers. Q.E.D.

Notes

¹This algorithm does not adequately impose sequential rationality when there are multiple competitive equilibria associated with a given policy. As will become clear, our proposed definition of equilibrium can address this problem.

²The difficulty is in modeling the behavior of private agents in the game so that the equilibrium outcomes are the same as the competitive equilibrium outcomes. In addition, the particular set of equilibrium outcomes depends on seemingly small changes in the structure of the game. (These issues are discussed in Shubik 1982.)

³The value of this utility function coincides with each consumer's utility level when all consumers choose the same actions. In what follows, we consider symmetric equilibria. Alternatively, the government can be thought of as maximizing the sum of consumers' utilities. For the latter approach, see Section VI.

⁴Note that we have defined the government's choice as a pair of numbers but the consumers' second-stage allocations as a function. Since the government's optimal policy varies with the first-stage decisions of consumers, why don't we define the government's choice as a function of first-stage decisions? Indeed, in an earlier version of this paper we defined sustainable equilibrium precisely that way. It turns out, however, that we don't need to define the government's optimal decision for all first-stage outcomes, because no single consumer perceives that the government will change its policies if that consumer individually changes a decision. However, the government perceives that its policy choices will alter second-stage decisions. Therefore, consumers' second-stage decisions must be described as functions, whereas the government's decision is described as a pair of numbers. (Also, see Section VI, where we show that deviations by consumers can be ignored in a game.)

⁵It should be clear that a large number of rather bizarre (subgame perfect) Nash equilibrium outcomes exist for this game. Thus, it is not quite precise to say that dynamic consistency is equivalent to subgame perfection.

⁶In particular, Fudenberg and Maskin (1986) show that with sufficiently little discounting, any vector of average payoffs that is better than mutual minimax can be supported by a perfect equilibrium. In our model, this is not true. By normalizing $U(w,0)$ to be zero, it is clear the the mutual minimax payoffs are $-M$. (Each player saves nothing and doesn't work, and the government cannot meet its budget constraint.) In our model, regardless of the discount factor, no average utility that is lower than autarky (some positive number) can be supported. The key difference, of course, is that here private agents are anonymous (or competitive). (Technically, our model doesn't satisfy Fudenberg and Maskin's "full-dimensionality" condition.)

References

- Abreu, Dilip. "On the Theory of Infinitely Repeated Games with Discounting." Econometrica 56 (March 1988): 383-96.
- _____. "Extremal Equilibria of Oligopolistic Supergames." J. Econ. Theory 39 (June 1986): 191-225.
- Atkeson, Andrew G. "Time Consistency of Policy in a Macro Model." Mimeographed. Stanford, Calif.: Grad. School Bus., Stanford Univ., 1988.
- Barro, Robert J., and Gordon, David B. "Rules, Discretion, and Reputation in a Model of Monetary Policy." J. Monetary Econ. 12 (July 1983): 101-21.
- Calvo, Guillermo A. "On the Time Consistency of Optimal Policy in a Monetary Economy." Econometrica 46 (November 1978): 1411-28.
- Chari, V. V., and Kehoe, Patrick J. "Sustainable Plans and Debt." Research Department Working Paper no. 354. Federal Reserve Bank of Minneapolis, 1989.
- Chari, V. V.; Kehoe, Patrick J.; and Prescott, Edward C. "Time Consistency and Policy." In Modern Business Cycle Theory, ed. Robert J. Barro, pp. 265-305. Cambridge, Mass: Harvard Univ. Press, 1989.
- Fischer, Stanley. "Dynamic Inconsistency, Cooperation, and the Benevolent Dissembling Government." J. Econ. Dynamics and Control 2 (February 1980): 93-107.
- Friedman, James W. "A Non-cooperative Equilibrium for Supergames." Rev. Econ. Studies 38 (January 1971): 1-12.
- Fudenberg, Drew, and Maskin, Eric. "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information." Econometrica 54 (May 1986): 533-54.
- Fudenberg, Drew, and Tirole, Jean. "Perfect Bayesian and Sequential Equilibria: A Clarifying Note." Mimeographed. Cambridge, Mass.: MIT, 1988.

Green, Edward J. "Noncooperative Price Taking in Large Dynamic Markets." J. Econ. Theory 22 (April 1980): 155-82.

_____. "Continuum and Finite-Player Noncooperative Models of Competition." Econometrica 52 (July 1984): 975-93.

Kydland, Finn E., and Prescott, Edward C. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans." J. P. E. 85 (June 1977): 473-91.

_____. "Dynamic Optimal Taxation, Rational Expectations and Optimal Control." J. Econ. Dynamics and Control 2 (February 1980): 79-91.

Lucas, Robert E., Jr., and Stokey, Nancy L. "Optimal Fiscal and Monetary Policy in an Economy without Capital." J. Monetary Econ. 12 (July 1983): 55-93.

Ramsey, Frank P. "A Contribution to the Theory of Taxation." Econ. Journal 37 (March 1927): 47-61.

Rogoff, Kenneth. "Reputational Constraints on Monetary Policy." In Bubbles and Other Essays, ed. Karl Brunner and Allan H. Meltzer. Carnegie-Rochester Conference Series on Public Policy 26 (1987): 141-82. Amsterdam: North-Holland.

Shubik, Martin. Game Theory in the Social Sciences: Concepts and Solutions. Cambridge, Mass.: MIT Press, 1982.