CHICAGO JOURNALS

Sustainable Plans

V. V. Chari and Patrick J. Kehoe

Federal Reserve Bank of Minneapolis and University of Minnesota

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 optimality conditions for the government and for private agents. We provide a complete characterization of the sustainable equilibrium outcomes for a variant of Fischer's model of capital taxation. We also relate our work to recent developments in the theory of repeated games.

I. Introduction

This paper describes a framework for analyzing the optimal design of government policy in dynamic general equilibrium models with competitive private agents and with governments lacking commitment technologies. In environments in which societies have a commitment technology to bind the actions of future governments, the policy design problem is well understood. The government chooses a policy, namely a sequence of event-contingent functions, once and for all, and then consumers make their decisions sequentially in a competitive fashion. In an environment with commitment, the optimal policy, together with the resulting competitive equilibrium, is called a "Ramsey equilibrium."

In environments in which societies have no such ability to bind future policy choices, the policy design problem is less well understood. Kydland and Prescott (1977) argued that in such environments the sensible way to set up the policy design problem was to formulate the decision problems of both the government and the private agents sequentially, requiring that choices be optimal at each point in time. For a finite-horizon model, they showed how to compute the optimal policy using backward induction. For infinite-horizon models, taking the limit of the backward induction solution recovers only one of what may be a much larger set of policies and allocations that are consistent with sequential optimality of the government and private agents. In particular, taking the limit of the backward induction solution rules out the possibility of trigger-type equilibria.

This paper studies several issues related to the analysis of policy design in infinite-horizon models. We first show how to formulate 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. We develop an equilibrium concept in which both the government and private agents' decision problems are sequential. We call this type of equilibrium a "sustainable equilibrium" to distinguish it from others in the literature (see, e.g., Lucas and Stokey 1983; Atkeson 1988). We illustrate this equilibrium in an infinite-horizon version of Fischer's (1980) optimal taxation model. 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 0 competitive equilibrium. Second, it must satisfy a simple set of inequalities. We use these inequalities to show that with sufficiently little discounting, even the Ramsey allocations are sustainable. We then relate our equilibrium concept to that of perfect Bayesian equilibrium in an appropriately defined game.

The novel feature of our approach is that it blends features of classical competitive analysis and game theory. In our model the government is strategic in that it takes into account that its choice of policy will affect the choices of private agents. Thus in making its decisions the government must forecast how its policy choice will affect the future behavior of private agents. To pose this forecasting problem, we define policies and allocations as functions of the history of decisions of the government. This is a break from the classical competitive approach, which defines equilibria as functions of exogenous events. However, the model, as we set it up, is not a standard game either because the histories do not include the past actions of all agents; in particular, they do not include those of consumers. It is also not a standard game because utilities are not defined for all possible choices of the private agents or the government. For example, we do

not define utilities for situations in which either the consumers or the government violates the budget constraints or in which identical private agents take different actions. It is possible to set this model up as a game. Indeed, we show how to map the environment into a game and demonstrate that the symmetric perfect Bayesian outcomes of that game coincide with the sustainable outcomes of our model. Thus our notion of equilibrium captures the notion of perfection games. An inspection of the cumbersome notation of the game, however, makes it clear that our approach is much simpler.

Before turning to the model, we briefly summarize how our work relates to the three distinct literatures on which it builds: 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; Lucas and Stokey 1983). In this literature, the decisions of private agents and the government depend on a small number of state variables; in our formulation, these decisions can depend on a much larger set of state variables, namely, the complete history of past policy. This difference explains why in our formulation the set of equilibria without commitment is much larger than others, possibly including the Ramsey allocations and policies.

Second, we build on ideas developed in the theory of repeated games, particularly in the oligopoly literature (see Friedman 1971; Fudenberg and Maskin 1986; Abreu 1988). 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, such as the folk theorem, cannot be applied.

Third, we build on the recent literature on time consistency in macroeconomic models (see Barro and Gordon 1983; Backus and Driffill 1985; Rogoff 1987). In this literature it often appears as if the government plays a game against a coalition of noncompetitive private agents, who may have objectives different from those of the government. In our model the government maximizes the welfare of private agents, who behave competitively.

II. An Infinite-Horizon Economy

Consider a simple infinite-horizon version of a model similar to the one in Fischer (1980). The economy contains a large number of iden-

tical consumers and a government. There is a linear production tech-
nology, for which the marginal product of capital is a constant $R > 1$
and for which the marginal product of labor is one. In each period
$t, t = 0, 1, \ldots$, consumers make decisions at two distinct points. At the
first stage of t, consumers are endowed with ω units of the consump-
tion good out of which they consume c_{1t} and save k_t. At the second
stage, they consume c_{2t} and work l_t units. Second-stage income, net of
taxes, is $(1 - \delta_t)Rk_t + (1 - \tau_t)l_t$, where δ_t and τ_t denote the time t tax
rates on capital and labor. For simplicity we assume that first-stage
consumption and second-stage consumption are perfect substitutes.
We also assume that capital cannot be stored between periods. A
consumer's preferences are given by the discounted value of the util-
ity per period, $U(c_{1t}, c_{2t}, l_t)$, where the discount factor β satisfies $0 <
\beta < 1$.

In each period t the government sets proportional tax rates on
capital and labor income to finance an exogenously given amount of
second-stage per capita government spending. We assume, through-
out, that $g > (R - 1)\omega$. This assumption implies that in any equilib-
rium, the government must tax labor. We also assume that it is feasi-
ble to finance government spending with only a tax on labor.

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. Con-
sumers then choose a sequence of allocations for all time. In contrast,
in the no-commitment version, the government and the consumers
make decisions sequentially. 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 $\pi = (\pi_0, \pi_1, \ldots)$ denote an infinite sequence of tax rates
starting at time 0. For each period t, let $x_t = (x_{1t}, x_{2t})$ be the allocations
for the first and second stages of t, with $x_{1t} = (c_{1t}, k_{1t})$ and $x_{2t} = (c_{2t}, l_{2t})$. Let $x = (x_0, x_1, \ldots)$ denote the infinite sequence of such alloca-
tions. For this environment, a policy for the government is an infinite
sequence of tax rates π. An allocation rule is a sequence of functions $f
= (f_0, f_1, \ldots)$ that maps government policies into sequences of al-
locations.

A *Ramsey equilibrium* is a policy π and an allocation rule f that satisfy
the following conditions: (i) For every policy π', the allocation rule
$f(\pi')$ maximizes $\Sigma_{t=0}^{\infty} \beta^t U(c_{1t} + c_{2t}, l_t)$ subject to $c_{1t} \le \omega - k_t$ and $c_{2t} \le
(1 - \delta_t)Rk_t + (1 - \tau_t)l_t$; (ii) the policy π maximizes $\Sigma_{t=0}^{\infty} \beta^t U(c_{1t}(\pi)$

$+ c_{2t}(\pi)$, $l_t(\pi)$) subject to $g \leq \delta_t R k_t(\pi) + \tau_t l_t(\pi)$ for $t = 0, \ldots, \infty$.

We denote a Ramsey equilibrium by the pair (π, f). In the Ramsey equilibrium, 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^r, x^r).

PROPOSITION 1. *The Ramsey outcome.*—The Ramsey outcome (π^r, x^r) has first-stage allocations $c_{1t}^r = 0$ and $k_t^r = \omega$ and a capital tax rate $\delta_t^r = (R - 1)/R$. Second-stage allocations c_{2t}^r and l_t^r and the labor tax rate τ_t^r solve

$$U^r = \max U(c_2, l) \tag{1}$$

subject to

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

$$-\frac{U_l}{U_c} = 1 - \tau,$$

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

Proof. Consider the allocation rule for capital $k_t(\pi)$. If the tax rate on capital δ_t is strictly greater than δ^r, then consumers save zero; if $\delta_t = \delta^r$, consumers are indifferent among all levels of savings; and if $\delta_t < \delta^r$, consumers save their entire endowments. For now, assume that when $\delta_t = \delta^r$, the allocation rule specifies $k_t(\pi) = \omega$. The tax on capital acts like a lump-sum tax when it is selected at any level less than or equal to δ^r. 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 $\delta^r R \omega$; therefore, it is optimal to set δ_t so that $(1 - \delta_t)R = 1$. Faced with this tax, consumers save their entire endowments. Given these facts, the optimal tax problem reduces to choosing c_2, l, and τ to solve (1).

Now suppose that when $\delta_t = \delta^r$, $k_t(\pi)$ equals some number α with $0 < \alpha < \omega$. With such a rule, the government can increase its utility by setting δ_t arbitrarily close to but smaller than δ^r and by setting τ_t close enough to τ^r 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 $k_t(\pi)$ is inconsistent with equilibrium. Q.E.D.

Notice that the Ramsey 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 that 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 Sec. III we characterize conditions under which a date 0 competitive equilibrium can be supported by a sustainable equilibrium.)

B. No Commitment

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 at which the decision is made.

At the first stage of period t, faced with the history $h_{t-1} = (\pi_s | s = 0, \ldots, t - 1)$, each consumer chooses a first-stage allocation $f_{1t}(h_{t-1})$ and a contingency plan for setting future actions for all possible future histories. After the first-stage decisions have been made, the government, faced with the history h_{t-1}, sets time t tax rates $\sigma_t(h_{t-1})$ and chooses a contingency plan for setting future tax rates for all possible future histories. At the second stage of t, an individual's history is $h_t = (h_{t-1}, \pi_t)$. 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. (The reader may wonder why the histories do not include consumers' decisions. In an earlier version of this paper, we did define histories that way, but it turns out that we do not need to: no individual consumer perceives that the government or other consumers will change policies if that consumer changes his or her decision. See also Sec. V, where we show that deviations by consumers can be ignored in a game.)

To define a sustainable equilibrium, we now need to explain how policy plans induce future histories and how policy plans together with allocation rules induce future utilities. For any policy plan $\sigma = (\sigma_0, \sigma_1, \ldots)$, let $\sigma^t = (\sigma_t, \sigma_{t+1}, \ldots)$ denote a sequence of policy rules from time t onward. 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 induces future histories by $h_t = (h_{t-1}, \sigma_t(h_{t-1}))$, and so on. Given a history h_{t-1}, the policy plan σ and the allocation rule f induce future utilities

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

where future histories are induced by σ from h_{t-1}.

Consider the first stage of period t. Given some history h_{2t}, the consumer's problem is to choose a sequence of allocation rules to maximize (2) 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. For any history h_{t-1}, the consumer's problem at the second 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 (2) subject to its budget constraints

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

where, for all $s \geq t$, the future histories are induced by σ.

A *sustainable equilibrium* is a pair (σ, f) that satisfies the following conditions: (i) Given a policy plan σ, the continuation of the allocation rule f solves the consumer's problem at the first stage for every history h_{t-1}, and the continuation of this allocation rule solves the consumer's problem at the second stage for every history h_t; (ii) given an allocation rule f, the continuation of the plan σ solves the government's problem for every history h_{t-1}.

Note that consumers take the evolution of future histories as unaffected by their actions and, in this sense, behave competitively. The government recognizes the effect of its policies on the histories and thus on the decisions of private agents and, in this sense, does not behave competitively.

III. 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. When we start 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 construct a sustainable equilibrium that we call the *static equilibrium*. (This equilibrium is static in the sense that the allocation rules and policy plans do not depend on the past history.) 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 reverting to this static equilibrium after deviations. We use this result to show that an arbitrary sequence of policies and allocations is an out-

come 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 static equilibrium (σ^s, f^s), which is defined as follows. For any history h_{t-1}, the static policy plan sets $\sigma_t^s(h_{t-1})$ equal to the policy $\pi^s = (\tau^s, \delta^s)$, where $\delta^s = 1$ and τ^s is given in the solution to the problem

$$U^s = \max_{\tau,\, l} U(\omega + (1 - \tau)l, l) \qquad (3)$$

subject to $-U_l/U_c = 1 - \tau$ and $g \leq \tau l$.

The static allocation rule f^s 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 $k^s(h_{t-1}) = 0$ and $c_1^s(h_{t-1}) = \omega$. For the second stage, given any history h_t, this rule sets the allocations equal to $f_2^s(\pi_t)$ defined by the solution to the problem: choose c_2 and l to maximize $U(\omega + c_2, l)$ subject to $c_2 \leq (1 - \tau_t)l$.

It is immediate that (σ^s, f^s) is a sustainable equilibrium. In particular, given a policy σ^s that specifies a capital tax rate of one, it is optimal for consumers never to save. For the second stage the consumer's problem reduces to the static problem used to define f_2^s. Next, given that the consumer's allocation rule specifies zero savings for all future histories regardless of the past policies of the government, it is optimal for the government to tax capital at rate 1, and the optimal labor tax problem reduces to the static problem used to define τ^s. It is also immediate that the outcomes generated by (σ^s, f^s) are the unique sustainable outcomes when the horizon is finite.

By construction, it follows that in each period of the static equilibrium the realized level of utility is U^s. From proposition 1 we know that in each period of the Ramsey equilibrium the realized utility is U^r. We have the following lemma.

LEMMA 1. The utility level in the Ramsey equilibrium is strictly greater than the utility level in the static equilibrium; that is, $U^r > U^s$.

Proof. A comparison of the Ramsey tax problem (1) with the static tax problem (3) shows that the static tax problem is simply the Ramsey problem with a larger level of g, namely, $\hat{g} = g + (R - 1)\omega$. It follows that the value of the Ramsey problem is strictly greater than that of the static problem. Q.E.D.

In the next lemma we show that the static equilibrium is the worst sustainable equilibrium. Proving this is the key to our method of characterizing the set of sustainable allocations. Let us denote the value of utility in a sustainable equilibrium (σ, f) by $V_0(\sigma, f)$. We have lemma 2.

LEMMA 2. The autarky equilibrium is the worst sustainable equilib-

rium. That is, for any sustainable equilibrium (σ, f), $V_0(\sigma, f) \geq V_0(\sigma^s, f^s)$.

Proof. For a given sustainable equilibrium (σ, f), we construct a plan $\bar{\sigma}$ that satisfies $V_0(\sigma, f) \geq V_0(\bar{\sigma}, f) \geq V_0(\sigma^s, f^s)$. 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

$$U^d(k_t) = \max_{\delta, \tau, c_2, l} U(\omega - k_t + c_2, l) \qquad (4)$$

subject to

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

$$-\frac{U_l}{U_c} = 1 - \tau,$$

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

where k_t is given by $f_{1t}(h_{t-1})$. Next note that 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 $f_2(k_t, \pi_t)$ that solves the static problem

$$\max_{c_2, l} U(\omega - k_t + c_2, l) \qquad (5)$$

subject to $c_2 \leq (1 - \delta_t)Rk_t + (1 - \tau_t)l_t$.

Now by construction of $\bar{\sigma}$ and since f_{2t} solves the static problem, it is clear that $\bar{\sigma}$ is feasible for any such allocation rule. Thus optimality of the government implies $V_0(\sigma, f) \geq V_0(\bar{\sigma}, f)$.

We now show that $V_0(\bar{\sigma}, f) \geq V_0(\sigma^s, f^s)$. We argue that the utility realized under the plans $\bar{\sigma}$ and f is at least as high as under the static plans σ^s and f^s. Let $\bar{h}_t$ denote the history induced by $\bar{\sigma}$. For any t such that $f_{1t}(\bar{h}_{t-1})$ specifies zero savings, the time t utility coincides with that of the static plan. For any t such that $f_{1t}(\bar{h}_{t-1})$ specifies positive savings, the time t utility exceeds that of the static plan. In any such period the government will collect a strictly positive amount of revenue using what is essentially a lump-sum tax on capital, so welfare is higher. Since this argument holds for any period t, welfare under $(\bar{\sigma}, f)$ must be at least as high as it is under (σ^s, f^s), where all revenue is raised through the distortionary labor tax. Q.E.D.

In the next proposition—which is the paper's main result—we characterize the conditions under which an arbitrary sequence of allocations and policies is sustainable. To prove the proposition we use a modified version of the static plans called the *revert-to-static* plans. For an arbitrary sequence of policies (π, x), the revert-to-static policy plans specify continuation with the candidate sequences (π, x) as long as the specified policies have been chosen in the past. If there has ever been a deviation, the government's plan specifies to revert to the static plan

σ^s. For consumers the allocation rule specifies, immediately after a deviation, to follow the second-stage allocation rules defined in (5) and to revert to the static allocation rules f^s in all subsequent periods. We then have proposition 2.

PROPOSITION 2. *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:

$$\sum_{s=t}^{\infty} \beta^{s-t} U(c_{1s} + c_{2s}, l_s) \geq U^d(k_t) + \frac{\beta}{1-\beta} U^s, \qquad (6)$$

where $U^d(k_t)$ is defined in (4).

Proof. Suppose, first, that (π, x) is the outcome of a sustainable equilibrium (σ, f). Consumer optimality requires that (π, x) maximize consumer welfare at date 0. Government optimality implies that (π, x) satisfies the government's budget constraint at date 0. Thus (π, x) is a date 0 competitive equilibrium. Next, at time t, given a history h_{t-1}, a deviation to the plan $\bar{\sigma}$ defined in lemma 2 is feasible. Under this deviation, the time t utility is $U^d(k_t)$, as defined in lemma 2, and for any $s > t$, lemma 2 guarantees that the time s utility is at least U^s. Clearly, then, the utility of the government must be at least as large as the right side of (6) for every period t. Thus conditions i and ii hold.

Next, suppose that some arbitrary pair of sequences (π, x) satisfies conditions i and ii. We show that the associated revert-to-static plans constitute a sustainable equilibrium. Consider histories under which there have been no deviations from π up until time t. Since (π, x) is a date 0 competitive equilibrium, it is obvious that its continuation from time t is optimal for consumers. 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 (6). Since the policy plan $\bar{\sigma}$ was constructed to maximize time t utility for any k_t, the maximal utility attainable under any deviation at t is simply the right side of (6). Hence, given that the assumed inequality holds, sticking with the specified plan is always optimal.

Now consider histories for which there has been a deviation before time t. The revert-to-static rules specify (σ^s, f^s) from date t onward. Such rules are clearly optimal for both the consumers and the government. For histories in which the first deviation is at time t, the second-stage allocation rules of consumers are optimal by construction. Q.E.D.

Proposition 2 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 sustain-

able equilibrium. It is worth noting that some competitive equilibria cannot be the outcome of any sustainable equilibrium. For instance, consider an equilibrium with the tax on capital identically equal to one and with the tax on labor inefficiently high (e.g., let the tax on labor be on the far side of the Laffer curve). Clearly, this equilibrium generates lower utility than the static equilibrium and, thus, is not sustainable. Notice that this equilibrium cannot be sustained for any discount factor in the unit interval.

With proposition 2, it can be shown 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. We have the following proposition.

PROPOSITION 3. *Sustainability of Ramsey allocations.*—There is some discount factor $\beta \in (0, 1)$ such that, for all $\beta \in [\underline{\beta}, 1)$, the Ramsey allocations are sustainable.

Proof. From proposition 2 it suffices to show that the inequality (6) holds for the Ramsey allocations. Thus to prove the result, it suffices to verify the inequality

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

Rearranging terms gives

$$\frac{\beta}{1 - \beta} (U^r - U^s) \geq U^d(k^r) - U^r. \tag{7}$$

From lemma 1, the left side of (7) is strictly positive. Thus there is some $\beta < 1$ such that this inequality holds for all $\beta \geq \underline{\beta}$. Q.E.D.

Two remarks about propositions 2 and 3 are warranted. First, it is immediate from the proof of proposition 3 that the Ramsey allocations are not sustainable for any $\beta \in (0, \underline{\beta})$. 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. That is, given any number U satisfying $U^s < U < U^r$, 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 static and the Ramsey allocations and by choosing the discount factor to be high enough, any such utility level can be sustained.

IV. An Example

We present an example to illustrate four features of sustainable outcomes and their associated utility levels that follow from propositions

FIG. 1.—The set of stationary sustainable utility levels

2 and 3. First, for low enough values of the discount factor, the only sustainable outcome is the static outcome. 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 the static and the Ramsey utilities 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 (6) reduce to the single inequality

$$U(c_1 + c_2, l) \geq (1 - \beta)U^d(k) + \beta U^s. \qquad (8)$$

To characterize the set of utilities that satisfy (8), it suffices to consider outcomes in which $\delta = (R - 1)/R$, the tax on labor is set optimally, and 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, $E(\beta) = \{U(k) | U(k) \geq (1 - \beta)U^d(k) + \beta U^s,$ for some $k \in [0, \omega]\}$. Let the utility function be $U(c_1 + c_2, l) = [(c_1 + c_2)^\alpha + \gamma(\bar{l} - l)^\alpha]^{1/\alpha}$ and let $\alpha = -0.3, \gamma = 1.2, \bar{l} = 100, \omega = 10, g = 25,$ and $R = 2$.

For this example, the set of stationary sustainable utility levels illustrates the four features (see fig. 1). First, for $\beta < 0.1, E(\beta) = U^s$. Second, $E(\beta) \subset E(\beta')$ for $\beta < \beta'$. Third, for $\beta \geq 0.1, U^r \in E(\beta)$.

Fourth, for $\beta \geq 0.48$, $E(\beta) = [U^s, U^r]$. 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 static utility are not sustainable (at least with stationary outcomes).

V. Anonymous Games

In this section we provide one rationalization of the equilibria considered in the previous sections, but in a game-theoretic context. We first 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 corresponds to the set of symmetric perfect Bayesian equilibria of a game with no commitment.

In the economies considered earlier, we modeled private agents as behaving competitively, in the sense that each private agent assumes that his decisions can affect neither the government's policies nor any other private agents' decisions. We capture this feature in a game by using two assumptions. First, we assume that there is a continuum of agents. Second, we assume that individuals observe only their own decisions and aggregate outcomes. A game with these features 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 $x = (k, l)$: $[0, 1] \rightarrow [0, \omega] \times R_+$. We denote the implied action of an individual agent i by $x(i) = (k(i), l(i))$. The single-period payoffs of agent i are

$$V_i(\pi, x(i), x)$$

$$= U(\omega - k(i) + (1 - \delta)Rk(i) + (1 - \tau)l(i), l(i)) + W(\delta, \tau, K, L),$$

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

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. Payoffs over strategies are defined as the discounted utility of the outcomes they induce.

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 the following 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; (ii) given the strategy profile f, the strategy π maximizes the government's payoff. Comparing this definition with the Ramsey equilibrium of Section II gives the following proposition.

PROPOSITION 4. *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 4 is given in Chari and Kehoe (1989). The requirement of subgame perfection is crucial in this proposition. Indeed, it is easy to see that the set of Nash equilibria is considerably larger than the set of subgame perfect equilibria. 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 follows that 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. 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 $h_{1t} = (x_s, \pi_s | s < t)$ and at the second stage it is $h_{2t} = (h_{1t}, x_{1t}, \pi_t)$. 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 and past policies. A history for the government at time t is $H_t = ((X_s, \pi_s | s < t), X_{1t})$. Players' histories correspond to information sets in the obvious 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

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

where the future actions are induced from h_{1t} by f and σ. The payoff for the government at time t is similarly defined.

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. (It should be clear that a large number of rather bizarre subgame perfect Nash equilibrium outcomes exist for this game. Thus it is incorrect to say that dynamic consistency is equivalent to subgame perfection.) An alternative is to consider a type of Bayesian equilibrium (see, e.g., 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, the expected utility of player i at the information set associated with history $h_{1t}(i)$ is $\int W_{it}(\sigma, f(i), f; h_{1t}) d\mu(h_{1t}|h_{1t}(i))$. We can similarly define the expected utility for the government at the information set associated with a history H_t and the payoffs for players at the second stage associated with a history $h_{2t}(i)$.

In the equilibria of Sections II and III, we used a representative agent to model the private agents. 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.

A *symmetric 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; (iii) the strategies of consumers are symmetric; and (iv) μ assigns probability one to symmetric histories.

To understand condition iv, 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 iv requires that $\mu(\cdot|h_{1t}(i))$ assign probability one to the symmetric history of the game associated with $h_{1t}(i)$, namely, to the history $h'_{1t} = (x'_s, \pi'_s|s < t)$ 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, this condition requires that, at any information set, player i believe that all the other private agents have behaved symmetrically in the past. Similarly, it requires that the government believe that all private agents have behaved symmetrically in the past.

PROPOSITION 5. *Equilibrium outcomes of the no-commitment game.*—The set of symmetric perfect Bayesian equilibrium outcomes of the no-commitment game is the same as the set of sustainable equilibrium outcomes.

Here we provide an intuitive explanation of the proposition. A formal proof is presented in Chari and Kehoe (1989). The essential difference between the definitions of a sustainable equilibrium and a symmetric perfect Bayesian equilibrium is that the latter requires optimality after histories with private deviations, whereas the sustainable equilibrium does not even consider such histories. To prove the proposition, we need to extend the functions that constitute a sustainable equilibrium to include a larger set of histories. For histories of the game in which no positive measure of agents have deviated, extend these functions for the government and the nondeviating private agents to be the same as if no private agents had deviated. Let the deviating private agents act optimally given their histories. For histories in which a positive measure has deviated, it really does not matter how we extend these functions as long as the continuation strategies by themselves form a Bayesian equilibrium. In particular, we can let them equal the analogues of the static sustainable equilibrium. Intuitively, the reason this extension works is that in our anonymous game the deviations of any single private agent do not influence the future behavior of other agents.

Symmetry and anonymity both play crucial roles in the proof of proposition 5. First, the set of perfect Bayesian equilibria is larger than the set of symmetric perfect Bayesian equilibria. For example, since consumers are indifferent among all saving levels when $(1 - \delta)R = 1$, we can have asymmetric equilibria in which some consumers save all their endowments 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 role of 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 use trigger strategies to support the equilibrium allocations obtained with lump-sum taxation. The strategies specify that in every period, agents save their entire endowments and supply the optimal amount of labor with lump-sum taxation 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 private actions are unobservable and have no effect on aggregate

outcomes implies 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.

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. These differences lead to different results. For example, Fudenberg and Maskin 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. When $U(\omega, 0)$ is normalized to be zero, it is clear that the mutual minimax payoffs are $-M$. (Each player saves nothing and does not work, and the government cannot meet its budget constraint.) In our model, regardless of the discount factor, no average utility that is lower than the static utility (some positive number) can be supported. Technically, our model gives rise to payoffs that do not satisfy Fudenberg and Maskin's "full-dimensionality" condition.

VI. Conclusion

We wrote this paper to address four related questions: (1) 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? (2) If so, precisely what is the equilibrium concept; in particular, what are the decision problems of private agents? (3) Is it possible to characterize *all* the equilibria? (4) 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. (For a further discussion of these issues, see Chari, Kehoe, and Prescott [1989].)

In some part of the macro literature/verbal tradition, we have heard expressed the (admittedly fuzzy) idea that Ramsey equilibria and time-consistent equilibria can be thought of as two different equilibrium concepts for a single policy game. These are supposed to correspond, respectively, to imperfect and perfect Nash equilibria of

that game. An important message of this paper is that the distinction between Ramsey equilibrium and time-consistent equilibrium is not perfection since both require it; rather, the distinction is that they are equilibria of two very different games.

In terms of comparing our approach with the standard game-theoretic one, it should be clear that ours is much simpler. Furthermore, since the sustainable equilibrium outcomes coincide with the equilibrium outcomes of the game, our approach does not miss anything essential in the game. We believe that the equivalence between the sustainable equilibria and the Bayesian equilibria of an appropriately defined anonymous game holds true for a wide variety of macroeconomic models. In particular, we have shown this equivalence for a model with debt (Chari and Kehoe 1989). Our characterization theorems should also apply to a variety of models.

References

Abreu, Dilip. "Extremal Equilibria of Oligopolistic Supergames." *J. Econ. Theory* 39 (June 1986): 191–225.

————. "On the Theory of Infinitely Repeated Games with Discounting." *Econometrica* 56 (March 1988): 383–96.

Atkeson, Andrew G. "Time Consistency of Policy in a Macro Model." Manuscript. Stanford, Calif.: Stanford Univ., Grad. School Bus., 1988.

Backus, David K., and Driffill, John. "Inflation and Reputation." *A.E.R.* 75 (June 1985): 530–38.

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." Working Paper no. 354. Minneapolis: Fed. Reserve Bank Minneapolis, Res. Dept., 1989.

Chari, V. V.; Kehoe, Patrick J.; and Prescott, Edward C. "Time Consistency and Policy." In *Modern Business Cycle Theory*, edited by Robert J. Barro. 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." Manuscript. Cambridge: Massachusetts Inst. Tech., 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.
Rogoff, Kenneth. "Reputational Constraints on Monetary Policy." In *Bubbles and Other Essays,* edited by Karl Brunner and Allan H. Meltzer. Carnegie-Rochester Conference Series on Public Policy, vol. 26. Amsterdam: North-Holland, 1987.