CONTINGENT PLANS, TIME CONSISTENCY AND PARETIAN RULERS

Daniel B. Klein
Department of Economics
UC-Irvine
Irvine, CA 92717

Brendan O'Flaherty
Department of Economics
Columbia University
New York, NY 10027

October 1991

Abstract: If there are no restrictions on the extent to which government plans can be made contingent, then there is some Paretian government with an optimal plan that never reaches a node at which deviation from the plan would yield a Pareto improvement. (At unreached nodes deviations from the optimal plan may be Pareto improving, so the need for commitment remains.) Examples like Fischer's [1980], where optimal plans are time inconsistent, arise because of restrictions on the contingency of government planning. Time inconsistency can be avoided if planning can be made sufficiently contingent.

Keywords: S-game, S-game set, optimal plan, time inconsistency, policy formation, Paretianism.

J.E.L. classification numbers: 025, 026.

The authors wish to thank Stergios Skaperdas for helpful comments.

1. INTRODUCTION

Consider Fischer's model of savings and taxation [1980]. In the first period citizens decide whether to consume or save, and in the second period they decide how much to work. At the beginning of the second period the government implements tax rates on labor income and accumulated savings to finance a public good. An uncommitted government will tax savings heavily because there is no second-period distortion there. But foreseeing this, citizens curtail savings in the first period. Everyone is worse off than when the government commits at the outset to a moderate tax rate on accumulated savings.

What's more, if the government conveys a commitment to moderate (and optimal) taxes on savings, inducing people to save, the economy reaches a point at which *everyone* prefers that the government dissolve its apparent commitment.

The fascination of the model arises from the juxtaposition of the following three elements:

a. The government's objective is the maximization of the utility of the representative agent. This strongly suggests commonality of interest.

b. The need for commitment. This suggests some measure of conflict.

c. The final unanimity on renegotiation (or inconsistency). This again suggests commonality of interest.

The apparent paradox between (a) and (b) is explored in Hillier, Klein, & O'Flaherty [1990]. There it is shown generally that if in a Stackelberg game all Paretian rulers would value commitment, then a prisoner's dilemma lurks in the citizen game induced by at least one of the ruler's plans. In terms of the Fischer model the explanation is as follows: Although only one citizen's utility function is written down, there is really a continuum of citizens just like this one. Since outcomes are always symmetric the government is maximizing the sum of welfares when it maximizes the utility of the representative citizen. Hence the scope for conflict. The government would like to prevent one citizen from bettering his lot if the improvement were more than offset by reductions in the welfare of other citizens. When a citizen chooses not to save, he is, in part, free riding on second period public good provision. By promising low taxes on accumulated savings the government induces unwitting cooperation in a prisoner's dilemma.

The goal of this paper is to provide a general account of item (c) -- why an economy can reach a point at which everyone would prefer a different finish. For the Fischer model we again can give a specific explanation. The question is, If everyone prefers the inconsistent finish off the optimal plan, why doesn't the government simply plan and announce that reversion plan from the start? Since the same citizen actions would promise an even greater payoff than they did with the original optimal plan, it would seem that every citizen would take the same actions.

To see why not, we need to analyze the Fischer model in more detail. The best outcome for the government is for all citizens to save, and then for them all to be taxed heavily on their savings. Under what circumstances will it be a jointly best response for them all to save, knowing that their savings will subsequently be taxed? They will all save if each believes that it is very bad not to save. Suppose the government can announce a plan of the following sort: *"If everyone saves, savings will be taxed at the second-period socially optimal rate; and if anyone fails to save everyone will be taxed at an extreme and senseless rate that will make us all miserable."*

Clearly the public will respond to this plan with unanimous saving, and after they have saved the government will have no reason to deviate from the announcement. The optimal plan is time consistent. (Of course, the plan is sequentially irrational. Time consistency permits sequential irrationality off the path of actual play. This paper is about time inconsistency; sequential irrationality, or what we call commitment dominance, is treated in Hillier, Klein, and O'Flaherty [1991].)

How then can play actually reach a node where everyone wants to change? The answer must be that the government cannot announce a plan like the one in italics. Notice how that plan differs from the type Fischer discusses: Fischer implicitly restricts the government to plans that specify a uniform tax policy for any history of savings. In other words, Fischer implicitly imposes a restriction that the government's action must be the same on a large set of nodes. We call such a restriction "imperfect policy formation."

What motivates imperfect policy formation? The standard game theoretic reason why a player would have to take the same action at different nodes -- namely, imperfect information -- does not apply: if a government cannot observe individual citizen's savings, how can it tax them? The appropriate motivation is to realize that the real issue is not government action but government commitment conveyance, and there may be limits on the complexity of the commitments that the government can convey. Although anyone who has observed the hundred or so volumes of the U.S. Code Annotated may have doubts about this motivation, there is much intuition and a long tradition behind the assumption that government may face coarseness constraints in laying down policy.

The purpose of this note is to generalize the result we have just shown for the Fischer model: if the government planning can be made perfectly contingent then there is some Paretian government whose optimal plans are time consistent. (The condition does not imply, however, that there is some Paretian government whose optimal plan is what we term "sequentially rational.") These results suggest that a government can avoid time inconsistency if it can refine its policy making.

Section 2 presents the Fischer model in more detail to give an entry to our results. Section 3 presents the machinery used in the paper, Section 4 gives our results, and Section 5 provides som discussion and Section 6 concludes.

2. THE FISCHER MODEL IN ARBORESCENT FORM

If the government conveys a commitment to tax rates before the savings decision the literature calls it the "rules" regime. The rules regime commit only to a tax plan that levies the same pair of tax rates at any "second period" node. That is, the government cannot make its tax plan contingent on individual citizen behavior.

Instead, the regime is called "discretionary" if it is common knowledge that the government reoptimizes given the history of citizen decisions. In our language, the discretionary regime is a ruler restricted to sequentially rational plans.

We know that the discretionary regime will opt to tax savings heavily because of the absense of distortion in levying from accumulated savings. Foreseeing this, citizens curtail savings in the first period. Under the rules regime everyone is better off because the government commits to moderate taxes on savings, inducing greater savings and hence more second-period consumption and public good provision.

In Figure 1 we depict the Fischer story with an arborescence, but we ask the reader to view this game-tree through an untraditional set of lenses. There are two citizens (players 2 & 3) and a utilitarian government (player 1). Although Fischer's model uses continuous variables, we need only depict the choice levels that arise under the various regimes. The figure shows the two citizens simultaneously choosing between the Low savings induced by the discretionary regime and the High savings induces by the rules regime. (The hoop depicts player 3's information set.) The combinations of their choices imply four nodes for the government, which, in keeping with the literature, form a single "policy set" (shown by the dotted line). This dotted line signifies that the rules regime must make its choice *uniform* over the nodes in the set.

[Figure 1 here.]

Each action at a government node represents a tax-rate pair for labor and accumulated savings. D subscripts are for "discretion;" R for "rules." From left to right the four actions are:

t_{D1} - the sequentially rational tax-rate pair when each citizen has chosen Low savings.

t_{D2}- the sequentially rational tax-rate pair when one citizen has chosen Low savings and the other has chosen High savings.

t_{D3} - the sequentially rational tax-rate pair when both citizens have chosen High savings.

t_R - the optimal tax-rate pair when the government enjoys commitment conveyance.

After tax rates are verified individuals make a labor decision, which we do not depict because a unique equilibrium set of labor decisions will be implied by each history.

The discretionary regime cannot convey a commitment and, hence, is restricted to sequentially rational play, shown in Figure 1 by the double arrows. The citizens respond with Low savings and they arrive at terminal node D.

The rules regime announces t_R (across the board), eliciting High savings from the citizens and outcome R, which Pareto dominates the discretionary outcome D. Notice the time inconsistency at government node z: if it could do so costlessly, the government would dissolve the apparent commitment to t_R and instead revert to t_{D3}. Ironically, the citizens would welcome this breach of contract.

Now, this irony stems from imperfect policy formation. If we removed the dotted line from the figure, making policy formation "perfect," the ruler could announce optimal plan (t_R, t_R, t_R, t_{D3}) and receive a payoff of 20. This plan would be *time consistent.* (It would, however, be "sequentially irrational.") This example is generalized in what follows. We show that if policy formation is perfect there is always some Paretian ruler whose optimal plan is time consistent. Imperfect policy formation lies at the heart of Paretian time inconsistency.

3. NOTATION AND DEFINITIONS

(2.1) An <u>S-game</u> is a four-tuple $\Sigma = (G, i, s, (U, C))$. S-games get their name from Stackelberg.

(2.2) <u>The reference game</u>, G, denotes an extensive form game.

(2.3) <u>The ruler, i, and the public</u>. Let I denote the index set of players of G, where $|I| = m+1$. Player i is called the <u>ruler</u>, and the set $I\backslash\{i\}$ of other players is called the <u>public</u>. For the remainder of this paper we let player 1 be the ruler (or i=1). The set of nodes assigned to the ruler is R; the set of nodes assigned to the public is P.

(2.4) <u>Information assumption on G</u>. Throughout this paper we assume: a) that every ruler information set is a singleton, and b) that at every information set in G all previous play by the ruler has been observed publicly. These are rather natural assumptions since the ruler announces her plan in advance of play and since she knows how the public will respond. These assumptions ensure that a subgame originates at every ruler node.

(2.5) <u>Plans</u>. Let B' denote the set of player i's behavior strategies in G. In the S-game Σ, a subset B of B' is the set of permissible <u>plans</u> for the ruler, with generic element b. Subsection (2.13) explains why the ruler may not have access to the complete set of behavior strategies, B'. Think of a plan as an announcement of what the ruler will do at each node $v \in R$.

(2.6) Public response function, s. Let S denote the set of behavior strategy m-tuples that exist in G for the public. Let s(b), the public response function, denote the unique element of S that is picked out when plan b is announced. For the remainder of the paper we assume that, for every b, s(b) is a subgame perfect equilibrium of the m-player game induced by b in G.

(2.7) Payoff functions. $h_j(b, d)$ denotes the payoff to player j when the ruler uses plan b and the public uses behavior strategy m-tuple $d \in S$. The ruler's payoff function is often written $h_R(b)$, which denotes the ruler's payoff when she plays b and the public plays s(b).

(2.8) Subgames. For each node $v \in R$ let G(v) denote the subgame whose origin is v. Let b(v) denote the behavior strategy induced by plan b on G(v), and denote the player j's payoff function on G(v) as $h_{jv}(.)$. Let s(b(v)) denote the behavior strategy m-tuple induced on G(v) by s(b).

(2.9) Sequential rationality at a node. For any node $v \in R$, let b_v denote the local strategy specified by b at v. Let $b|b_v'$ denote the ruler's behavior strategy that results if the local strategy assigned by b to node v is changed to b_v' while the local strategies assigned by b to other nodes in R remain unchanged.

A plan b is sequentially rational at node v if for every local strategy b_v' at v

$$h_{Rv}(\, b(v), s(b(v))\,) \;\geq\; h_{Rv}(\, b(v)|b_v', s(b(v)|b_v')\,)$$

A plan b is sequentially rational if this relation holds for every $v \in R$. Note that our usage of this term is distinct from that of Kreps & Wilson [1982].

(2.10) Time consistency. The concatenated behavior strategy for the ruler $\kappa(b, v, b')$ is the behavior strategy that results from behavior strategy b if the behavior strategy induced by b on G(v), $v \in R$, is changed to b'(v) while the local strategies assigned by b to other nodes in R remain unchanged. Similarly, the concatenated behavior strategy m-tuple for the public $\gamma(d, v,$

d') is the behavior strategy m-tuple that results from behavior strategy m-tuple d if the behavior strategy m-tuple induced by d on G(v), $v \varepsilon R$, is changed to d'(v) while the local strategies assigned by d to other vertices in P remain unchanged.

A plan b is time inconsistent iff there exists a $v \varepsilon R$ and some other plan b' such that

$$h_R(\kappa(b, v, b'), \gamma(s(b), v, s(b'))) > h_R(b). \quad (1)$$

A plan that is not time inconsistent is called "time consistent." We say that a ruler "faces "time inconsistency" iff every optimal plan is time inconsistent.

(2.11) Policy formation. In Fischer's saving-taxation model [1980], a government commitment to tax policy cannot take the form of any reaction function (with domain being the history of citizen saving decisions). In the time inconsistency literature, the government is restricted to plans that take the form of a single magnitude that applies uniformly across all possible citizen histories standing at the government's "time to act." We need to account for the possible coarseness of policy formation.

Policy formation is described by a pair (U, C). U is a partition of the ruler's nodes into eligible subsets. We use the term eligible in Selten's sense [1975, 26]: a set u of ruler nodes is eligible "if every play intersects u at most once, and if the number of alternatives at [v] is the same for every [v] ε u." We call these eligible subsets policy sets. Loosely speaking, a policy set is a set of nodes that the ruler must treat similarly when conveying a commitment. It may be useful to think of a policy set as a "point in time."

For any $u \varepsilon U$, let A_u be the set of all alternatives at nodes $v \varepsilon u$. C is a partition of all the alternatives at all the ruler nodes; specifically, C partitions these alternatives into subsets c of all the A_u. Each of these subsets, c, must be eligible in Selten's sense [1975, 26]: a subset c of A_u is eligible "if it contains exactly one alternative at [v] for every [node v] ε u."

Intuitively, a subset u is a policy set and a subset c is a choice at a policy set. The idea of c is to define the choosing of the "same action" at every node in a policy set, as, for example, the taxing authority may be restricted to announcing a single tax rate that will prevail in the second period, no matter what history (or node) actually transpires.

A plan b for the ruler is eligible iff for every policy set $u \varepsilon U$ either

(1) the set of alternatives it specifies is an element of C,

or, (2) it is sequentially rational at every node $v \varepsilon u$.

Loosely speaking, programmers can think of option (1) as the "open-loop" option and option (2) as the "closed-loop" option.

B is the set of eligible plans. We depict policy sets by connecting ruler nodes in a single policy set with a dashed line. If every policy set of an S-game is a singleton we say that the S-game satisfies <u>perfect policy formation</u>. Otherwise policy formation is <u>imperfect</u>.

As an illustration of policy formation, return to the Fischer model of Figure 1. All of the ruler's node are contained in a single policy set. Double arrows show option (2) ("closed-loop"), while single arrows are one example of option (1) ("open loop").

(2.12) <u>Paretianism</u>. We say that strategy combination (b, d) <u>strongly Pareto dominates</u> strategy combination (b', d') iff hj(b, d) > hj(b', d') for every citizen j. We shall say that a ruler is <u>Paretian</u> iff the ruler's payoff function respects the strict Pareto orderings of public welfare -- that is, $h_R(b, d) > h_R(b', d')$ whenever (b, d) strongly Pareto dominates (b', d'). Our results depend on our use of strict orderings -- otherwise ties can upset the results. (Such an example is available from the authors.)

(2.13) <u>S-game set</u>. This paper explores time inconsistency faced by Paretian rulers. Yet since Paretianism is no guide on distributional issues, a particular Paretian ruler may face particular problems that are not essential to Paretianism. For this reason we must couch our discussion in terms of time consistency faced by *every* Paretian ruler, given all the features of the S-game besides the ruler's payoff function. Like plugging different sockets on to the head of a rachet wrench, we must think of plugging different Paretian payoff functions on to the ruler's head of the S-game. Thus, Π is some S-game Σ except that blank slips of paper have been placed over the ruler's payoffs at the terminal nodes. We call Π an S-game set. Let $\Pi(h_R)$ denote the complete S-game that results when the ruler payoff function h_R is plugged on to Π.

(2.14) Optimal plans. A plan b* is optimal iff for every b, $h_R(b^*) \geq h_R(b)$.

(2.15) Potential time consistency. An S-game set Π is potentially time consistent if there is some Paretian ruler h_R who has at least one plan in $\Pi(h_R)$ that is both optimal and time consistent. A failure of potential time consistency means that all Paretian rulers are doomed to time inconsistency.

3. RESULTS

The Lemma says that perfect policy formation precludes Pareto improving concatenations.

Lemma: Let Σ be an S-game that has a Paretian ruler with an optimal plan b*. *If Σ has perfect policy formation, then there is no concatenation $\kappa(b^*,v,b')$ that strongly Pareto dominates b*.*

Proof: Suppose there were such a concatenation. That is, suppose there were a b' and a ruler node v such that (b*, s(b*)) was strongly Pareto dominated by

$$\big(\ \kappa(b^*, v, b'),\ \gamma(s(b^*), v, s(b'))\ \big)$$

In this proof we refer to this strategy profile as "(κ, γ)."

Let $h_{jv}(b(v), s(b(v)))$ denote the payoff to player j in the subgame G(v) when the strategies used are (b(v), s(b(v))).

Consider the S-game that results when we trim off the subgame G(v) and make v a terminal node with citizen payoff vector:

$$\big(\ h_{1v}(b^*(v), s(b(v))),\ \ h_{2v}(b^*(v), s(b(v))), \ldots,\ h_{mv}(b^*(v), s(b(v)))\ \big).$$

Upon assimilation of our notation it should be obvious that in this trimmed S-game the public response to b-minus-b(v) would be s(b)-minus-s(b(v)).

Now consider the trimmed S-game where the ruler node v shown in (κ, γ) becomes a terminal node with the citizen payoff vector:

$$\Big(\ h_{1v}(\ b'(v),\ s(b'(v))\ ),\ \ h_{2v}(\ b'(v),\ s(b'(v))\ ),\ \ldots,\ h_{mv}(\ b'(v),\ s(b'(v))\ )\ \Big).$$

The assumption that (κ, γ) strongly Pareto dominates (b*, s(b*)), in conjunction with the no-ties condition, implies that

$$h_{jv}(\ b'(v),\ s(b'(v))\ )\ >\ h_{jv}(\ b^*(v),\ s(b^*(v))\ ), \quad \text{for all } j = 1,2,\ldots m.$$

In this trimmed S-game, then, the public response to b*-minus-b*(v) is s(b*)-minus-s(b*(v)). In the untrimmed S-game, then, s(b*, v, b') is identical to γ(s(b*), v, s(b')). Therefore, since (κ, γ) strongly Pareto dominates (b*, s(b*)), so does $\big(\ \kappa$(b*, v, b'), s(κ(b*, v, b')) $\big)$. Therefore the ruler should strictly prefer using κ(b*, v, b') over using b*, contradicting the optimality of b* .

QED

The Lemma is of some direct interest. Its contrapositive says that if concatenated play off an optimal plan for a Paretian ruler yields a Pareto improvement, policy formation must be imperfect. Pareto improving concatenations are precisely what we encounter in the Fischer model or the Kydland & Prescott model of the Phillips Curve [1977], so immediately we know those models must involve imperfect policy formation.

But the main interest in the Lemma is its use in proving the following Proposition, which says that if policy formation is perfect there is some Paretian ruler whose optimal plan is time consistent.

Proposition: Let Π be some S-game set. *If Π satisfies perfect policy formation, then Π is potentially time consistent.*

Proof: Let h_R be a Paretian ruler and let b* be one of her optimal plans in the S-game $\Pi(h_R)$. If b* is time consistent we would have a ruler with a plan that is both optimal and consistent.

So suppose that b* is time inconsistent, that is, in at least one instance there is some $v \in R$ and some plan b' such that

$$h_R\big(\, \kappa(b^*, v, b'),\ \gamma(\, s(b^*), v, s(b'))\,\big) \quad > \quad h_R\,(\, b^*, s(b^*)\,) \tag{2}$$

In this proof we refer to any profile of concatenated strategies like that shown on the LHS of (2) as "some (κ, γ)," where the ruler node v and the reversion plan b' are not specific. (In contrast, the same b* is common to all (κ, γ).)

From the Lemma we know that no (κ, γ) strongly Pareto dominates (b*, s(b*)). Hence we can construct a Paretian ruler $h_{R'}$ such that the following two conditions hold:

(a) b* is an optimal plan for $h_{R'}$, since for any other plan b_1 the outcome from $(b_1, s(b_1))$ cannot strongly Pareto dominate (b*, s(b*)), for otherwise we have a contradiction of the combination of assumptions of h_R being Paretian and b* being optimal for h_R.

and, (b) for any (κ, γ) satisfying (2), $h_{R'}(\kappa, \gamma) \leq h_{R'}(b^*, s(b^*))$. This is achievable because no (κ, γ) satisfying (2) strongly Pareto dominates (b*, s(b*)).

We conclude that b* is both optimal and time consistent for some Paretian ruler $h_{R'}$.

QED

The Proposition tells us that, given a mild assumption, perfect policy formation implies that some Paretian ruler has a time consistent optimal plan. If perfect policy formation in a given situation is readily apparent, the feasibility of time consistency is easily checked. More importantly, this result tells us that, in any situation, time inconsistency might be circumvented, while preserving Paretianism, if policy making were more refined.

To get some idea of how little "sufficiently" might mean, return to the Fischer tree in Figure 1. The utilitarian ruler needs to extricate *only node z* from the policy set. The following plan would then be eligible: (t_R, t_R, t_R, t_{D3}). This plan would be optimal -- inducing outcome F -- and time consistent. Hence the degree of contingency that the ruler needs to achieve time consistency may be rather small. Aiming for such contingent policy making would seem to be a fruitful path toward resolution of the time consistency problem.

The Proposition establishes that imperfect policy formation lies at the heart of Paretian time inconsistency. We might ask if we can go further: Does imperfect policy formation imply time inconsistency? Certainly not. The imperfections in policy formation may be of a very inconsequential nature. We can, however, write down particular S-game sets with imperfect policy formation in which every Paretian ruler faces time inconsistency (example available from the authors).

4. DOES TIME CONSISTENCY MATTER? A LOOSE DISCUSSION

The time inconsistency literature tells us that governments may have a hard time committing to an optimal plan precisely because citizens know that, if they were to believe the plan, the government would have a "second period" incentive to cheat. We have shown that with a contingent plan this incentive to cheat would be eliminated. Does this elimination of inconsistency really enhance credibility? Just because contingent planning eliminates sequential irrationalities along the path, aren't the remaining sequential irrationalities off the path still enough to undermine credibility in the plan? Has all the focus in the literature on time consistency, as opposed to sequential rationality, been the proper focus?

Consider Klein's model of a flood plain [1990]. The utilitarian ruler needs to convey a commitment to the following announcement: "I will not bail you out if you move to the plain and suffer flooding." When people believe the announcement they do not move to the plain -- the plan is time consistent. Is this type of plan, by virtue of its time consistency, more likely to elicit citizen belief than the time inconsistent plan in the (noncontingent) Fischer model?

We do not see why sequentially irrational plans that are time consistent should be much more believable than sequentially irrational plans that are time inconsistent. Both types are sequentially irrational, so both require commitment conveyance. Commitments made off the path are generally just as essential as those made along the path. (For an interpretive discussion of commitment conveyance see Klein & O'Flaherty (forthcoming).) It seems to us that sequential irrationality, not the subsumed concept of time inconsistency, deserves center stage. (Klein & O'Flaherty [1991] establish the relationships between sequential rationality, time consistency, subgame perfection, and Nash equilibrium.)

Still, time consistency can't hurt. In the contingent planning case of the Fischer model, where the ruler can announce the time consistent plan, it would seem that this plan would be somewhat more likely to elicit high savings than the time inconsistent plan in the noncontingent case. If citizens have less than complete trust in the ruler's announced sequentially irrational moves, they will find high savings somewhat more attractive with the contingent plan. We have worked out a simple hypothetical belief structure which, for some levels of trust in the ruler's announcements, citizens will choose high savings in the contingent (time consistent) case but low savings in the noncontingent (time inconsistent) case. These considerations need further contemplation.

5. CONCLUSION

Time inconsistency for all Paretian rulers has been shown to depend on imperfect policy formation. This suggests that if government can refine its policy making sufficiently, by making policy more contingent (or history specific), time inconsistency can be avoided.

The programming literature that brought the notion of time inconsistency into economics implicitly uses the idea of imperfect policy formation. In that literature the "open loop" policy is one which specifies actions at a "point in time." In a game depiction this "point in time" corresponds to many possible histories and a whole set of nodes. Having to announce a uniform action across those nodes, or declare sequential reoptimization, is the added restriction on possible plans in Stackelberg games that we capture in our notion of "policy formation."

REFERENCES

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

Hillier, Brian, Daniel B. Klein, and Brendan O'Flaherty [1991]: "Universal Commitment Dominance and Prisoner's Dilemma," Columbia University, Economics Discussion Paper no. 538.

Klein, Daniel B. [1990]: "The Microfoundations of Rules vs. Discretion," *Constitutional Political Economy* 1, 1-19.

____________ and Brendan O'Flaherty [1991]: " Time Inconsistency and Related Concepts," ms.

______________________________(forthcoming): "A Game-Theoretic Rendering of Promises and Threats," *Journal of Economic Behavior and Organization.*

Kreps, David M. & Robert Wilson [1982]: "Sequential Equilibrium," *Econometrica*, 50, 863-94.

Kydland, Finn and Edward Prescott [1977]: "Rules Rather than Discretion: The Inconsistency of Optimal Plans," *Journal of Political Economy*, 85, 473-493.

Selten, R. [1975]: "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games," *International Journal of Game Theory*, 4, 25-55.

Figure 1

The savings-taxation model with two citizens.

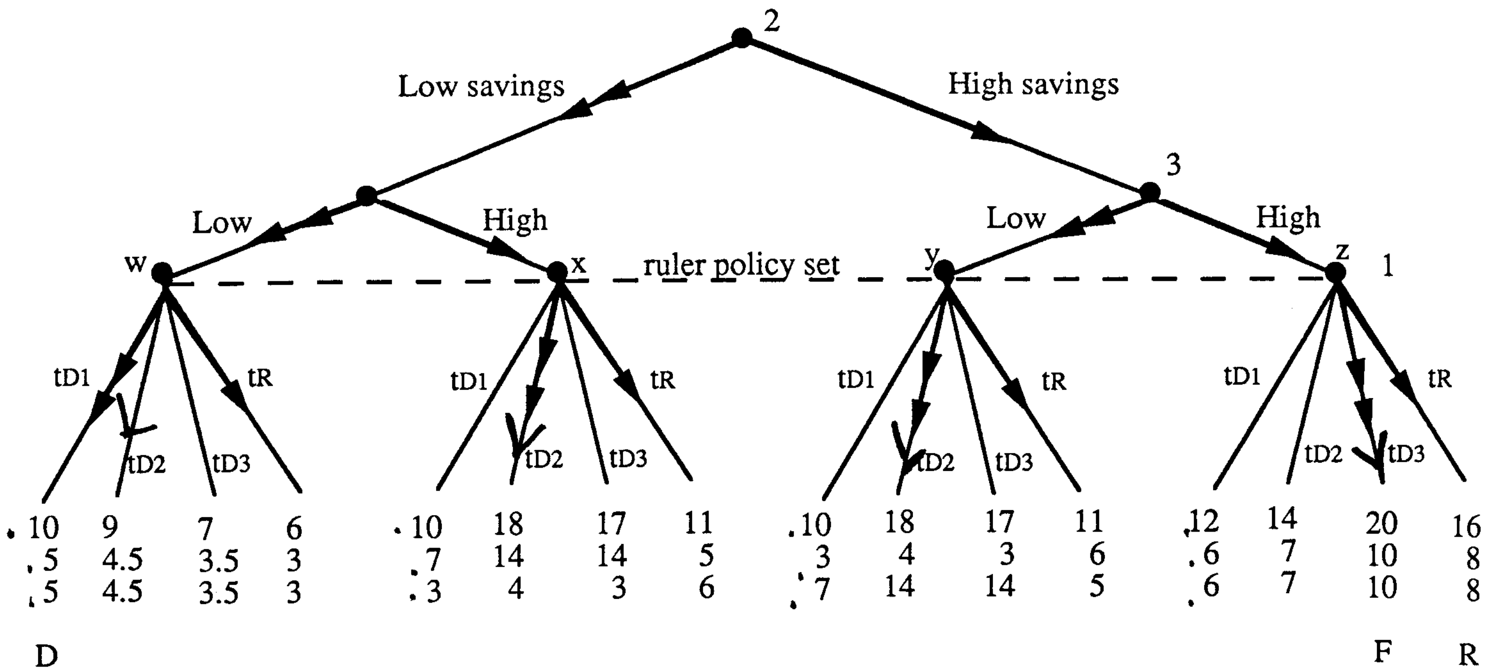

Payoffs are listed in the order: player 1 (**ruler**)
player 2
player 3

Double arrows show strategies under sequential rationality.

Single arrows show strategies under the optimal plan. The optimal plan is time inconsistent.

Notice that if the ruler had perfect policy formation (blot out the dashed line), she could announce tR at w, x, and y and tD3 at z. That plan would be time consistent and arrive at node F.