H. Stutzer

Federal Reserve Bank of Minneapolis
Research Department Working Paper

TIME CONSISTENCY OF OPTIMAL PLANS: A PRIMER

Michael J. Stutzer

Working Paper 223
PACS File 2850

Revised November 1983

Time Consistency of Optimal Plans: A Primer

Introduction

A planning problem consists of the following elements:

1. A designated agent, termed the primordial planner, who chooses a (possibly infinite) planning horizon, the initial time of which is chosen, without losses of generality, to be $t = 0$. The final time of the horizon is denoted T, with $T = \infty$ for an infinite horizon.
2. A set of present and future sequences (or continua) of decisions, termed plans. A subset of these plans are determined by the primordial planner to be feasible, given his perceptions of the problem's constraints.
3. A system linking the feasible plans to outcomes they cause.
4. Performance criteria used by the primordial planner to rank the desirability of feasible plan outcomes. A highest rank feasible plan (if one exists) is termed an optimal, or equivalently, a primordial plan.

As the optimal plan is implemented, other planners may want to change it. They may have fundamental differences with the primordial planner about one or more of the planning elements above. For example, new instruments may become available which change the constraints determining the feasible plans. Unforeseeable new information may become known about the system. Or, society may demand that the planner fundamentally change the performance criteria initially adopted. In any of these events, it

seems appropriate that the "optimal" plan be changed to reflect the new elements.

It is possible, though, that future planners will be in agreement with the primordial planner about the latter three planning elements. Because they evaluate the plans starting at later dates than did the primordial planner, though, their planning horizons are not the same as the primordial planner's. Future planners will only concern themselves with the impacts their proposed plans have on their own present and future, ignoring their own pasts. This myopic (from the viewpoint of the primordial planner) behavior on the part of future planners may lead them to deviate from the primordial plan, even if they all agree with the primordial planner on the other planning elements 2-4. From a normative point of view, if one views the primordial planner's choice of horizon as the appropriate choice, then this myopic behavior is undesirable. Can it somehow be avoided?

If the primordial planner cannot force the future planners to implement the optimal plan, he can only hope that future planners will, of their own free will, choose to implement it. A planning problem which has the property that its optimal plan will be implemented by all future planners, each of whom is free to change it but won't do so, is termed a planning problem with a consistent optimal plan. Even if a planning problem does not have consistent optimal plans, it may be true that some future planner at time $\tau < T$ will freely choose to implement that part of the optimal plan under his control. A planning problem whose optimal plan has this property is termed a problem with a τ-consistent optimal plan.

As we will see, only certain problems with very special structures have consistent or τ-consistent optimal plans. Strotz[1] and Burness[2] have shown that these structures also require that, in any problem which involves discounting the future, the discount function must be of a very special form. In economic planning problems, Kydland and Prescott,[3] Calvo,[4] and Fischer[5] have shown that these special structures require that expectations be adaptive rather than being rational or being, in some other sense, determined by future policy decisions. Finally, we will also show that problems with what are usually termed nonseparable performance criteria do not possess consistent optimal plans.

These results are presented in a unified way through the use of a discrete time optimal control model. The model is broad enough to encompass all of the above special structures, yet simple enough to be understood by people unfamiliar with dynamic optimization techniques. In the process, though, the reader will become acquainted with the so-called Bellman Principle of Optimality, which is a necessary condition characterizing optimal plans in planning problems with certain special structures. We will see that, when discounting is not employed, the problems whose optimal plans are characterized by Bellman's Principle are precisely those problems whose optimal plans are consistent. When discounting is employed, though, this is no longer the case, and a different test for consistency is required.

In problems whose optimal plans are not consistent, the primordial planner must either find some way to force future

planners to implement her plan, or adopt some other strategy. If the former course of action is impossible, the primordial planner may choose to propose some other plan which will be freely implemented by future planners. Such a plan is termed consistent, and in general there may be many consistent, albeit suboptimal, plans. In problems not involving discounting, plans which are consistent are precisely those plans derived via the use of the Bellman Principle, regardless of whether or not it characterizes the optimal plan. As mentioned above, these consistent plans are suboptimal unless the problem has a structure in which Bellman's Principle does characterize the optimal plan.

The Optimal Planning Problem

The most general optimal planning problem considered here can always be placed in the form:

$$(1) \qquad \max_{D_0,\ldots,D_T} U(X_0,D_0,\ldots,X_T,D_T,X_{T+1},t)$$

$$\text{s.t. } X_{t+1} = f_t(X_t,D_0,\ldots,D_t,\ldots,D_T);\ X_0 \text{ given},\ t=0,\ldots,T$$

$$D_t \in C_t(X_t).$$

For each t, the primordial planner chooses a vector D_t from a choice set $C_t(X_t)$ to solve (1), producing an optimal or primordial plan $(D_0^*,\ldots,D_T^*)$. The vector valued state equations f_t determine the evolution of the state vector X_t. Note that we permit the possibility that future decisions may affect current states. This simple setup is broad enough to exhibit all known causes of inconsistency while avoiding complications concerning the existence and

computation of optimal plans inherent in both infinite horizon and stochastic planning problems. An illustrative example follows.

Example 1: Capital Budgeting

Consider the following problem faced by a firm's primordial planner, who chooses a plan to allocate a fixed capital budget K among T competing projects. Project t starts in period t and, for simplicity, is assumed to last one period. A dollar invested in the project starting (and ending) in period t earns a dividend of B_t. An investment of D_t dollars in period t is then assumed to yield dividends of B_tD_t. In addition, retained earnings in period t, denoted X_t, are taxed in that period at the rate P_t.

The undiscounted return to the firm in period t is then:

$$U_t(X_t,D_t) = D_tB_t - P_tX_t, \quad t=0, \ldots, T. \tag{2}$$

At time t = 0, the firm's primordial planner correctly perceives a future interest rate series, $(i_1,\ldots,i_T)$. In deriving her plan, she discounts the return in period t by the factor $1/(1+i_1) \cdot (1+i_2) \cdot \ldots \cdot (1+i_t)$. Defining the discount function

$$r(t,\tau) = 1/(1+i_{\tau+1})\,(1+i_{\tau+2}) \ldots (1+i_t), \qquad r(0,0) = 1 \tag{3}$$

the primordial planner's maximand is then:

$$\sum_{t=0}^{T} r(t,0)U_t(X_t,D_t). \tag{4}$$

To complete the specification of the primordial planner's problem, we derive the state equation for retained earnings X_t. Starting at t = 0, the planner invests D_0 from the total capital fund K, leaving before tax retained earnings of $K - D_0$. After tax retained earnings of $X_1 = (1-P_0)(K-D_0)$ are then left for future investment purposes. In period t = 1, investment of D_1 leaves before tax retained earnings of $X_1 - D_1$, and after tax earnings of $X_2 = (1-P_1)(X_1-D_1)$ are available for future investment purposes. Continuing in this manner, we see that:

$$(5) \qquad X_{t+1} = (1-P_t)(X_t-D_t) \triangleq f_t(X_t,D_t),\ t=0,\ \ldots,\ T;\ X_0 = K.$$

The primordial planner's problem is then:

$$(6) \qquad \max_{D_0,\ldots,D_T} \sum_{t=0}^{T} r(t,0)U_t(X_t,D_t) \triangleq U(X_0,D_0,\ldots,X_T,D_T,X_{T+1},t)$$

$$\text{s.t. } X_{t+1} = f_t(X_t,D_t),\ t=0,\ \ldots,\ T;\ X_0 = K;\ D_t \in C_t(X_t) = [0,X_t]$$

where f_t is given by (5).

Classification of Optimal Planning Problems

There are several subclasses of problem (1) that prove useful in the following. One subclass of problem is termed <u>non-anticipatory</u>, and constitutes those problems in which the current state vector depends only on current and past decision vectors, i.e.,

$$(7) \qquad X_{t+1} = f_t(X_t,D_0,\ldots,D_t),\ t=0,\ \ldots,\ T.$$

By redefining the current state X_t to include the past decision vectors $D_0, \ldots, D_{t-1}$, a nonanticipatory problem's state equations can always be written in the form

$$(8) \qquad X_{t+1} = f_t(X_t, D_t), \quad t=0, \ldots, T.$$

The most general nonanticipatory problem is illustrated in Figure 1. Another important subclass of problems are the monotonically separable problems, in which the objective function U is separable and monotone in the following sense:

$$(9) \qquad U(X_0, D_0, \ldots, X_T, D_T, X_{T+1}) =$$

$$g[U_0(X_0, D_0), U_1(X_1, D_1), \ldots, U_T(X_T, D_T), U_{T+1}(X_{T+1})] =$$

$$g_1[U_0(X_0, D_0), g_2(U_1(X_1, D_1), \ldots, U_T(X_T, D_T), U_{T+1}(X_{T+1}))]$$

and

for any fixed value of U_0, g_1 is a monotonically non-decreasing function of g_2.

A common type of monotonically separable problem is the additive separable problem:

$$(9a) \qquad U(X_0, D_0 \ldots, X_T, D_T, X_{T+1}) = \sum_{t=0}^{T} U_t(X_t, D_t) + U_{T+1}(X_{T+1})$$

diagrammed in Figure 2; and the discounted additive separable problem

$$(9b) \qquad U = \sum_{t=0}^{T} r(t,0) U_t(X_t, D_t) + r(T+1,0) U(X_{T+1})$$

where $r(t,0)$ is a real valued discount function, giving the primordial planner's discount factor for decisions implemented at time t.

Definition of Consistent Plans

Roughly speaking, a consistent plan is a sequence of decisions $(\hat{D}_0,\ldots,\hat{D}_T)$ having the property that, for each time τ, $\hat{D}_\tau$ is the first decision in a plan which maximizes the "objective remaining" after $\hat{D}_0$, ..., $\hat{D}_{\tau-1}$ have been determined. At each τ, the above maximization treats the state evolution $\hat{X}_1$, ..., $\hat{X}_\tau$ parametrically, and calculates the impacts of $\hat{D}_\tau$ on the future states $\hat{X}_{\tau+1}$, ..., $\hat{X}_T$ and utilities U_τ, ..., U_T. More precisely, we adopt two formal definitions within our planning framework; one for undiscounted problems and the other for discounted problems.

For undiscounted problems, a plan $(\hat{D}_0,\ldots,\hat{D}_T)$ is consistent if and only if for each $0 \leq \tau \leq T$, $\hat{D}_\tau$ is the first component in a solution of the following problem:

$$(10) \qquad \max_{D_\tau,\ldots,D_T} U(X_0,\hat{D}_0,\hat{X}_1,\hat{D}_1,\ldots,\hat{X}_{\tau-1},\hat{D}_{\tau-1},\hat{X}_{\tau-1},\hat{X}_\tau,D_\tau,X_{\tau+1},$$

$$D_{\tau+1},\ldots,X_T,D_T,X_{T+1})$$

$$\text{s.t.} \quad X_{t+1} = f_t(\hat{X}_t,\hat{D}_0,\ldots,\hat{D}_{t-1},D_t,\ldots,D_T);\ t=\tau,\ \ldots,\ T$$

with X_0 given and $(\hat{X}_1,\ldots,\hat{X}_T)$ is the state evolution resulting from $(\hat{D}_0,\ldots,\hat{D}_T)$ and the state equations f_t.

This property can be weakened by requiring only that $\hat{D}_\tau$ be the first component of a solution vector to (10) for some particular $\tau > 0$. This weaker property is termed τ-consistency. The plan is consistent if and only if it is τ-consistent for all $\tau = 1, \ldots, T$.

For discounted problems with the performance index (9b), the formulation (10) will not do. This is because (10) assumes that the maximand of a planner at time t is

$$(11) \qquad \sum_{t=t}^{T} r(t,0)U_t(\hat{X}_t, D_t) + r(T+1,0)U_{T+1}(X_{T+1}),$$

which implies that the planner at t discounts his own current and future utilities as heavily as did the primordial planner t periods earlier. A more realistic assumption treats any future planners and the primordial planner symmetrically by assuming that the planner at t discounts future utilities $U_{t+1}, \ldots, U_T$ in the same way that the primordial planner discounts her own future utilities $U_1, \ldots, U_{T-t}$. For example, if it is assumed that the primordial planner does not discount her current utility U_0, then we also must assume that the planner at t does not discount his own current utility U_t. Denoting the discount factor for U_t of a planner at time τ by $r(t,\tau)$, we say a plan $(\hat{D}_0, \ldots, \hat{D}_T)$ is discounted consistent if and only if for each τ, $0 < \tau \leq T$, $\hat{D}_\tau$ is the first component in a solution vector to the following problem:

$$(12) \qquad \max_{D_\tau,\ldots,D_T} \sum_{t=\tau}^{T} r(t,\tau)U_t(X_t,D_t) + r(T+1,\tau)U_{T+1}(X_{T+1})$$

$$\text{s.t. } X_{t+1} = f_t(\hat{X}_t,\hat{D}_0,\ldots,\hat{D}_{t-1},D_t,\ldots,D_T),\ t=\tau,\ \ldots,\ T.$$

If, for some particular $\tau > 0$, $\hat{D}_\tau$ is the first component in a solution vector to (9b), we say that the plan is τ-discounted consistent.

Of course, the optimal plan $(D_0^*,\ldots,D_T^*)$ in either an undiscounted or a discounted problem may or may not be consistent or discounted consistent. In the following sections, we explore the relationship between consistency and optimality.

The Relationship Between Consistency and Optimality for Undiscounted Problems

Bellman's Principle of Optimality

Richard Bellman formulated the basis for his optimization technique, called Dynamic Programming, in his 1957 book with the same title.(6) The basis for the technique, called the Principle of Optimality, is a necessary condition characterizing optimal plans for certain types of planning problems. It is reprinted below:

> "An optimal policy has the property that whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision."

In our jargon, the applicability of Bellman's Principle implies that the optimal plan $(D_0^*,\ldots,D_T^*)$ in an undiscounted problem is 1-

consistent. This is because it says that $(D_1^*,\ldots,D_T^*)$ solves (10) for $\tau = 1$, so its first component D_1^* is the first decision in the optimal plan. If Bellman's Principle applies at each $\tau > 0$, considering the initial state to be X_τ^* and the initial decision to be $D_{\tau-1}^*$, then the optimal plan is consistent.

As shown by Mitten,[7] and restated in Nemhauser,[8] Bellman's Principle applies at each τ for nonanticipatory, monotonically separable problems, which proves that they have consistent optimal plans. The application of Bellman's Principle in the important special case of nonanticipatory, additive separability is demonstrated below. The problem is:

$$\max_{D_0,\ldots,D_T} U = U_0(X_0,D_0) + \ldots + U_t(X_t,D_t) + \ldots + U_T(X_T,D_T) + U_{T+1}(X_{T+1}) \tag{13}$$

s.t. $X_{t+1} = f_t(X_t,D_t)$, X_0 given, $t=0, \ldots, T$, $D_t \in C_t(X_t)$, and is diagrammed in Figure 2.

To apply Bellman's Principle, we first note that the nonanticipatory nature of the state equation allows us to recursively compute X_{t+1} as follows:

$$\begin{aligned} X_1 &= f_0(X_0,D_0) \\ X_2 &= f_1(X_1,D_1) = f_1(f_0(X_0,D_0),D_1) \\ X_3 &= f_2(X_2,D_2) = f_2(f_1(f_0(X_0,D_0),D_1),D_2) \\ &\vdots \\ X_{t+1} &= f_t(X_t,D_t) = \\ &\quad f_t(f_{t-1}(f_{t-2}(\ldots(f_0(X_0,D_0),D_1),D_2),\ldots),D_t). \end{aligned} \tag{14}$$

Thus, any future state depends solely on decisions made prior to that time and the initial state X_0, which is clear from a glance at Figure 2. By substituting the above relations into the additively separable objective function U, we see that the problem becomes:

$$(15) \qquad \max_{D_0,\ldots,D_T} U_0(X_0,D_0) + U_1(f_0(X_0,D_0),D_1) + U_2(f_1(f_0(X_0,D_0),D_1),D_2) + \ldots + U_T(X_T,D_T) + U_{T+1}(X_{T+1})$$

with X_0 given and D_t chosen from the set $C_t(X_t)$, t=0, ..., T.

It is the special structure of (13) that permits the application of Bellman's Principle. Noting that D_1 doesn't occur in U_0, (13) can be solved by:

$$(16) \qquad = \max_{D_0}[U_0(X_0,D_0) + \max_{D_1,\ldots,D_T} U_1(X_1^*,D_1) + \ldots + U_T(X_T,D_T) + U_{T+1}(X_{T+1})],$$

$$\text{s.t. } D_0 \in C_0(X_0) \qquad D_t, \in C_t(X_t).$$

$$X_0 \text{ given} \qquad X_1^* = f_0(X_0,D_0)$$

$$X_{t+1} = f_t(X_t,D_t) \quad t = 1, \ldots, T$$

This is precisely what Bellman's Principle says, i.e., that the remaining decisions $(D_1^*,\ldots,D_T^*)$ must constitute an optimal plan with regard to the state $(X_1^*=f_0(X_0,D_0^*)$ resulting from the first decision D_0^*. Also, this ability to decompose the optimization problem shows that the optimal plan solving (10) is 1-consistent. A planner at time 1 who reconsiders the optimal plan

$D_1^*, \ldots, D_T^*$ computed by the primordial planner would solve (10) for $\tau = 1$, taking $X_1^* = f_0(X_0, D_0^*)$ as given. Pictorially, the planner at time 1 ignores the part of Figure 2 to the left of X_1 in calculating his plan. This is the same problem (so produces the same solution) as the inner maximization in (16) computed by the primordial planner. In fact, at each $\tau > 0$, Bellman's Principle still applies to (10), whose solution is thus τ-consistent for all τ, i.e., consistent. To see this, note from (16) that its inner maximization can be similarly decomposed into two parts:

$$
\begin{aligned}
&\max_{D_1,\ldots,D_T} \left[U_1(X_1,D_1) + \sum_{t=2}^{T} U_t(X_t,D_t) + U_{T+1}(X_{T+1})\right] \\
&= \max_{D_1}\{U_1(X_1,D_1) + \max_{D_2,\ldots,D_T} U_2(X_2^*,D_2) + \sum_{t=3}^{T} (U_t(X_t,D_t) + \\
&U_{T+1}(X_{T+1}))\}, \quad D_t \in C_t(X_t) \\
&\text{s.t. } D_1 \in C_1(X_1) \qquad D_t \in C_t(X_t) \\
&\qquad X_1 \text{ given} \qquad\quad X_2^* = f_1(X_1,D_1) \\
&\qquad\qquad\qquad\quad X_{t+1} = f_t(X_t,D_t) \quad t = 2, \ldots, T
\end{aligned}
\tag{17}
$$

because U_1 does not depend on the future decisions $D_2, \ldots, D_t$. Thus, Bellman's Principle applies again, and the optimal plan is also 2-consistent. Continuing in this manner, the original T+1 variable optimization problem (13) can be decomposed into T+1 single variable optimization problems, yielding the dynamic programming decomposition

$$
\begin{aligned}
&\max_{D_0,\ldots,D_T} U = \max_{D_0}\{U_0(X_0,D_0) + \max_{D_1}[U_1(X_1,D_1) + \ldots + \\
&\max_{D_T}\{U_T(X_T,D_T) + U_{T+1}(X_{T+1}))\}]\ldots]
\end{aligned}
\tag{18}
$$

where $X_{t+1} = f_t(X_t,D_t)$, $t=0, \ldots, T$ and $D_t \in C_t(X_t)$. So, Bellman's Principle applies at each $\tau > 0$, and the optimal plan in the undiscounted, additive separable, nonanticipatory problem is consistent. A planner at any time $\tau > 0$ thus ignores the part of Figure 2 to the left of X_τ in computing his plan. As was mentioned earlier, Bellman's Principle will also apply at each τ in the more general class of undiscounted monotonically separable, nonanticipatory problems maximizing (9) subject to (8). Optimal plans for these problems are then also consistent. To summarize, Bellman's Principle applies to undiscounted, monotonically separable, nonanticipatory problems. In any undiscounted problem, the applicability of Bellman's Principle at each τ implies that the optimal plan is consistent.

Discounted Consistency for Discounted Problems

The discounted, additive separable, nonanticipatory problem of maximizing (9b) subject to (8) can also be solved by applying Bellman's Principle at each τ. This is easily seen by noting that redefining U_t to be $r(t,0)U_t(X_t,D_t)$ brings (9b) into the form of (13). Then, one applies the proof just given.

This fact, though, has no bearing on whether or not the optimal plan is discounted consistent. While the primordial plan satisfies (10) for each τ, discounted consistency requires the satisfaction of (12) for each τ. In a continuous time setting, Burness, generalizing a result of Strotz, found necessary and sufficient conditions that must be satisfied by $r(t,\tau)$ to ensure discounted consistency of the optimal plan. His proof involves the use of the calculus of variations. Our discrete time formu-

lation allows us to derive the discrete time analogue of his result using simple calculus.

Discrete Time Burness Theorem

A necessary and sufficient condition for the discounted consistency of an optimal plan maximizing (9b) subject to (8) is that

$$(19) \qquad \frac{r(t,\tau)}{r(s,\tau)} = \frac{r(t,\tau')}{r(s,\tau')}, \text{ for all } s,\ t,\ \tau,\ \tau'.$$

A simple corollary of this is that if the discount function is given by (3), then the optimal plan is discounted consistent.

Proof. We first prove necessity. The final decision D_T^* of the optimal plan is found by solving

$$(20) \qquad \max_{D_T} r(T,0)U_T(X_T^*,D_T) + r(T+1,0)U_{T+1}(f_T(X_T^*,D_T))$$

where X_T^* is the state at time T of the primordial plan. Assuming an interior maximum, solving (20) yields:

$$r(T,0)\frac{\partial U_T}{\partial D_T}(X_T^*,D_T^*) + r(T+1,0)\frac{\partial U_{T+1}}{\partial f_T}\frac{\partial f_T}{\partial D_T}(X_T^*,D_T^*) = 0$$

or

$$(21) \qquad \frac{r(T+1,0)}{r(T,0)} = -\frac{\frac{\partial U_T}{\partial D_T}(X_T^*,D_T^*)}{\frac{\partial U_{T+1}}{\partial f_T}\frac{\partial f_T}{\partial D_T}(X_T^*,D_T^*)}$$

A planner at time τ would solve (12), with his final decision D_T, computed by solving:

$$(22) \qquad \max_{D_T} r(T,\tau)U_T(\hat{X}_T,D_T) + r(T+1,\tau)U_{T+1}(f_T(\hat{X}_T,D_T))$$

where $\hat{X}_T$ is the state at time T and is taken parametrically. Solving (22) yields:

$$(23) \qquad \frac{r(T+1,\tau)}{r(T,\tau)} = -\frac{\frac{\partial U_T}{\partial D_T}(\hat{X}_T,\hat{D}_T)}{\frac{\partial U_{T+1}}{\partial f_T}\frac{\partial f_T}{\partial D_T}(\hat{X}_T,\hat{D}_T)}.$$

Consistency of the optimal plan requires that the sequence of future planners at each time τ chooses the same decision that the primordial planner did. This means that $\hat{D}_0 = D_0^*$, $\hat{D}_1 = D_1^*$, ..., $\hat{D}_T = D_T^*$, which, through the state equations, implies that $\hat{X}_1 = X_1^*$, $\hat{X}_2 = X_2^*$, ..., $\hat{X}_{T+1} = X_{T+1}^*$. In particular, consistency of the primordial plan requires $\hat{X}_T = X_T^*$ $\hat{D}_T = D_T^*$, so that the right-hand sides of (21) and (23) are equal. Thus, we see that

$$(24) \qquad \frac{r(T+1,0)}{r(T,0)} = \frac{r(T+1,\tau)}{r(T,\tau)}, = 1, \ldots, T,$$

i.e., that the ratio $\frac{r(T+1,\tau)}{r(T,\tau)}$ is independent of τ.

Next, consider the primordial planner's choice D_{T-1}^*. This is formed by solving:

$$(25) \qquad \max_{D_{T-1}} \Big[r(T-1,0)U_{T-1}(X_{T-1}^*,D_{T-1}) + \max_{D_T}\big[r(T,0),U_T(f_{T-1}(X_{T-1}^*,D_{T-1}),D_T) + r(T+1,0)U_{T+1}(f_T(f_{T-1}(X_{T-1}^*,D_{T-1}),D_T))\big]\Big]$$

whose solution satisfies the condition that the partial derivative of (25) with respect to D_{T-1} equals zero, or

$$(26)\qquad r(T-1,0)\,\frac{\partial U_{T-1}}{\partial D_{T-1}}\,(X^*_{T-1},D^*_{T-1}) + r(T,0)\,\frac{\partial U_T}{\partial f_{T-1}}\,\frac{\partial f_{T-1}}{\partial D_{T-1}}$$

$$(X^*_{T-1},D^*_{T-1})$$

$$+\ r(T+1,0)\,\frac{\partial U_{T+1}}{\partial f_T}\,\frac{\partial f_T}{\partial f_{T-1}}\,\frac{\partial f_{T-1}}{\partial D_{T-1}}\,(X^*_{T-1},D^*_{T-1}) = 0.$$

Dividing by $r(T,0)$ and rearranging leaves:

$$(27)\qquad \frac{r(T-1,0)}{r(T,0)}\,\frac{\partial U_{T-1}}{\partial D_{T-1}} = -\,\frac{\partial U_T}{\partial f_{T-1}}\,\frac{\partial f_{T-1}}{\partial D_{T-1}}$$

$$+\,\frac{r(T+1,0)}{r(T,0)}\,\frac{\partial U_{T+1}}{\partial f_T}\,\frac{\partial f_T}{\partial f_{T-1}}\,\frac{\partial f_{T-1}}{\partial D_{T-1}}$$

where all derivatives are evaluated along the optimal plan.

The planner at any time τ who reconsiders this choice of D_{T-1} would solve:

$$(28)\qquad \max_{D_{T-1}}\ [r(T-1,\tau)U_{T-1}(X^*_{T-1},D_{T-1}) +$$

$$\max_{D_T}[r(T,\tau)U_T(f_{T-1}(X^*_{T-1},D_{T-1}),D_T)$$

$$+\ r(T+1,\tau)U_{T-1}(f_T(f_{T-1}(X^*_{T-1},D_{T-1}),D_T))]].$$

Consistency requires its solution to be the same as that of (25), so a modified condition (27) results, with 0 replaced by τ:

$$(29)\qquad \frac{r(T-1,\tau)}{r(T,\tau)}\,\frac{\partial U_{T-1}}{\partial D_{T-1}} = -\,\frac{\partial U_T}{\partial f_{T-1}}\,\frac{\partial f_{T-1}}{\partial D_{T-1}} +$$

$$\frac{r(T+1,\tau)}{r(T,\tau)}\,\frac{\partial U_{T+1}}{\partial f_T}\,\frac{\partial f_T}{\partial f_{T-1}}\,\frac{\partial f_{T-1}}{\partial D_{T-1}}\ .$$

As we have just seen, $\frac{r(T+1,\tau)}{r(T,\tau)}$ is independent of τ. Therefore, (29) implies that $\frac{r(T-1,\tau)}{r(T,\tau)}$ must also be independent of τ, including, by (27), $\tau = 0$. But then, of course, $\frac{r(T+1,\tau)}{r(T-1,\tau)}$ must also be independent of τ.

Considering the primordial and future planners' problems for finding D_{T-2}, and imposing consistency, we similarly find that $\frac{r(T-1,\tau)}{r(T-2,\tau)}$ is independent of τ. This implies that the ratios $\frac{r(T,\tau)}{r(T-2,\tau)}$ and $\frac{r(T+1,\tau)}{r(T-2,\tau)}$ are also independent of τ. Continuing the process for D_{T-3}, ..., D_0, it is tedious, but simple, to find that it is necessary that a discounted consistent, optimal plan in a nonanticipatory, additive separable problem satisfies (19).

The proof of the sufficiency of (19) for the discounted consistency of the optimal plan is straightforward, and will be omitted here.

The Relevance of Strotz-Burness

The problem both Strotz and Burness use to discuss this result is a problem where the primordial planner, who also plays the role of all future planners, is planning her own future consumption stream. If the planner's preference ordering over future consumption streams is not representable by a discounted separable utility (9b) with $r(t,\tau)$ of the necessary form (19), then the primordial plan is not consistent. Strotz views the lack of consistency as damaging to a normative theory of behavior, which prescribes that the plan which should be implemented is the primordial plan. From this viewpoint, the consumer who doesn't implement her primordial plan is myopic. Strotz says (1, p. 173)

> "An individual, who because he does not discount all future pleasures at a constant rate of interest finds himself continuously repudiating his past plans, may learn to distrust his future behavior, and may do something about it."

Strotz proposes that such an individual may choose either of two courses:

1. At time zero, the individual chooses to precommit her future decisions irrevocably, or
2. the individual will decide to successively recalculate the optimal plan.

The latter would produce a plan in which, for each τ, the chosen $\hat{D}_\tau$ is that found by solving (12) for τ. By definition, this procedure always generates a discounted consistent plan, which is not the primordial plan unless $r(t,\tau)$ is of the necessary form (20).

Strotz's alarm is unjustified in problems like those he discussed, where there is only a single decision maker, i.e., the primordial planner. Unless the primordial planner has changed her mind about the appropriate horizon, or about the appropriate performance criterion, state equations and constraints operative over that horizon, she never would reconsider the primordial plan. Its consistency (or lack of) is then a moot point. If the primordial planner did change her mind about one or more of the above planning problem elements, then she would not necessarily consider the primordial plan to be optimal anymore. The fact that it is also inconsistent in this event does not imply that the

primordial planner is myopic. Rather, the economist who posits that such a decision maker will implement the primordial plan has made an incorrect assumption.

Many economic modellers of a single decision maker's behavior have assumed that the primordial plan will be implemented. These modellers assume that the single decision maker (e.g., a firm or a consumer) maximizes some objective (e.g., profits or utility) over some fixed horizon, usually chosen to be infinite. Examples include the dynamic competitive firm models treated by Sargent.(9) Among theories of the firm, the most widely accepted behavioral assumption is that firms maximize discounted profits, using the discount function (3) of our capital budgeting example. Among theories of the consumer, there is less agreement on the admissible forms of the consumer's intertemporal utility function. However, the assumption commonly employed in the recent rational expectations models of Sargent and others is that a consumer's intertemporal preferences are representable by a discounted, additively separable, time invariant utility function:

$$(30) \qquad \sum_{t=0}^{\infty} \left(\frac{1}{1+i}\right)^t U(C_t) \triangleq \sum_{t=0}^{\infty} r^t U(C_t)$$

where $r = \frac{1}{1+i}$ is a "constant rate of time preference."

The corollary to the Burness Theorem implies that nonanticipatory models of the firm with a discount function (3) will be consistent. Thus, even if the firm's management planner changes its decision horizon and reconsiders the primordial plan, the primordial plan will still be implemented. Nonanticipatory

models of the consumer with a discounted, additively separable utility of the form (30) will also be consistent, as noted by Strotz. This is because the constant rate discount function used in them is:

$$(31) \qquad r(t,\tau) = r^{t-\tau} = \left(\frac{1}{(1+i)}\right)^{t-\tau} = \frac{1}{\underbrace{(1+i)\ (1+i)\ \ldots\ (1+i)}_{t-\tau \text{ times}}}$$

and is thus of the same form as (3), with $i_t = i$, for all t. Of course, even if a consumer's utility function is additively separable, it is possible that the consumer's discount function will not be of the form (19). In this event, an economist who admits the possibility that the consumer will change his horizon and reconsider his primordial plan, must then also admit that the primordial plan is inconsistent and will not be implemented.

Strotz's claim that inconsistency of primordial plans leads to undesirable myopic behavior has more validity in models with future planners other than the primordial planner. Their decision horizons differ from that of the primordial planner. In these models, even if all future planners agree with the primordial planner's choice of performance criterion, and with her assessment of the state equations and constraints, they still may not implement the primordial plan unless $r(t,\tau)$ satisfies (19). If one believes that the primordial planner's choice of horizon is, in some sense, the correct choice, then the inconsistency certainly seems to cause undesirable myopic behavior.

Optimal growth models which maximize social welfare over long or infinite horizons must admit the existence of future

planners. Inconsistency of their optimal plans is likely, for they are often proposed for use in planning environments where it is impossible to precommit the behavior of future planners, and because there is seldom any good reason to believe that discounting will be of the form (19). Several theorists assume that the social welfare function is of the form (30), in which case the optimal plan is consistent. But the price that these models pay is the same paid by models of individual agents which assume this form. These theorists must rule out the possibility that the planner's preference orderings may not be representable by social welfare functions of that form. Koopmans[10] has found axioms characterizing preference ordering which are representable by (30). Whether or not these axioms are plausible in positive models, or desirable in normative models, has never been determined.

Anticipatory Behavior as a Cause of Inconsistency

Anticipatory behavior as a pervasive cause of inconsistency has been stressed by Prescott and Kydland, Fisher, and the introductory article in Lucas and Sargent.[11] The concept can be fully illustrated by the simple, additive separable, anticipatory problem illustrated in Figure 3. The planner's problem is to:

$$(32) \qquad \max_{D_0,D_1} U_0(X_0,D_0) + U_1(X_1,D_1) + U_2(X_2)$$

$$\text{s.t. } X_1 = f_0(X_0,D_0,D_1),\ X_2 = f_1(X_1,D_1).$$

Substituting the state equations in the objective function, problem (32) is to find D_0^* and D_1^* solving

$$(33) \qquad \max_{D_0,D_1} U_0(X_0,D_0) + U_1(f_0(X_0,D_0,D_1),D_1) + U_2(f_1(X_1,D_1)).$$

The solution of (32) will, in general, be inconsistent. The problem is undiscounted, so consistency requires that (33) be decomposed by the Bellman Principle:

$$(34) \qquad \max_{D_0} [U_0(X_0,D_0)+\max_{D_1}[U_1(X_1^*,D_1)+U_2(f_1(X_1^*,D_1))]]$$

where (X_1^*,X_2^*) is the state trajectory determined by f_0, f_1, and the optimal plan D_0^*, D_1^*.

However, (34) does not follow from (33), so (33) is not 1-consistent. This is because a future planner's choice of D_1 at time period 1, taking the state X_1^* as given, will solve the inner maximization in (34), i.e.,

$$(35) \qquad \max_{D_1} [U_1(X_1^*,D_1)+U_2(f_1(X_1^*,D_1))],$$

producing a decision $\hat{D}_1$ which, unlike the optimal plan's D_1^*, takes no account of the indirect impact D_1 has on U_1 through $X_1 = f_0(X_0,D_0,D_1)$. Thus $\hat{D}_1 \neq D_1^*$, so the optimal plan is inconsistent. This is obvious from Figure 3, where we see that a future planner at time 1 who accepts X_1^* as given, will, of course, ignore the effect D_1 has on X_1. The future planner worries only about the effects D_1 has on U_1, X_2, and D_2.

Prescott and Kydland argue that the assumption of perfect foresight or rational expectations in macroeconomic planning models introduces such anticipatory phenomena into them, which cause the primordial plans to be inconsistent. In our paradigm, a macroeconomic planning model has the government choosing policy variables D_0, ..., D_T to maximize some performance cri-

terion U subject to the behavioral rules of firms and consumers, as given in the form of state equations. These rules take the general form:

$$(36) \qquad X_{t+1} = f_t\,(X_t, D_0, \ldots, D_t, D^e_{t+1}, \ldots, D^e_T),\ t=0,\ \ldots,\ T$$

which means that the relevant behavior of firms and consumers in the next period is some function of past behavior and government decisions as well as their expectation of future government decisions $D^e_{t+1}, \ldots, D^e_T$. The assumption of perfect foresight in the deterministic model above is that:

$$(37) \qquad D^e_{t+j} = D_{t+j},\ j=1,\ \ldots,\ T-t;\ \text{for each } t=0,\ \ldots,\ T-1$$

which creates the most general possible planning problem (1) considered here. Unlike the previous example, the state X_t in each period t depends on future decisions. Then, not only is the optimal plan inconsistent, it is not even τ-consistent for any τ.

Furthermore, this negative result still holds when the assumption of perfect foresight is weakened. For example, consider the myopic perfect foresight case illustrated in Figure 4. There, the behavioral state at time t depends not on all future decisions, but only on that decision made in period t+1:

$$(38) \qquad X_{t+1} = f_t(X_t, D_t, D^e_{t+1}) = f_t(X_t, D_t, D_{t+1}).$$

The same thing that happens in Figure 3 happens in each period, so that an optimal plan for an economy with myopic perfect foresight is neither consistent nor τ-consistent for any $\tau > 0$.

In fact, any systematic relationship between expected and actual future decisions will cause inconsistency. In more sophisticated stochastic models where future expectations are random variables, the rational expectations hypothesis that D^e_{t+j} equals the mean of D_{t+j} will cause inconsistency for the same reason. In fact, even systematically biased, but still future determined expectations $D^e_{t+j} = g$ (mean of D_{t+j}), where g is some function, should cause inconsistency.

Prescott and Kydland give some microeconomic examples of anticipatory phenomena, which fit into the simple two-period model of Figure 3. For example, they mention the case of optimal government patent policies. There, it is optimal to offer a number of patents (D^*_0) initially, to induce an optimal level of inventive activity (X^*_1). But in period 1, <u>after these inventions are created</u>, it seems best to remove all the patents, leaving no ($\hat{D}_1 = 0$) patents, rather than leaving $D^*_1 = D^*_0$ patents, so that the inventions can be produced by competitive markets rather than by monopolies. Inventors who expected this to happen (i.e., $D^e_1 = \hat{D}_1 = 0$) would not have undertaken the optimal level of inventive activity, (X^*_1), though. So, the optimal level of patents that should remain, $D^*_1 \neq \hat{D}_1 = 0$, and the optimal plan is inconsistent.

Both Prescott and Kydland and Fischer discuss the inconsistency of optimal taxation plans in worlds where firms have future determined expectations. However, they offer radically different remedies for this problem. The former advocate Strotz's first course, i.e., that future decisions be precommitted by the primordial planner. They propose to do this by having the pri-

mordial planner substitute simple policy rules, wherein future planners are bound to use predetermined rules to make decisions only nominally under their control, for the discretion future planners would otherwise have in making decisions. Although these rules are clearly suboptimal from the primordial planner's point of view, Prescott and Kydland's simulations show that these rules appear to outperform the suboptimal solution which would otherwise be implemented at the discretion of future planners. However, as Fischer pointed out, there will generally still be an incentive for future planners to break these suboptimal (from their viewpoint) rules. Of course, one might argue that the primordial planner could adopt constitutional or other binding constraints to force the future planners to follow these policy rules. But if that were feasible, why couldn't the primordial planner adopt constraints to prevent the optimal plan itself from ever being changed?

Rather than advocate the substitution of simple rules for discretion, Fischer contends that the use of discretion on occasion may be warranted. The patent example cited earlier can be used to illustrate Fischer's reasoning. If the future planners repeatedly revoke patents granted earlier to investors, then eventually all future inventors would expect this to happen and would cease inventing. Repeatedly revoking patents in a rational expectations environment is thus undesirable. However, the occasional revocation of patents which would have produced vast monopoly profits will probably not reduce total inventive activity very much, and will produce much social benefit through lowered

monopoly profits. Fischer raises the possibility that randomization of decisions may be of use in devising such a strategy. In the patent example, the government might randomly cancel patents, using some preannounced and expected probability distribution to do so. The resulting outcome of this stochastic plan may be better than that of a simple, deterministic rule, particularly if the inventors most willing to bear the risk are also the most likely to produce valuable inventions. Once again, though, why wouldn't future planners have an incentive to change a previously adopted randomization rule?

All of the problems created by time consistency due to anticipatory phenomena are the result of a fundamental paradox. This paradox is due to the logical inconsistency of the following three assumptions common to analysts' models of optimal planning with anticipatory phenomena:

(A1) Assume that agents correctly anticipate future planning actions (i.e., decisions, policy rules, randomization rules, etc.) which are relevant to their own welfare.

(A2) Assume that future planners can observe the agents' past anticipations of current and future planning actions.

(A3) Assume that future planners have the discretion to implement actions which were not anticipated in the past by agents.

Time inconsistency implies that planners endowed with the powers of (A2) and (A3) will actually implement actions which were not anticipated by past agents, thus violating the "rational expectations" assumption (A1). Alternatively, if an analyst insists that (A1) and (A2) must be valid, then (A3) must be violated, i.e., the analyst cannot admit the possibility that future planners have the discretion to change the primordial plan. In this case, future planners have no "free will" to act if agents can correctly anticipate these actions beforehand. They are preordained to follow the prior anticipations of agents which, due to time inconsistency, would not be realized if planners did have "free will." Because of this, it is difficult to understand the logic of estimating the effects of "policy regime changes" in models (see, e.g., (11)) that purport to have rational expectations operating both before and after regime changes.

Nonseparability As a Cause of Inconsistency

Even in problems without anticipatory elements, a nonseparable performance criterion can lead to inconsistency of the optimal plan. To see this, consider the simple, undiscounted two-period problem below:

$$\max_{D_0,D_1} U(X_0,D_0,X_1,D_1)$$

(39)

$$\text{s.t. } X_1 = f_0(X_0,D_0),\ X_0 \text{ given } D_t \in C_t(X_t)$$

Consistency requires that the Bellman Principle applies to the optimal solution, i.e, that for any fixed X_0:

$$(40) \qquad \max_{D_0,D_1} U(X_0,D_0,X_1,D_1) = \max_{D_0} \max_{D_1} U(X_0,D_0,f_0(X_0,D_0),D_1).$$

Unfortunately, it is well known that the joint maximization of a function of two variables cannot always be decomposed into the two one-variable problems that consistency requires it to be. As mentioned earlier, Mitten[7] has found simple sufficiency conditions (9) on the form of U which guarantee that Bellman's Principle applies. The additive separable case (5) satisfies these conditions, as does the multiplicative form

$$(41) \qquad U(X_0,D_0,X_1,D_1,\ldots,X_T,D_T) = \prod_{t=0}^{T} U_t(X_t,D_t)$$

when the U_t's are constructed so that $U_t > 0$ for all $D_t \in C_t(X_t)$. As was argued earlier, it is not at all clear that optimal growth problems will, or should, have additively separable utility functions. If the utility function is not even monotonically separable, then the primordial plan will, in general, be inconsistent, even though discounting is not employed.

Conclusions

We have seen that any of the following elements will lead to inconsistency of primordial plans:

1. Discounting with a discount function not satisfying (19).
2. Anticipatory elements, such as expectations which are, in some sense, determined by future policy decisions.
3. Nonseparable performance criteria.

We have argued that discounting may not be of the appropriate form in optimal growth problems. Anticipatory elements seem likely in any economic problem where controlled agents have both an incentive to forecast future decisions, and the skills required to forecast in some way which is systematically related to the actual future decisions. Nonseparable performance criteria may occur in problems involving intertemporal preferences over future consumption streams, unless one is confident that the restrictive axioms needed to justify additive separable utilities characterize actual or desirable behavior.

Figure 1 Nonanticipatory Problem

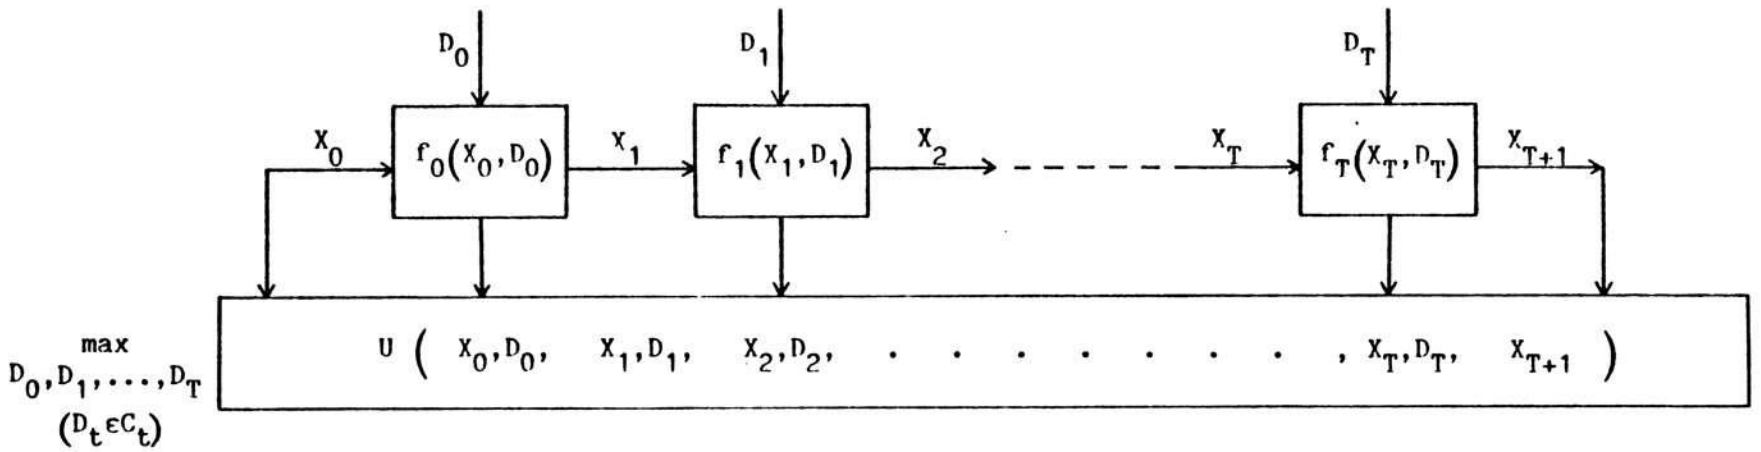

Figure 2 Additive Separable Nonanticipatory Problem

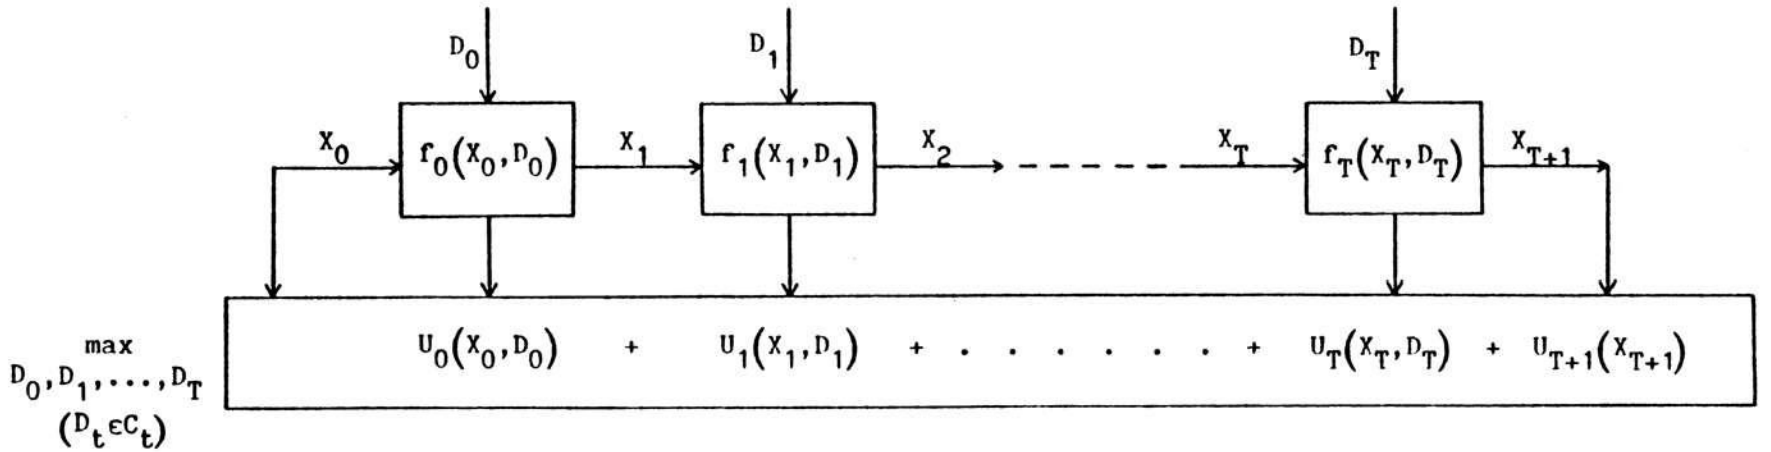

Figure 3 Additive Separable Anticipatory Problem

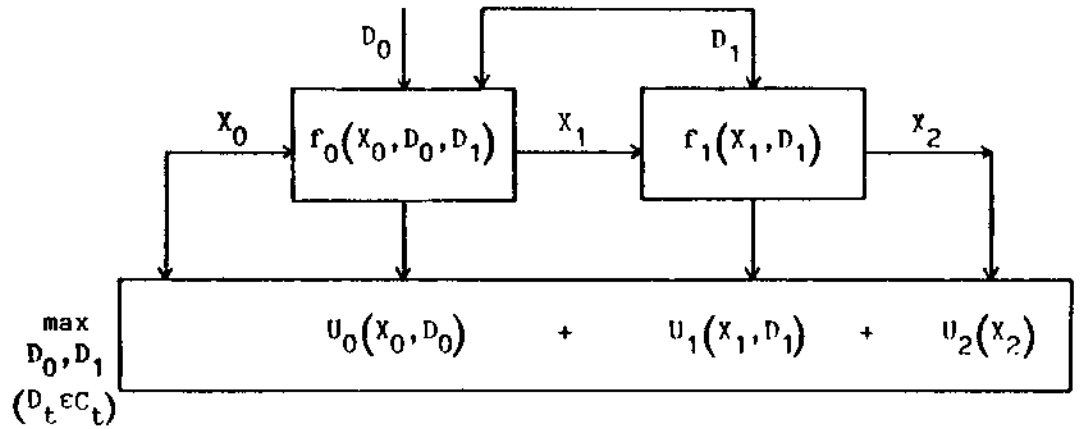

Figure 4 Myopic Perfect Foresight: $\left(D_{t+1}\right)^e = D_{t+1}$

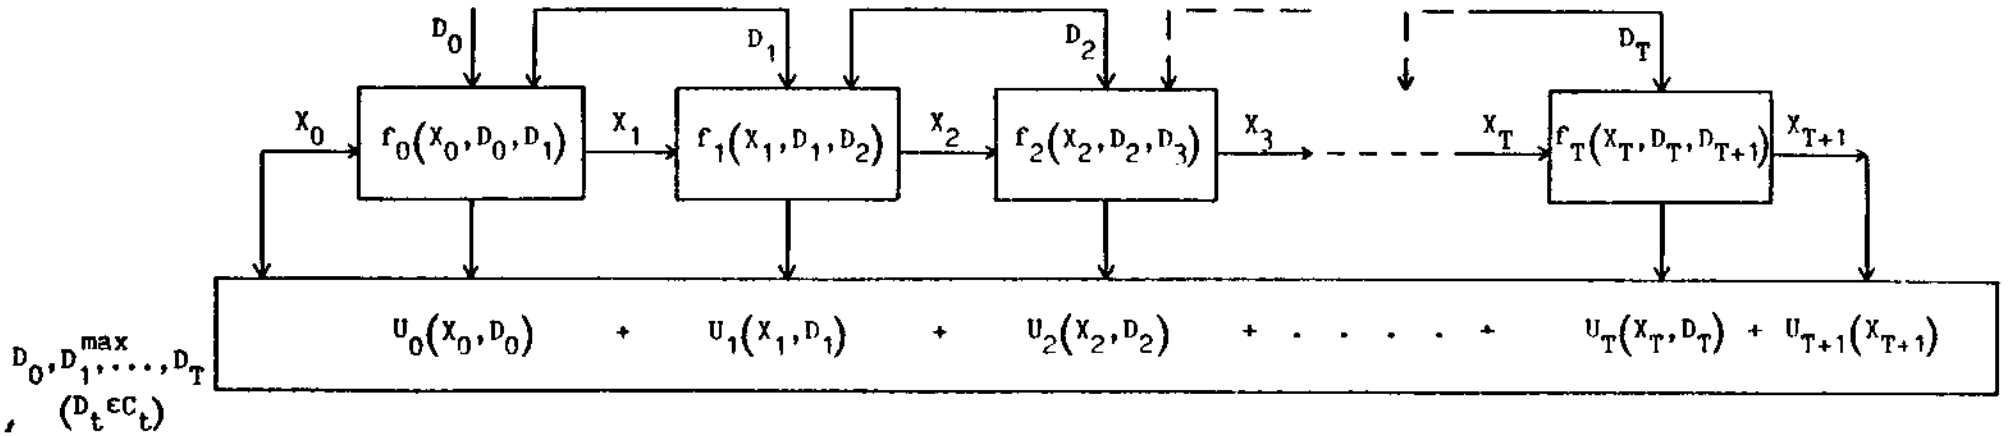

References

(1) R. Strotz, "Myopia and Inconsistency in Dynamic Utility Maximization," Review of Economic Studies 23, (1956), pp. 165-180.

(2) H. S. Burness, "A Note on Consistent Naive Intertemporal Decision Making and an Application to the Case of Uncertain Lifetime," Review of Economic Studies 43, (1976), pp. 547-549.

(3) F. Kydland and E. Prescott, "Dynamic Optimal Taxation, Rational Expectations, and Optimal Control," Journal of Economic Dynamics and Control 2, (1980), pp. 79-91.

(4) G. Calvo, "On the Time Consistency of Optimal Policy in a Monetary Economy," Econometrica 46, (1978), pp. 1411-1428.

(5) S. Fischer, "Dynamic Inconsistency, Cooperation and the Benevolent Dissembling Government," Journal of Economic Dynamics and Control 2, (1980), pp. 93-107.

(6) R. Bellman, Dynamic Programming, Princeton University Press, 1957.

(7) L. Mitten, "Composition Principles for Synthesis of Optimal Multistage Process," Operations Research 12, (1964), pp. 610-619.

(8) G. Neuhauser, Introduction to Dynamic Programming, Wiley, 1966.

(9) T. Sargent, Macroeconomic Theory, Academic Press, 1979.

(10) T. Koopmans, "Stationary Ordinal Utility and Impatience," Econometrica 28, (1960), pp. 287-309.

(11) R. Lucas and T. Sargent, Rational Expectations and Econometric Practice, University of Minnesota Press, 1981.