Review of Economic Studies (1986), LIII, 877-882 0034-6527/86/00530877$02.00

On the Existence of Markov-Consistent Plans under Production Uncertainty

B. DOUGLAS BERNHEIM
Stanford University and National Bureau of Economic Research

and

DEBRAJ RAY
Stanford University

Strotz (1956) and Pollak (1968) were among the first to study the behaviour of an economic agent whose preferences change over time. They suggested that such an agent would choose a "consistent plan", which they described as "the best plan that he would actually follow". A Markov-consistent plan has a particularly simple structure: current decisions are independent of past decisions, except insofar as past decisions affect the current values of state variables. Unfortunately, Markov-consistent plans do not generally exist. In this paper, we demonstrate that the existence problem dissappears for finite horizon problems when one introduces even a small amount of smooth uncertainty into production.

1. INTRODUCTION

Strotz (1956) and Pollak (1968) were among the first to study the behaviour of an economic agent whose preferences change over time. They suggested that such an agent would choose a "consistent plan", which they describe as "the best plan that he would actually follow". In related work (see Phelps and Pollak (1968), as well as Peleg and Yaari (1973)) it has been observed that this situation is formally equivalent to one in which decisions are made by a sequence of heterogeneous planners. Indeed, Goldman (1980) notes that the notion of a consistent plan corresponds to that of a subgame perfect equilibrium (Selten (1965, 1975)) in the game played by such a sequence of planners.

The existence of consistent plans has proved to be difficult to establish, even in relatively simple environments. Most of the literature in this area concerns simple aggregative planning models where, in each period t, the t-th planner (or the t-th incarnation of the single planner) chooses to divide a given endowment between consumption and investment. The return to this investment forms the endowment of the next planner. For this class of models, Peleg and Yaari (1973) investigated the existence of what we shall refer to as "Markov-consistent plans"[1] (these correspond to the notion of Markov-perfect equilibria). A Markov-consistent plan is a consistent plan in which the planner's current choice depends upon the history of choices only through his current endowment. Unfortunately, Markov-perfect equilibria do not generally exist, as Peleg and Yaari demonstrated through the analysis of an example.

In particular, they consider a sequence of four planners, indexed $t = 0, \ldots, 3$. The problem is one of "cake-eating"; i.e. the consumptions (c_i, $i = 0, \ldots, 3$) of the four planners must exhaust an exogeneously given initial stock. The planners' preferences are as follows:

$$u_3(c_3) = c_3$$

$$u_2(c_2, c_3) = \min\left(2c_2, \frac{c_2+3}{2}\right) + c_3$$

$$u_1(c_1, c_2, c_3) = \min\left(2c_1, \frac{c_1+3}{2}\right) + c_3$$

$$u_0(c_0, c_1, c_2, c_3) = (c_0 c_1 c_2)^{1/2} + c_1.$$

For this model, no Markov-consistent plan exists. Furthermore, while the utility functions described above are not strictly concave, this is not the source of non-existence: the problem is more subtle.[2]

Despite this negative result, subsequent studies have made significant progress on the existence question. One approach proceeds by restricting attention to certain subclasses of consistent planning problems. Berheim and Ray (1983), as well as Leininger (1983), consider models in which the preferences of the t-th planner depend only upon consumption in period t (which he himself chooses), and consumption in period $t+1$ (which is chosen by his successor). These simplifications provide the needed structure to assure the existence of Markov-consistent plans.

A second approach looks at more general policy functions. In particular, one might allow arbitrarily complex history-dependence, rather than restricting current choices to depend only on current endowments (i.e. establish the existence of consistent plans, rather than Markov-consistent plans). In fact, Goldman (1980) has provided a very general existence result for the finite horizon case (it subsumes the environment for Peleg and Yaari's counterexample). Recently, Harris (1985) has adopted Goldman's techniques to demonstrate existence for infinite horizon problems.

Despite these advances, the lack of a sufficiently general theorem on the existence of *Markov*-consistent plans represents a serious lacuna in the literature. Presumably, economists study such models not only to establish existence theorems, but in order to draw concrete conclusions about behaviour. While the existence question is certainly important, it is only a preliminary to further study. Unfortunately, it is very difficult to investigate the properties of consistent plans when one allows for complex history dependence of policy functions. In this respect, it is extremely helpful to focus attention on Markov-consistent plans, for two reasons. First, the requirement of Markov-consistency simplifies the structure of equilibria. Simple equilibria may be more likely to arise in practice, and the properties of such equilibria are certainly more amenable to study (see, for example, Bernheim and Ray (1985a)). Second, Markov-consistent plans will undoubtedly turn out to be very useful in studying the properties of consistent plans with more complex history dependence. Consider for a moment the literature on subgame perfect equilibria in repeated games. Since such games lack a "state" variable, the set of Markov-perfect equilibria coincides exactly with the set of perfect equilibria formed by taking a sequence of solutions to the static game. The standard approach is to enforce some profile of choices by retreating to these Markov-perfect equilibria to punish deviations (we note that better punishments are often available—see Abreu (1982)). Similarly, one might enforce choices in general dynamic games by retreating to Markov-perfect equilibria.[3] Thus, the study of Markov-consistent plans may well help us to understand the structure of consistent plans in general.

The primary difficulty in establishing the existence of Markov-consistent plans is that the objective function of a particular planner need not be continuous, even given reasonably well behaved utility functions, and strategies for his successors. The point of this paper is simple: a small amount of uncertainty smooths these discontinuities, thereby eliminating the primary obstacle to existence. The introduction of uncertainty has proved useful for establishing existence in other contexts (see, for example, Anderson and Sonnenschein (1982)). However, we must emphasize that this procedure is particularly natural for the consistent planning problem, since uncertainty can be introduced through the production technology, and is likely to exist there.

We establish the existence of Markov-consistent plans for finite horizon problems under extremely general conditions. It is worth emphasizing that the proof of our main result is both short and conceptually simple. A comparison with Goldman's technique clearly suggests that the introduction of production uncertainty vastly simplifies the problem of consistent planning.

Finally, we note two points. First, the approach taken here permits a simple existence proof of consistent plans in non-Markov environments. Second, it is possible to obtain a similar result for infinite horizon economies under two additional assumptions—there is a single, aggregate good, and utility is separable in one's own consumption, and the consumption of one's descendents. Since these assumptions are significantly more restrictive than those employed here, and since the proof of our infinite horizon result nevertheless involves considerably more technical machinery, we present this partial extension elsewhere (Bernheim and Ray (1985*b*)).

2. THE MODEL

Consider a sequence of $T+1$ generations, labelled $0, \ldots, T$. Generation t receives an endowment vector, $y_t \in \mathbb{R}^n$, and chooses an investment, $x_t \in \mathbb{R}^n$. Consumption is determined as a residual. Note that our framework allows for disaggregated commodity bundles. The well-being of generation t depends upon the sequence of choices. Specifically, we assume

Assumption 1. The utility of generation t is given by some continuous function, $u_t : \mathbb{R}^{2n(T+1-t)} \to \mathbb{R}$, which we write as follows:

$$u_t(y_t, x_t; \ldots ; y_T, x_T).$$

Remark 1. Implicitly, we assume that each generation's well being is independent of its ancestors' choices. Trivially, this assumption could be weakened to require separability between ancestors' choices and descendents' choices, where "descendents" is understood to include generation t itself. Further weakening of this assumption is clearly impossible: if ancestors' choices affect t's ordinal preferences over descendents' choices, the use of Markov policy functions will, in general, be suboptimal.

Remark 2. Note that we write utility as a function of endowments and investments. This subsumes the case where utility depends only on the stream of consumptions, but is not limited to this case. In particular, generations may also directly enjoy the possession of wealth, or the act of making transfers.

Remark 3. Note that utility is defined over possibly negative arguments. Conceptually, this is unimportant, since we rule out negative realizations below. However, it is convenient for the purposes of our proof.

The investment vector chosen by each generation determines the endowment of its successor up to a random disturbance, ω_t, which is realized independently of prior disturbances from some state space, Ω_t.[4] Specifically,

$$y_{t+1} = G_t(x_t, \omega_t). \tag{1}$$

It is always possible to decompose G_t as follows:

$$G_t(x_t, \omega_t) = g_t(x_t) + z_t, \tag{2}$$

where $z_t \in \mathbb{R}^n$ is random, and where the distribution of z_t depends upon x_t. We will find it convenient to use this additive formulation, and to make assumptions directly on the distribution of z_t. These correspond to more primitive assumptions on G_t and the distribution of ω_t, but we will not pursue these implications here.

Specifically, we assume

Assumption 2. g_t is a continuous function for all t.

Assumption 3. For each x_t, the probability distribution of z_t is atomless, and can be described by a continuous density function $f_t(\cdot, x_t)$.

Assumption 4. If $z_t \in \operatorname{supp}[f_t(\cdot, x_t)]$, then $g(x_t) + z_t \geqq 0$.[5]

Assumptions 2 and 3 are self-explanatory. Assumption 4 implies that realized endowments are always non-negative.

Generation 0 is exogenously endowed with y_0, taken to lie in some compact set. We will say that the sequence $\langle y_t, x_t, z_t\rangle_{t=0}^T$ is *feasible* if, for all t,

$$0 \leqq x_t \leqq y_t \tag{3}$$

$$y_{t+1} = g_t(x_t) + z_t \tag{4}$$

$$z_t \in \operatorname{supp}[f(\cdot, x_t)]. \tag{5}$$

If x_t(resp. y_t, z_t) appears as an element in some feasible sequence, we will say that x_t(resp. y_t, z_t) is *feasible.* Denote by X_t(resp. Y_t) the set of feasible x_t(resp. y_t).

Assumption 5. $\bigcup_{x_t \in X_t} \operatorname{supp}[f_t(\cdot, x_t)]$ is bounded

Assumption 5 implies that the set of feasible z_t is bounded.

A (Markov) *strategy* for generation t is a function $K_t: Y_t \to \mathbb{R}^n_+$ such that for all $y_t \in Y_t$, $0 \leqq K_t(y_t) \leqq y_t$. A sequence of strategies $\langle K_t^*\rangle_0^T$ is a *Markov-perfect equilibrium* (Markov-consistent plan) if, for each t and $y_t \in Y_t$, $K_t^*(y_t)$ solves[6]

$$\begin{aligned} &\max_{x_t} E_{z^t} u_t(y_t, x_t; \dots; y_T, x_T) \quad \text{where } y_\tau = g_{\tau-1}(x_{\tau-1}) + z_{\tau-1}, \ \tau = t+1, \dots, T, \\ &\qquad x_\tau = K_\tau^*(y_\tau), \qquad \tau = t+1, \dots, T, \end{aligned} \tag{6}$$

and $z^t \equiv (z_t, \dots, z_T)$.

3. EXISTENCE

We provide a preliminary result, from which our existence theorem follows immediately.

Theorem. *Suppose* $(K_1, \dots, K_T)$ *are measurable functions used as strategies by generations* 1 *through T. Then, under Assumptions* 1 *through* 5, *an optimal strategy for generation* 0 *exists. Moreover, this strategy can be chosen to be measurable.*

Remark. Throughout, we have in mind the notion of measurability induced by the Borel σ-algebra.

Proof. Define the functions $\Phi^T, \ldots, \Phi^0$ recursively as follows.

$$\begin{aligned}\Phi^T(y_0, x_0; \ldots; y_T, x_T) &\equiv u_0(y_0 x_0; \ldots; y_T, x_T)\\ \Phi^t(y_0, x_0, \ldots, y_t, x_t) &\equiv E_{z_t}\Phi^{t+1}(y_0, x_0, \ldots, y_t, x_t, g_t(x_t)+z_t, K_{t+1}(g_t(x_t)+z_t)).\end{aligned} \tag{7}$$

Notice that generation 0's problem can be written as: for each y_0, choose x_0 to maximize $\Phi^0(y_0, x_0)$. We will argue that Φ^0 is well-defined and continuous. This is established by induction. First, Φ^T is continuous. Now suppose that Φ^{t+1} is continuous; we will show that Φ^t is well-defined and continuous. First, notice that for all $h^t \equiv (y_0, x_0, \ldots, y_t, x_t)$, we may write

$$\Phi^t(h_t) = \int_{\mathbf{R}^n} \Psi^{t+1}(h_t, z_t)\,d\mu(z_t) \tag{8}$$

where

$$\Psi^{t+1}(h_t, z_t) \equiv \Phi^{t+1}(y_0, x_0; \ldots; y_t, x_t; g_t(x_t)+z_t, K_{t+1}(g_t(x_t)+z_t))f_t(z_t, x_t) \tag{9}$$

and $\mu(\cdot)$ is n-dimensional Lebesgue measure.

Now suppose that we have some sequence $h_t^m \to h_t$. Write

$$\begin{aligned}\Phi^t(h_t^m) &= \int_{\mathbf{R}^n} \Psi^{t+1}(h_t^m, \xi_t)\,d\mu(\xi_t)\\ &= \int_{\mathbf{R}^n} \Phi^{t+1}(y_0^m, x_0^m, \ldots, y_t^m, x_t^m, g_t(x_t)+z_t, K_{t+1}(g_t(x_t)+z_t))f_t(z_t+g_t(x_t)\\ &\quad - g_t(x_t^m), x_t^m)\,d\mu(z_t)\end{aligned} \tag{10}$$

(the last equality follows from a change of variables, taking $\xi_t = z_t + g_t(x_t) - g_t(x_t^m)$).

Now we use Lebesgue's dominated convergence theorem, as follows. Let

$$\gamma(z_t) = \Psi^{t+1}(h_t, z_t)$$

$$\gamma^m(z_t) = \Phi^{t+1}(y_0^m, x_0^m, \ldots, y_t^m, x_t^m, g_t(x_t)+z_t, K_{t+1}(g_t(x_t)+z_t))f_t(z_t+g_t(x_t)-g_t(x_t^m), x_t^m).$$

By continuity of Φ^{t+1}, f_t, and g_t, $\gamma^m(z_t) \to \gamma(z_t)$ for all z_t. Since the distribution of z_t has compact support, $\gamma^m(z_t)$ is zero outside of a compact set (see Assumption 6). Further, it is bounded, since both Φ^{t+1} and f_t are bounded (u_t and f_t are continuous functions, and the feasible set is bounded). Thus, the conditions of Lebesgue's dominated convergence theorem are satisfied, and

$$\Phi^t(h_t^m) = \int_{\mathbf{R}^n} \gamma^m(z_t)\,dz_t \to \int_{\mathbf{R}^n} \gamma(z_t)\,dz_t = \Phi^t(h_t). \tag{12}$$

So Φ^t is continuous, which completes the induction step. We know that Φ^0 is continuous in y_0 and x_0. Let

$$\boldsymbol{K}_0(y_0) = \{\arg\max_{0 \le x_0 \le y_0} \Phi^0(y_0, x_0)\}. \tag{13}$$

By the maximum theorem, $\boldsymbol{K}_0(\cdot)$, as y_0 varies over its domain, is an upperhemicontinuous correspondence. Since its graph is closed, it is measurable. By von-Neumann's measurable selection theorem, there exists some measurable function K_0, with $K_0(y_0) \in \boldsymbol{K}_0(y_0)$ for all y_0. ||

To establish the existence of Markov-consistent plans, we employ an obvious induction step on the number of generations. Thus, we have our main result:

Proposition. *Under Assumptions* 1 *through* 5, *any finite period planning game has a Markov-perfect equilibrium.*

First version received May 1985; *final version accepted May* 1986 (*eds*)

This research was supported by National Science Foundation Grant SES-84-04164 at the Institute for Mathematical Studies in the Social Sciences, Stanford University, Stanford, California.

NOTES

1. Actually, they used the term "dynamic programming solution", a bit inappropriate in a model of changing planners.
2. In brief, the problem arises becauses a "well-behaved" environment does *not* guarantee a well-defined maximization problem for each agent. The feasible set for an agent depends *also* on the policies to be followed by other agents. Structure on such policies cannot be imposed exogenously.
3. Following the insight of Abreu (1982), it might be argued that Markov equilibria may not be the best "punishment equilibria" available, so why study them? In reply, we note that Markov equilibria will certainly support *some* collusive history-dependent equilibria (for appropriate discount factors), so the latter may be analyzed. Our choice of Markov equilibria is on grounds of their innate simplicity (see Bernheim and Ray (1985*a*)).
4. Unless the ω_t's are realized independently, the environment will not be a Markov environment. One can easily adapt our proof to the case where the ω_t's are correlated, but then generation t's equilibrium strategy will, in general, depend upon the enire history of random disturbances.
5. supp $[f(\cdot)]$ denotes the support of the density function f.
6. For any random variable q, E_q denotes expectation taken over q.

REFERENCES

ABREU, D. (1982), "Repeated Games with Discounting: A General Theory and an Application to Oligopoly" (mimeo, Department of Economics, Princeton University).

ANDERSON, R. and SONNENSCHEIN, H. (1982), "On the Existence of Rational Expectations Equilibrium", *Journal of Economic Theory*, **26**, 261-279.

BERNHEIM, B. D. and RAY, D. (1983), "Altruistic Growth Economies: I. Existence of Bequest Equilibria" (Technical Report No. 429, Institute for Mathematical Studies in the Social Sciences, Stanford University).

BERNHEIM, B. D. and RAY, D. (1985*a*), "Altruistic Growth Economies: II. Properties of Bequest Equilibria" (mimeo, Department of Economics, Stanford University).

BERNHEIM, B. D. and RAY, D. (1985*b*), "Markov Perfect Equilibrium in Altruistic Growth Economies" (mimeo, Department of Economics, Stanford University).

GOLDMAN, S. (1980), "Consistent Plans", *Review of Economic Studies*, **48**, 533-537.

HARRIS, C. (1985), "Existence and Characterization of Perfect Equilibrium in Games of Perfect Information", *Econometrica*, **53**, 613-628.

PHELPS, E. and POLLAK, R. (1968), "On Second-Best National Saving and Game-Equilibrium Growth", *Review of Economic Studies* **35**, 185-199.

POLLAK, R. (1968), "Consistent Planning," *Review of Economic Studies*, **35**, 201-208.

SELTEN, R. (1965), "Spieltheoretische Behandlungeines Oligopolmodells mit Nachfragetragheit", *Zeitschrift fur die Gesamte Straatiswissenscahft*, **121**, 301-324.

SELTEN, R. (1975), "Reexamination of the Perfectness Concept for Equilibrium Point in Extensive Games", *International Journal of Game Theory*, **4**, 25-55.

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